

STRUCTURE OF A PAPER

RENDERING RESEARCH PAPER STRUCTURE

1. Teaser
2. Abstract
3. Introduction
4. Related work
5. Algorithm or System
6. Evaluation
7. Conclusions & Discussion
8. Bibliography
9. Appendices



TEASER

Sample-based Monte Carlo Denoising using a Kernel-Splatting Network

MICHAËL GHARBI, Adobe and MIT CSAIL

TZU-MAO LI, MIT CSAIL

MIIKA AITTALA, MIT CSAIL

JAAKKO LEHTINEN, Aalto University and NVIDIA

FRÉDO DURAND, MIT CSAIL

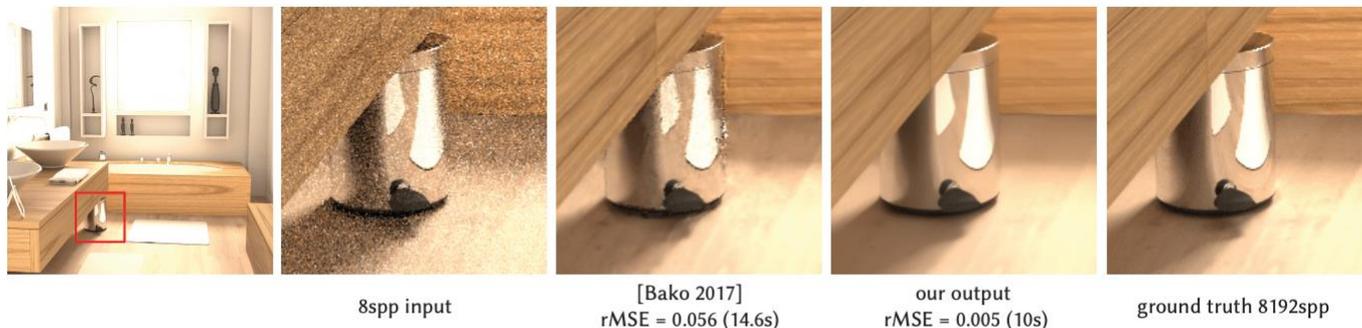


Fig. 1. State-of-the-art pixel-based Monte Carlo denoising algorithms (right) struggle with very noisy inputs rendered with a low sample count (left). Our method (middle) works with the *samples* directly, it uses a *splatting* approach, and is trained using deep learning. This makes it possible to appropriately handle various components of the illumination (indirect lighting, specular reflection, motion blur, depth of field, etc) more effectively.

TITLE

Mechanosensation of cyclical force by PIEZO1 is essential for innate immunity

Angel G. Solis, Piotr Bielecki, Holly R. Steach, Lokesh Sharma, Christian C. D. Harman, Sanguk Yun, Marcel R. de Zoete, James N. Warnock, S. D. Filip To, Autumn G. York, Matthias Mack, Martin A. Schwartz, Charles. S. Dela Cruz, Noah W. Palm, Ruaidhrí Jackson & Richard A. Flavell
Nature volume 573, pages 69-74 (2019)

<https://www.nature.com/articles/s41586-019-1485-8>

TITLE & ABSTRACT

Wide BVH Traversal with a Short Stack

K. Vaidyanathan S. Woop C. Benthin

Intel Corporation

Abstract

Compressed wide bounding volume hierarchies can significantly improve the performance of incoherent ray traversal, through a smaller working set of inner nodes and therefore a higher cache hit rate. While inner nodes in the hierarchy can be compressed, the size of the working set for a full traversal stack remains a significant overhead. In this paper we introduce an algorithm for wide bounding volume hierarchy (BVH) traversal that uses a short stack of just a few entries. This stack can be fully stored in scarce on-chip memory, which is especially important for GPUs and dedicated ray tracing hardware implementations. Our approach in particular generalizes the restart trail algorithm for binary BVHs to BVHs of arbitrary widths. Applying our algorithm to wide BVHs, we demonstrate that the number of traversal steps with just five stack entries is close to that of a full traversal stack. We also propose an extension to efficiently cull leaf nodes when a closer intersection has been found, which reduces ray primitive intersections by up to 14%.

TITLE & ABSTRACT

Wide BVH Traversal with a Short Stack

K. Vaidyanathan S. Woop C. Benthin
Intel Corporation

The problem/challenge/goal

Prior limitation

Key idea

Abstract

Compressed wide bounding volume hierarchies can significantly improve the performance of incoherent ray traversal, through a smaller working set of inner nodes and therefore a higher cache hit rate. While inner nodes in the hierarchy can be compressed, **the size of the working set** for a full traversal stack remains a significant overhead. **In this paper we introduce an algorithm for wide bounding volume hierarchy (BVH) traversal that uses a short stack** of just a few entries. This stack can be fully stored in scarce on-chip memory, which is especially important for GPUs and dedicated ray tracing hardware implementations. Our approach in particular **generalizes the restart trail algorithm for binary BVHs to BVHs of arbitrary widths**. Applying our algorithm to wide BVHs, we demonstrate that the number of traversal steps with **just five stack entries** is close to that of a full traversal stack. We also propose an extension to efficiently cull leaf nodes when a closer intersection has been found, which **reduces ray primitive intersections by up to 14%**.

Deliverable

How well it works

INTRODUCTION

- Very readable, high-level
- Background & problem statement
- Seminal related work citations
- Description of constraints
- The paper's approach to the problem
- **Contributions of this paper**

This is what you need for your 1st pass

1 INTRODUCTION

Monte Carlo (MC) integration is an essential tool in light transport simulation [Pharr et al. 2016; Veach 1997] and other fields of science and engineering [Kalos and Whitlock 2008]. An inherent problem of MC integration is its slow convergence, which is why numerous variance reduction schemes have been proposed, notably importance sampling. Its extension, known as multiple importance sampling (MIS) [Veach and Guibas 1995], is particularly versatile as it enables combining different sampling techniques in a robust way to form better MC estimates...

Our work focuses on weighting functions for MIS. We derive...

We provide further theoretical insights into...

Our practical contribution consists in proof-of-concept applications of the optimal weighting scheme in light transport...

RELATED WORK

- Mini survey paper
- Categorize the previous work and position this paper
- Explain differences between this and previous papers
- More detail on what is completely new

2 RELATED WORK

Monte Carlo Methods. Kajiya and Von Herzen [1984] were the first to use path tracing for numerically estimating radiative transfer in volumes [Chandrasekhar 1960]..., these methods are far from reaching interactive frame rates when used on the highly scattering materials that we target...

We explore a new approach based on approximating the cloud geometry by a hierarchical descriptor and predict local illumination using a deep neural network....

...All these methods are either interactive, or produce high-fidelity images, but none of them achieve both concurrently...

Neural Networks. Deep neural networks (see Bengio et al. [2013]; LeCun et al.[2015] for a comprehensive review) are able to efficiently model complex relationships between input and output variables in a highly non-linear manner... **We use** a hierarchical feature and feed its levels into the network progressively..

[Deep scattering: rendering atmospheric clouds with radiance-predicting neural networks](#), Kallweit et al., SIGGRAPH'17

ALGORITHM OR SYSTEM

- Main body of the new technique
- Includes both derivation and solution
- May span multiple sections

- Maybe a small part of this is the new contribution you're looking for!
- ...and it may depend on pieces that *only* appear in other work, not even here

4. Sampling the Projected Area of a Hemisphere

In this section, we derive an area-preserving parameterization that we use to sample the projected area of the hemisphere.

4.1. Orthonormal Basis

We start by constructing an orthonormal basis (V_h, T_1, T_2) (see Figure 4), where T_1 is in the tangent plane orthogonal to $Z = (0, 0, 1)$:

$$T_1 = \frac{Z \times V_h}{\|Z \times V_h\|} = \frac{(-y_v, x_v, 0)}{\sqrt{x_v^2 + y_v^2}}, \quad (7)$$

$$T_2 = V_h \times T_1. \quad (8)$$

```
// Section 4.1: orthonormal basis (with special case if cross product is zero)
float lensq = Vh.x * Vh.x + Vh.y * Vh.y;
vec3 T1 = lensq > 0 ? vec3(-Vh.y, Vh.x, 0) * inversesqrt(lensq) : vec3(1, 0, 0);
vec3 T2 = cross(Vh, T1);
```

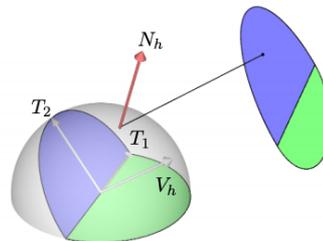


Figure 4. Orthonormal basis for sampling the projected area of the hemisphere.

4.2. Parameterization of the Projected Area

Shape of the projected area. Figure 5 shows the shape of the projected area of the hemisphere. It is the signed sum of the projected areas of the two half disks. The projected area of the half disk located in the tangent plane (in green) is proportional

EVALUATION (A.K.A. RESULTS)

- Quantitative evaluation
- Comparison to previous work and “ground truth”
- Measured performance or asymptotic analysis
- Look at images & read captions on first reading pass
- Beware of measurement differences between papers

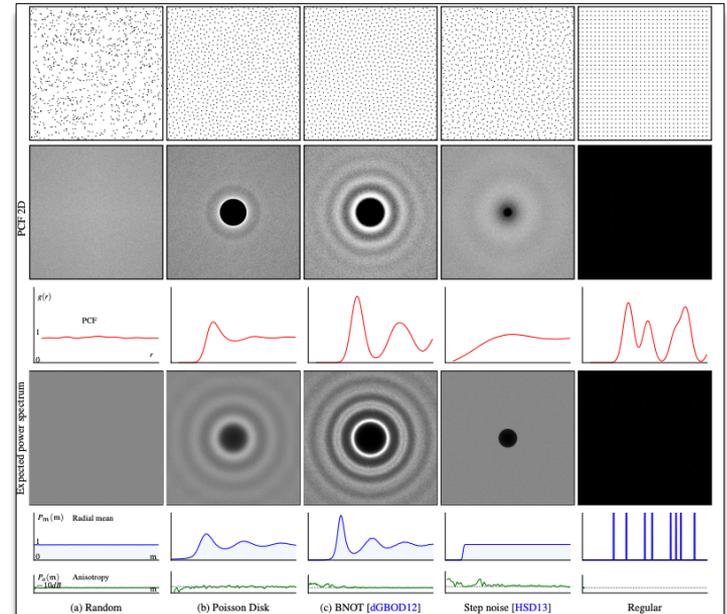


Figure 5: An example distribution generated by different underlying point processes (top row) and its corresponding 2D PCF (second row) are shown for distributions varying from purely random to a regular grid pattern. Third row shows the 1D PCF version where the r axis is normalized with the maximum possible distance between pairs of points for this square domain and the number of points. The expected power spectrum (details in Section 3.2.1) of each sampling pattern is shown in the fourth row. The corresponding radial statistics include radial mean (fifth row) and radial anisotropy (sixth row) along the radial frequency m .

and second order correlations, which are essential to analyze and develop point patterns for further error reduction.

3.1.2. Pair Correlation Function

Point process statistics provide provably informative measures for point patterns, i.e. the characteristics of a point pattern are uniquely and fully defined once these statistics are provided. As mentioned in Section 2, for our purpose of analyzing error in integral estimation, we only need first and second order product densities. We will see

why this is the case in Section 4. Historically, such spatial measures of correlations have come quite late into the rendering literature, but have turned out to be very powerful, as they contain all information on a point pattern via the underlying point process.

For general point patterns, in particular if we do not assume stationarity, the second order correlation captured by $\rho(\vec{x}, \vec{y})$ is a high dimensional function that is hard to visualize or analyze. Instead, the literature on rendering focuses on *stationary* patterns and hence the

CONCLUSIONS & DISCUSSION

- Often begins with a skippable paper summary
- Expansion of conclusions based on experimental results that first appeared in the abstract or title
- Valuable high-level, subjective or philosophical discussion of what the authors learned from this research
- Suggestions for future research (great ideas for your own work)



Burley et al., The Design and Evolution of Disney's Hyperion Renderer, ACM ToG 2018

© 2018 DISNEY

READING PROCESS

STUDY TITLE & TEASER

- What is this paper about? What do the words in the title really mean?
- How does the best-case result in the teaser compare to prior art and the “ground truth” goal?
- Is this paper likely the beginning, middle, or end of this line of research?
- Check for backward references in the ACM Digital Library... see how those authors describe this paper in *their* related work, and maybe recursively start reading those papers.
- When was it published (is there something newer I should read?)
- [Do these authors and this venue have a reputation for clarity, objectivity, full disclosure, and practicality?]
- Are there presentation slides or a video available that I can use for an easier overview?

The Rendering Equation

James T. Kajiya

SIGGRAPH'86

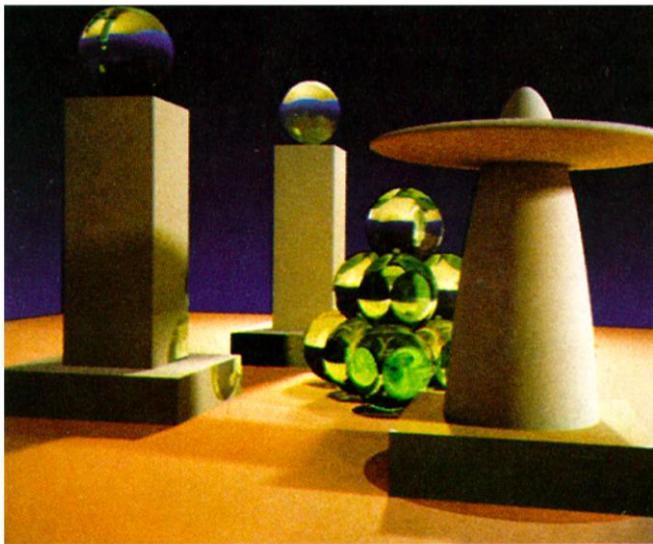


Figure 6. A sample image. All objects are neutral grey. Color on the objects is due to caustics from the green glass balls and color bleeding from the base polygon.

READ THE ABSTRACT (1/2)

What kind of paper is this?

- New problem
- New system solution
- New theory solution
- New data set
- Survey
- Position

READ THE ABSTRACT (2/2)

- **What** problem is being addressed?
- **Why** does that problem matter?
- **Why** is a new solution needed?
- What is the key idea to the new solution?
- What is the main limitation/drawback/constraint?
- How well does it work?

The Rendering Equation

Abstract.

We present an integral equation which generalizes a variety of known rendering algorithms. In the course of discussing a monte carlo solution we also present a new form of variance reduction, called Hierarchical sampling and give a number of elaborations shows that it may be an efficient new technique for a wide variety of monte carlo procedures. The resulting rendering algorithm extends the range of optical phenomena which can be effectively simulated.

The Rendering Equation

Deliverables

Key idea

Abstract.

We present an **integral equation** which **generalizes** a variety of known rendering algorithms. In the course of discussing a monte carlo solution we also present a **new form of variance reduction**, called Hierarchical sampling and give a number of elaborations shows that it *may be* an *efficient* new technique for a wide variety of monte carlo procedures. **The resulting rendering algorithm** extends the **range of optical phenomena** which can be effectively simulated.

The problem

LOOK FOR CONTRIBUTIONS IN THE INTRODUCTION

Skip to the end of the introduction, where there is usually an explicit description of the contributions.

Skim for any term of art definitions in italics or bold—those must be important concepts.

(Introduction)

The technique we present subsumes a wide variety of rendering algorithms and provides a **unified context** for viewing them as more or less accurate approximations to the solution of a **single equation**. That

...

points. Equation (1) differs from the radiosity equation of course because, unlike the latter, **no assumptions are made about reflectance characteristics of the surfaces involved.**

This diagram also points out an alternative algorithm for conventional distributed ray tracing. Rather than shooting a branching tree, just shoot a path with the rays chosen probabilistically. For scenes with much reflection and refraction, this cuts down vastly on the number of ray object intersections to be computed for a given pixel and performs a remarkable speed up of ray tracing for very little programming work. However, for this new fast form of ray tracing—called *path tracing*—we have found that it is very important to maintain the correct proportion of reflection, refraction, and shadow ray types contributing to each pixel. Rather than choosing the ray type randomly, there are two alternatives. First, keep track of the number of each type shot. Make sure the sample distribution of ray types closely matches the desired distribution by varying the probability of each type so that it is more certain that the sample distribution matches. This is the approach we have actually implemented. A second approach is to let the ray types be chosen randomly but to scale the contribution of each ray type by the ratio of desired distribution to the resulting weighted sample distribution.

SKIM CONCLUSIONS

- After all of this work, what insight did the authors leave with?
- **Is this important enough to study the rest of the paper in detail?**
- Are there limitations disclosed here?

As an approximation to Maxwell's equation for electromagnetic eq. (1) does not attempt to model all interesting optical phenomena. It is essentially a geometrical optics approximation. We only model time averaged transport intensity, thus no account is taken of phase in this equation—ruling out any treatment of diffraction. We have also assumed that the media between surfaces is of homogeneous refractive index and does not itself participate in the scattering light. The latter two cases can be handled by a pair of generalizations of eq. (1). In the first case, simply by letting $g(x, x')$ take into account the eikonal handles media with nonhomogenous refractive index. For participating propagation media, a integro-differential equation is necessary. Extensions are again well known, see [Chandrasekar 1950], and for use in a computer graphics application [Kajiya and von Herzen 1984]. Elegant ways of viewing the eikonal equation have been available for at least a century with Hamilton-Jacobi theory [Goldstein 1950]. Treatments of participatory media and of phase and diffraction can be handled with path integral techniques. For a treatment of such generalizations concerned with various physical phenomena see [Feynman and Hibbs 1965]. Finally, no wavelength or polarization dependence is mentioned in eq. (1). Inclusion of wavelength and polarization is straightforward and to be understood.

Limitations

Future work/
extensions

Limitations

SKIM RESULTS, EMPHASIZING DATA

- How well does this perform? Absolute and scaling
- How robust is the method across input variation?
- How does the *worst* case quality differ from *best* case?
- When do failure cases occur?

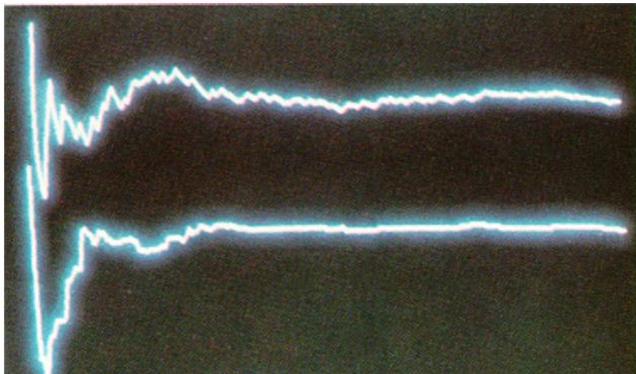


Figure 3. Convergence of naive monte carlo vs. hierarchical integration. Shown are integral estimates as a function of number of samples cast. Naive monte carlo is the top curve.



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.

SKIM *RECENT* RELATED WORK

How does this differ from the most closely-related previous work?

- Restrictions
- Performance
- Robustness
- Quality

SKIM THE ALGORITHM/SYSTEM SECTION

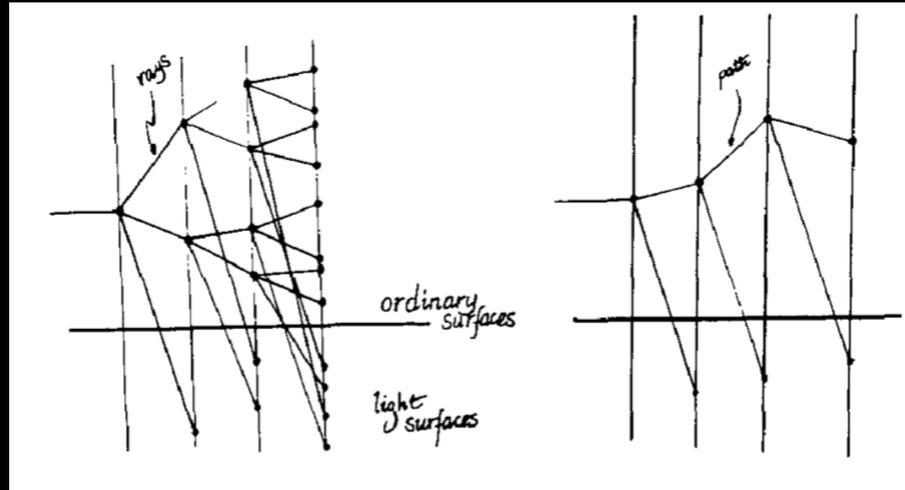
- Read the section *titles*, but skip the content
- Find the key listing or equation (usually at the end)
- Decipher the notation
- What is the magic “aha!” step?
- Look for parameters, limitations, assumptions, and dependencies

The path tracing algorithm

1. Choose a point \mathbf{z}' 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{z}'' and calculate the geometrical factor $g(\mathbf{x}, \mathbf{z}')$.
 - 3.2 Calculate the reflectance function $\rho(\mathbf{x}, \mathbf{z}', \mathbf{z}'')$ and multiply by $\epsilon(\mathbf{z}', \mathbf{z}'')$.
 - 3.3 Add this contribution to the pixel intensity.

Note that calculating the emittance and scattering factors is simply a matter of consulting texture maps and lighting models. Calculating

Distribution ray tracing (prior work): exponential time!



Path tracing:
linear time!

RE-READ FRONT MATTER

- Abstract, Introduction, Related Work
- Does my interpretation of the abstract change?
- Consider the positioning of the Introduction.
 - Does it give me insight into their approach or the problem in general?
 - Does it motivate the problem?
- For the most closely related work, follow the reference and read *their* abstracts and teasers, maybe recursively reading those full papers.

RE-READ RESULTS

- Read the Result text in full to really understand the evaluation.
- Compare to results in previous (or future) papers...is there some case that you aren't being shown?
- Failure and limitation figures are *good* signs. Be worried about papers that *don't* disclose these.

READ THE ALGORITHM SECTION

- *Now study the Algorithm/System main body.*
- Ensure that you understand *every* aspect of the notation (you may have to check other sections, other papers, books, etc.!)
 - Pay attention to hats, subscripts, superscripts, stars, etc. Beware of similar-looking symbols: W/ω , $O/0$, $L/1$, V/ν , etc.
 - What size are matrices?
 - What are the units?
 - What do functions return?
- **If you're trying to really understand this paper, then rederive the code or equations as you progress. The paper will skip steps to save space. You shouldn't.**

READ THE CONCLUSIONS. THINK.

- Think about
 - What are the key insights that are perpetual?
 - What parts are less important technical details?
 - Does this paper
 - pave the way to solve the problem in a new way?
 - or just propose one solution that could have been done anyway?
- **What makes the paper important or not-so-important?**

APPENDIX AND SUPPLEMENTALS

- Often contain more results and detailed mathematical derivations and discussions
- You do not necessarily need to read them
 - You might already understand/convinced while reading the main body of the paper
 - The authors are usually asked to make the main paper self-contained without referring to supplementals
- Refer to them when you have some doubt/questions



Wald et al., Embree: A Kernel Framework for Efficient CPU Ray Tracing, ACM ToG 2014

SUMMARY

MANY MOTIVATIONS

Many reasons you might read a paper: learn scientific result, learn structure, meta-insights, review.

Many benefits from reading primary sources.

Consider the authors' constraints and audiences when interpreting their words.

SUMMARY

1. Read in multiple passes
2. Skip around (in a specific way) to decode the paper
3. Evaluate critically as an active reader
4. *Rederive key equations/code*
5. Follow forward and backward references

HOW TO PRESENT

PRINCIPLES

- You paper = the product
- You presentation = **the advertisement**
- Focus on *inspiring interest*

PRINCIPLES

- Goal is to convey the messages/ideas you have to your target audience
- Keep it simple
- Think about your audience
- Everything you show should matter
- One message per slide
- Focus on **what** and **why**

PRINCIPLES

- At the end of your talk, audience know
 - **what** the talk was about
 - **why** is it interesting
 - **what** is novel
- Relevant audience might be interest to read the paper
 - The “**how**”.

PRINCIPLES

- Example presentation:
<https://benedikt-bitterli.me/restir/>

Recommended Presentation Structure

1. Background of the problem (what's the problem? why do we care?)
2. Previous work (keep it high-level and position the paper)
3. Contributions of the paper ← *add your own emphasis here*
4. Details of the problem
5. Key ideas (use selected equations only when needed)
6. Show results ← *critique results and/or evaluation methodology*
7. Conclude with your own understandings
8. Set up and facilitate discussion

PREFER MINIMAL SLIDES

- Avoid clip art and irrelevant images
- Simple template and font
- Good, readable colours
- Target fewer than 25 words/page
- Target one important image or equation per slide
- Don't overuse animation when it's not important
- Keep the number of equations as few as possible
- Avoid showing many numbers at once (e.g., table)
- Not too many bullet points

DEFINE AND EXPLAIN EVERYTHING

- Figures
- Equations
- Plots
- Vocabulary (MCMC, PDF, BDPT, MLT, etc.)
- Variables
- Axes
- Numbers
 - If you are not going to explain, then don't put them!

TIME CONSIDERATIONS

- “Real” presentations: know your time limits, and stay within the limits!
- In class: ~30min.

USE REFERENCES

- Cite equations and figures as needed
- Distinguish the paper from your editorial content

MYTH OF NUMBER OF SLIDES PER MINUTE

- Don't trust any tip or recommendation like this:
“You should aim for X slides for Y minutes talk”
 - Very much wide-spread, but wrong advice
 - Usually lead to “dense” slides that are difficult to understand
- Instead, check if you have one message per slide
 - “Information density” per minute is important
 - If you have animation, each slide might be very “sparse”

MYTH OF TABLE OF CONTENTS

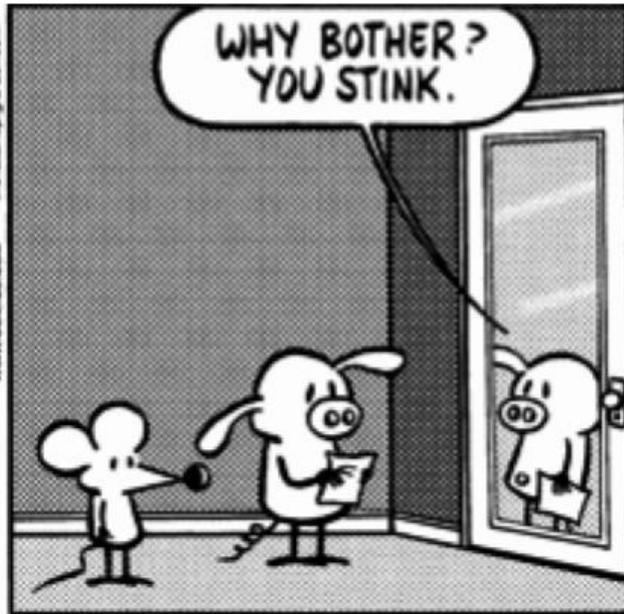
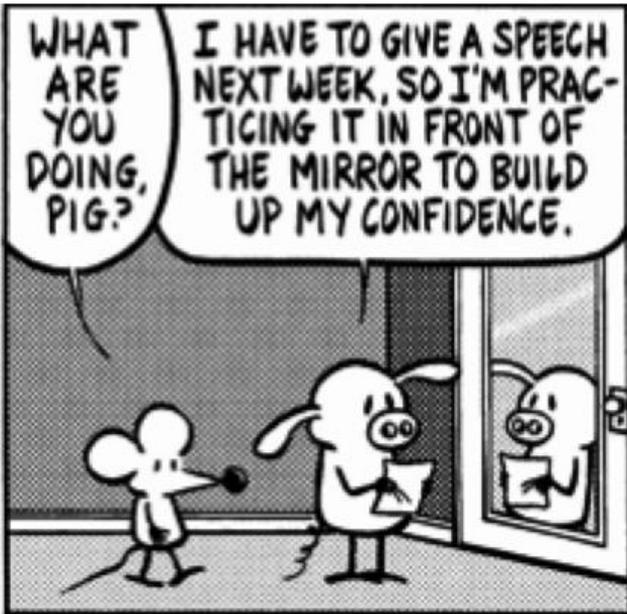
- Don't trust any tip or recommendation like this:
“Your first slide should be the table of contents”
 - Very much wide-spread, but wrong advice, again!
 - Usually irrelevant for tech talks and audience forget
- Instead, have “title only” slides time to time to remind what you are going to talk
 - “Background”, “Previous Work”, “Method”, “Results”
 - Heads-up for audience about what you gonna say

MYTH OF “THANK YOU” SLIDE

- Don't trust any tip or recommendation like this:
“Your last slide should say “Thank you for listening”.”
 - Very much wide-spread, but wrong advice
 - Same for acknowledgements – show it before the last slide
 - Of course you thank audience, but is this slide useful?
- Instead, conclude with a summary, take-home message, etc.
 - Audience will stare at the final slide for a while and think about questions, so make it useful for them

PRESENTING

- Test your slides on the computer
- Vary tone, pace, and volume



DISCUSSION

- If you are not presenting:
 - Come prepared to discuss!
 - Actively read the paper, understand the **what** and **why**.
 - Have an idea about the **how**.
 - Prepare several questions.