Form and Shape for Lighting Virtual Worlds

Graham Hazel

Form & Deformation in Art, Toys and Games, 1st December 2017

Shape influences lighting at all scales

Microstructure of surface affects appearance (= interaction with light)

Larger scale features produce shadows and reflections

Photography and cinematography use shape of light sources as a tool to sculpt image

How can we use understanding of shape to render virtual images?

Scientific progress?

Phase 1 – Abstraction

- Simplify problem to a handful of essential variables

Phase 2 – Statistics

- Simplify problem as a large collection of simple processes

Phase 3 – Complex system

- True problem is a large collection of complex processes, with feedback loops and emergent properties

Quantitative Theory of Light

Phase 1 – Publication of *Photometria* by Lambert (1760)

First coherent set of mathematical definitions of photometric quantities (eg “albedo”)

Proved formula for idealised area light (form factor for radiative transfer)

Completely ignored for ~100 years…

Computer graphics

Phase 1 – Phong shading model (1975)

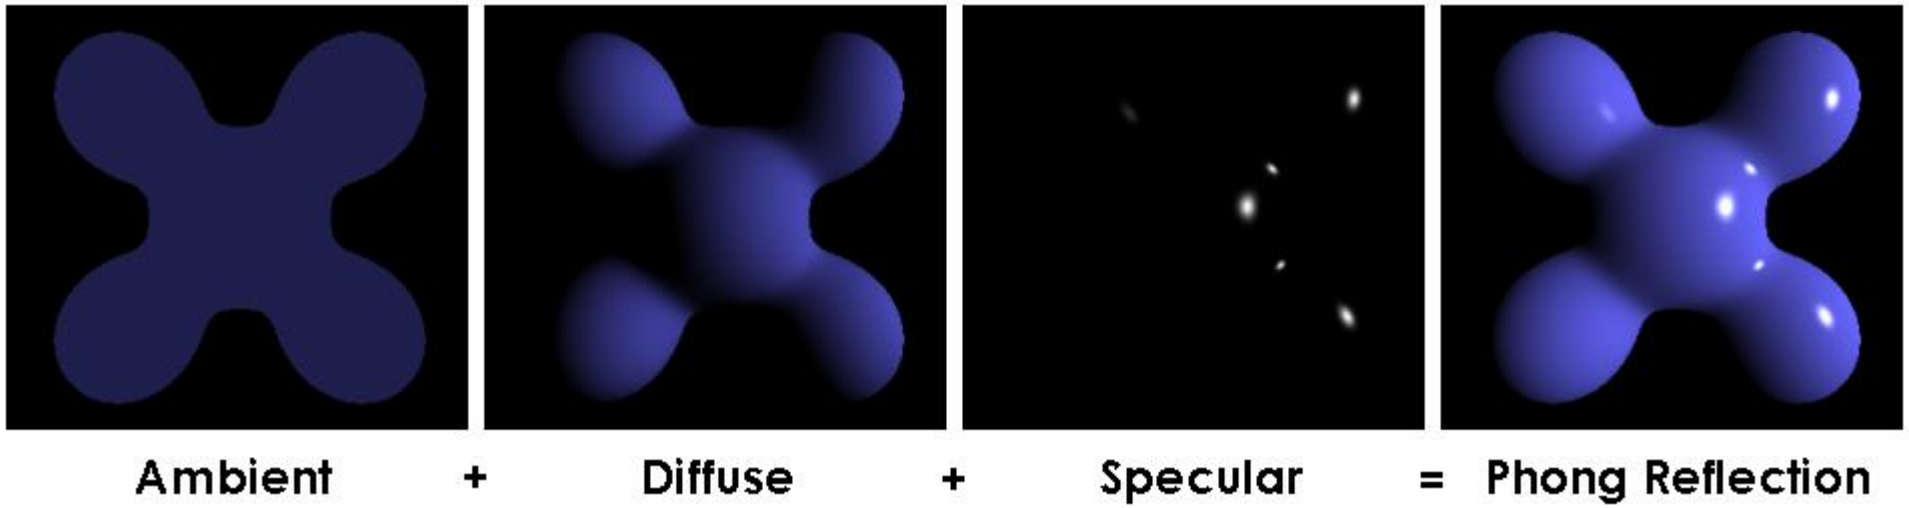

Parameters are *ambient colour*, *diffuse colour*, *specular colour* and *specular power*

Phong model is particularly good for plastic…

All lighting effects hand-authored

In *Toy Story*, all illumination is the direct effect of an authored light source

Effect of light bouncing around the scene must be “faked” by hand

There’s a red point light inside every tomato in *Ratatouille* (2007)

Very labour intensive, doesn’t scale well

Statistics to the rescue!

Physically Based Rendering (2004)

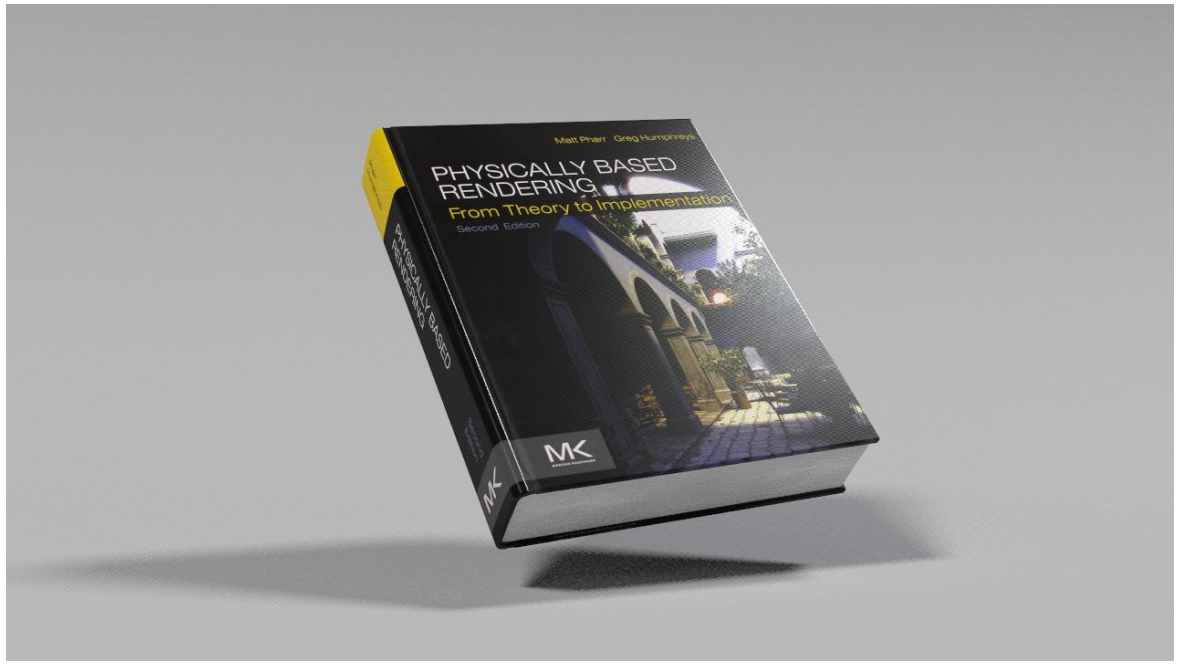

This book has deservedly won an Academy Award. I believe it should also be nominated for a Pulitzer Prize — Donald Knuth

Combine photometry with statistics

Accurate measurements of

- Reflectivity (“everything is shiny”)
- Diffusion (subsurface scattering)
- Transparency and translucency
- Fresnel reflectivity (“everything has Fresnel”)
- Metallicity

Statistical treatment of microsurface structure

Ensure all models conserve energy

What is “Metallicity”?

Conductors (metals) absorb transmitted light and reflect more light directly

In insulators (dielectrics), transmitted light bounces around and some is diffusely reflected

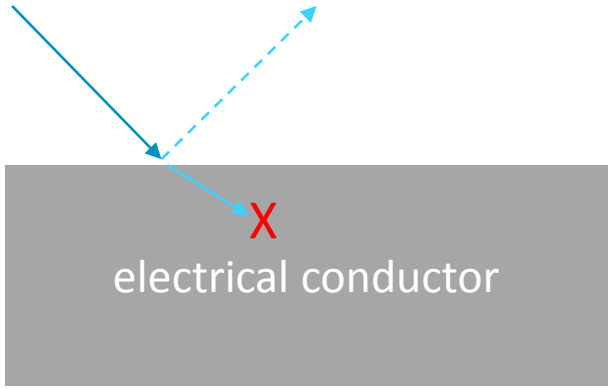

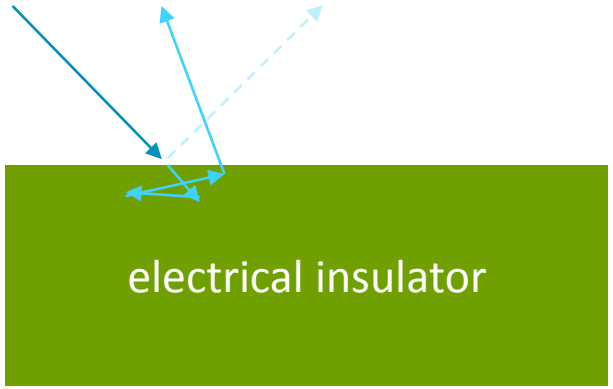

Metallicity is more-or-less a yes/no question: is this a metal?

(Q: What about semiconductors? A: Don’t render semiconductors…)

Rough surfaces reflect light in more directions

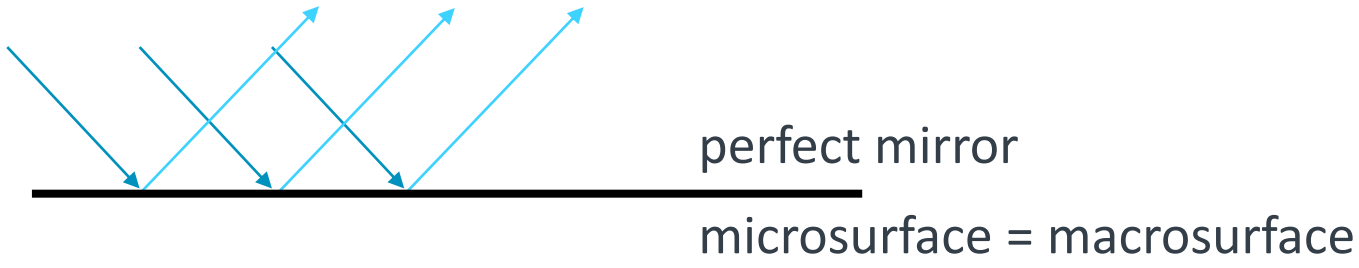

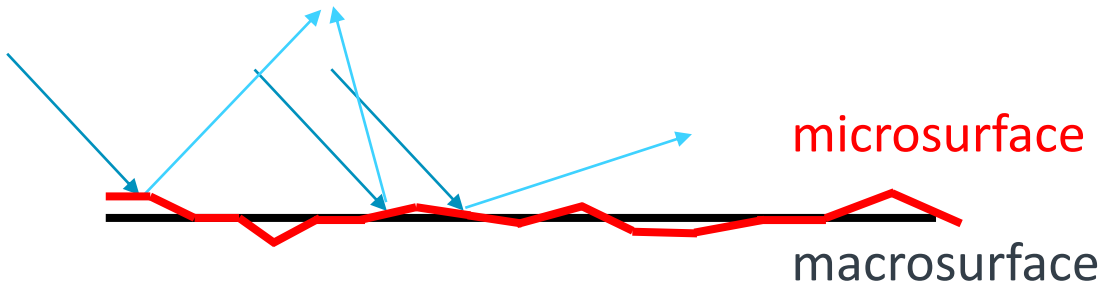

Strategy:

- Statistically model microsurface as flat *microfacets*
- Normals of microfacets vary
- Distribution controlled by *roughness* parameter

Effect of roughness

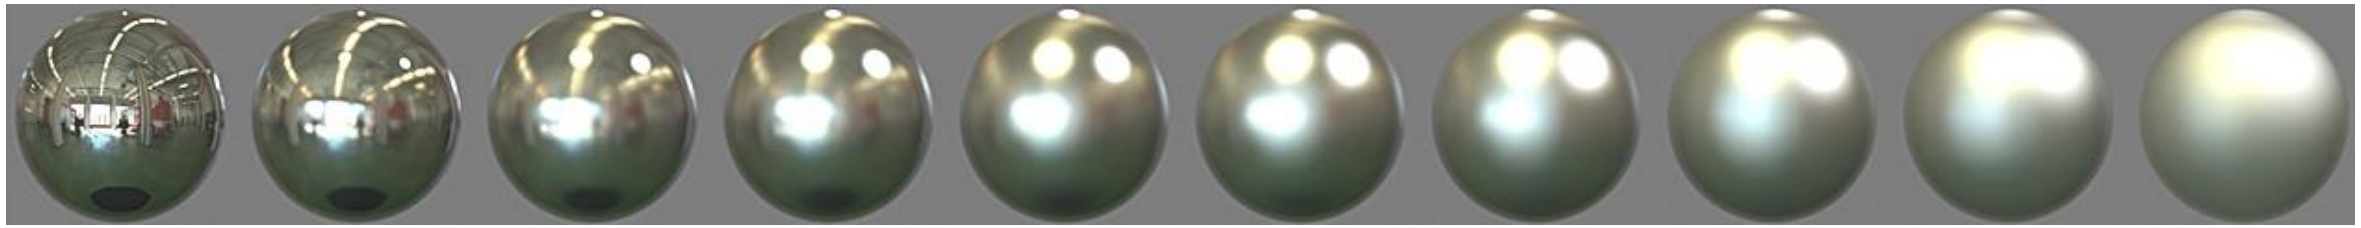

smooth rough

Wide variety of realistic materials

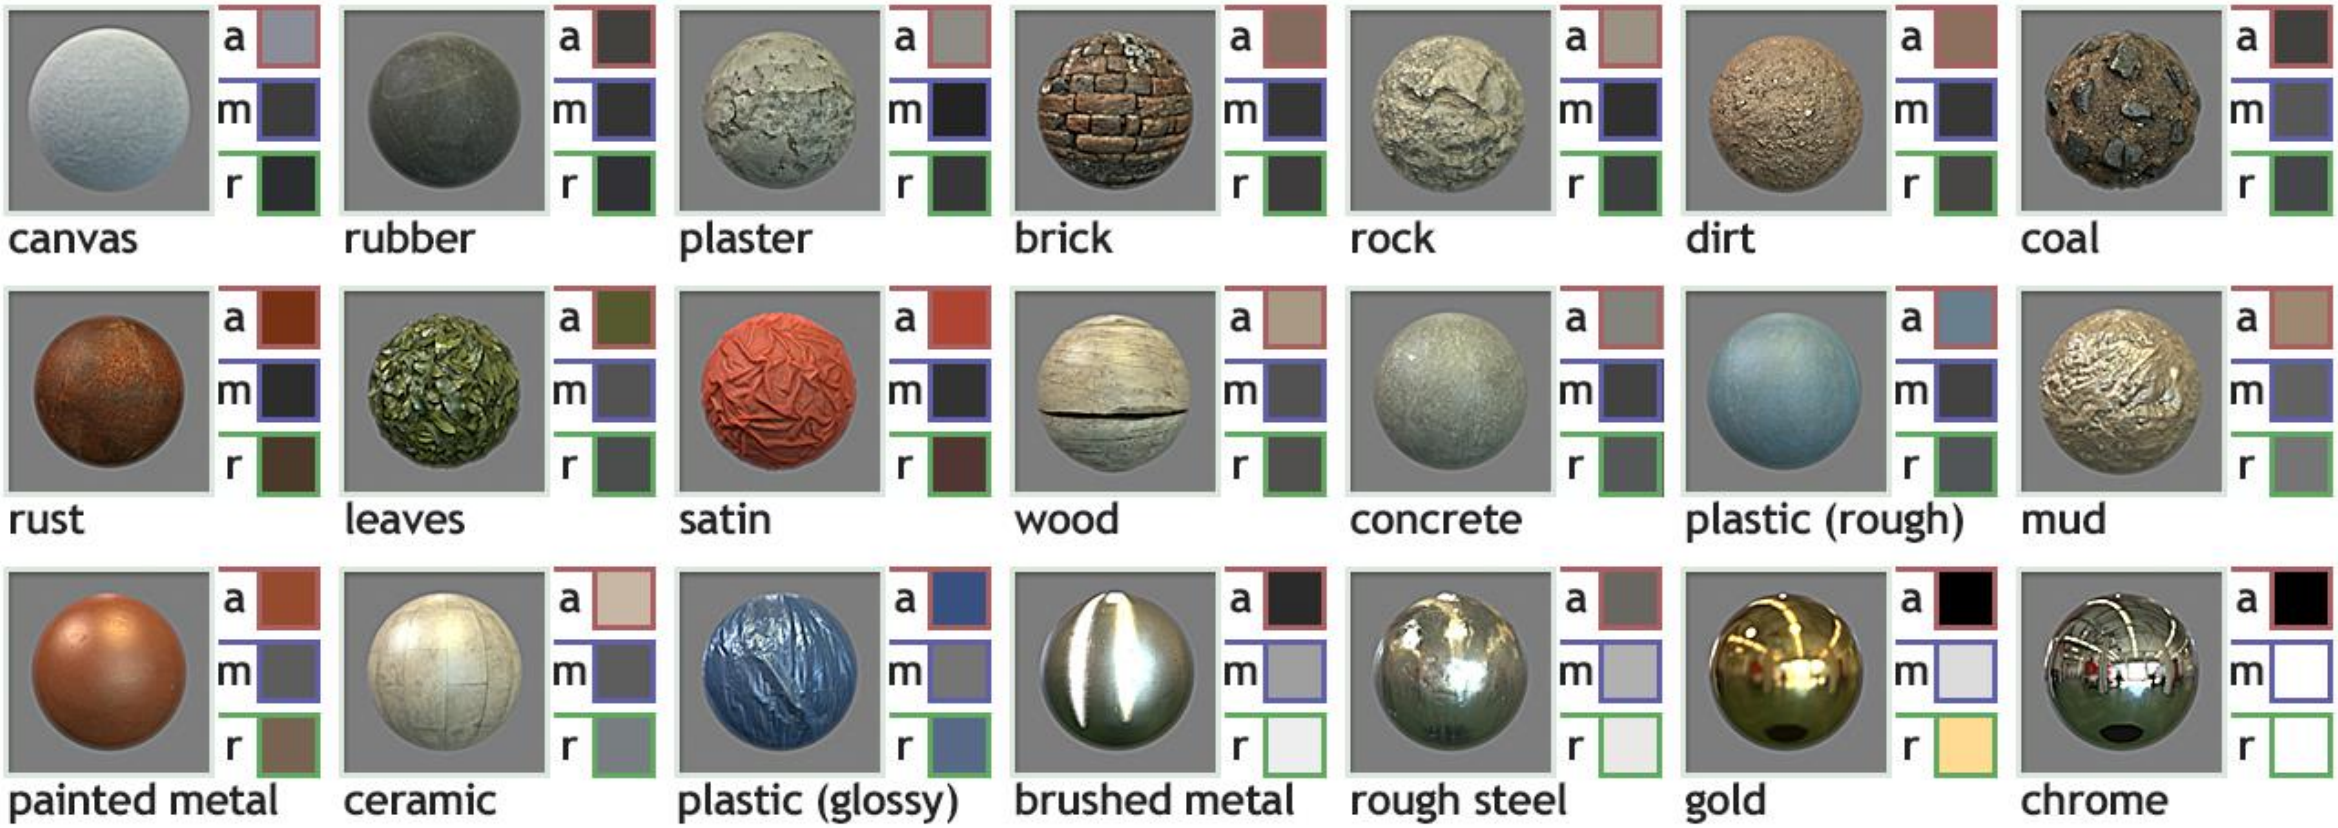

Powerful model of light<->surface interaction

Complete separation of surface and material authoring from lighting

- Same objects can be reused in different lighting conditions and “it just works”

Hides statistical complexity in the mathematical models

Exposes intuitive parameters to artists

…restricted to physically plausible values!

Normal maps – “milli-scale” surface shape

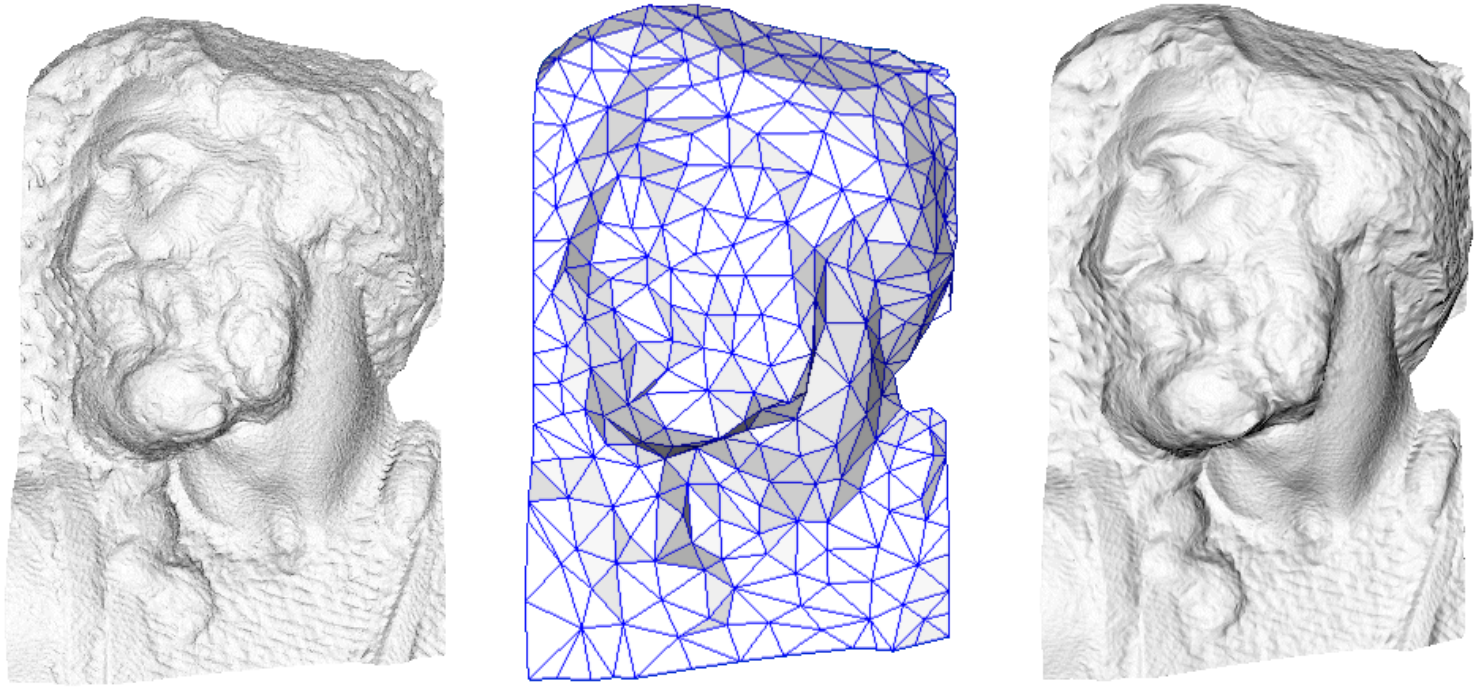

original mesh
4M triangles

simplified mesh
500 triangles

simplified mesh
and normal mapping
500 triangles

Form of compression

Represent small visible details as deviations from the large-scale surface

Can be invisible to artist

…but not to the renderer!

Need to account for multiple bounces of light

Have complete surface<->light model

Illumination of a surface depends on light coming from all directions

Incoming light is from distant environment, or from another surface

…in which case the incoming intensity is another complicated integral

Apply statistical methods + Moore's Law

Full rendering equation is a large recursive integral

Use Monte Carlo methods and a lot of computing power

Animated movies are now "path-traced" using physically-based methods

As of 2014 (*Big Hero 6*) Disney Animation had a top-100 supercomputer (55,000 CPUs)

What if we want an interactive system?

In 2006, Geomerics set out to solve the global illumination problem

…for gaming platforms (Xbox 360, PlayStation 3, PC)

…in real time

…for all games


STAR WARS
BATTLEFRONT
EA


FPS: 68
TEMP: 43 ºC
PNG: 78 MS
RTT: 122 MS
IND: 53 MS
:20
PREPARAD VUESTRAS DEFENSAS
A
DEFENDER
200 / 200
0
5/5
MAYÚS IZQ
E


0-0
0:00
REAL MADRID C.F.
ATLÉTICO DE MADRID
REAL MADRID
FIFA 18

Strategy: understand shape at global scale

Goal: encode intra-scene visibility in smallest possible bitrate

Build a geometrical hierarchy of world based on position, normal direction and visibility

Geometry which is “seen together” is grouped together

Geometrical hierarchy

Level 0

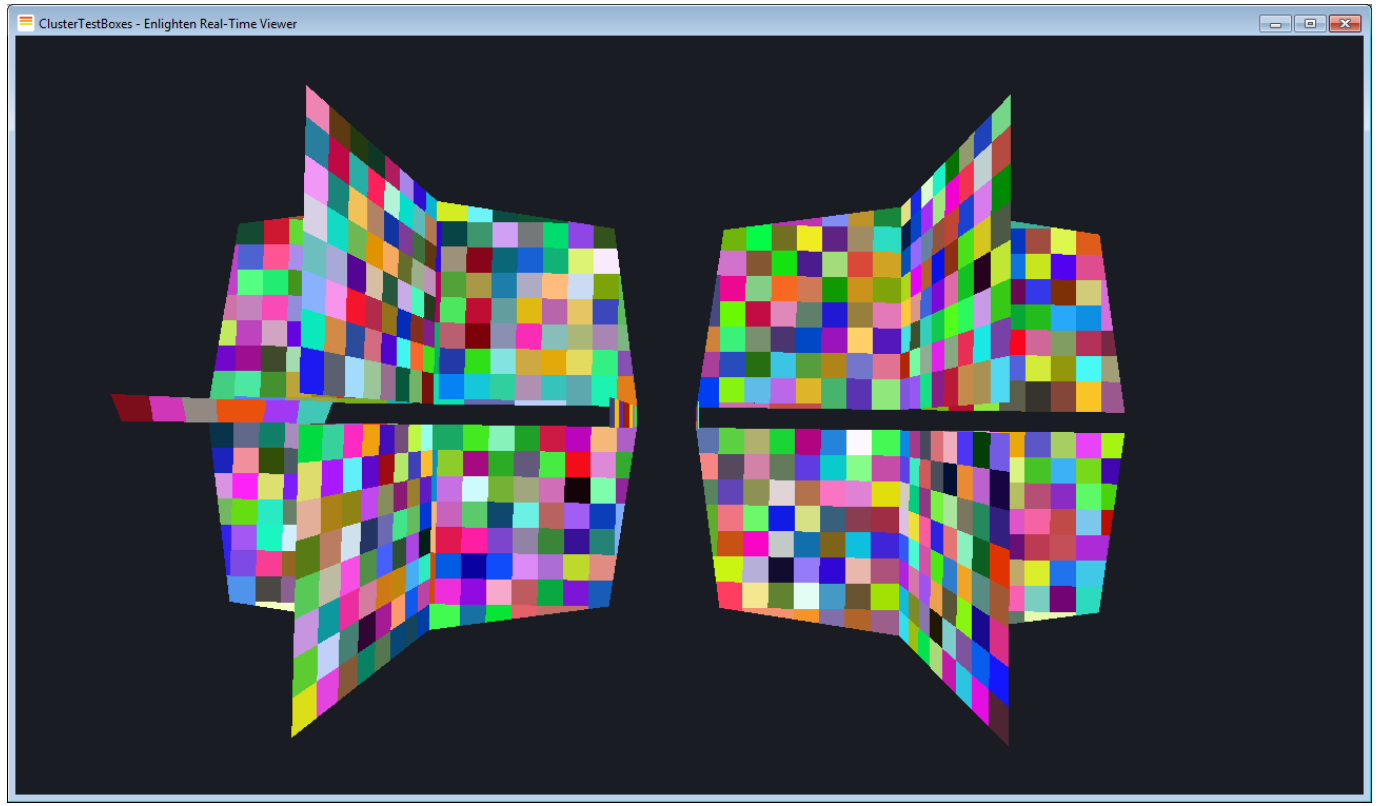

Geometrical hierarchy

Level 5

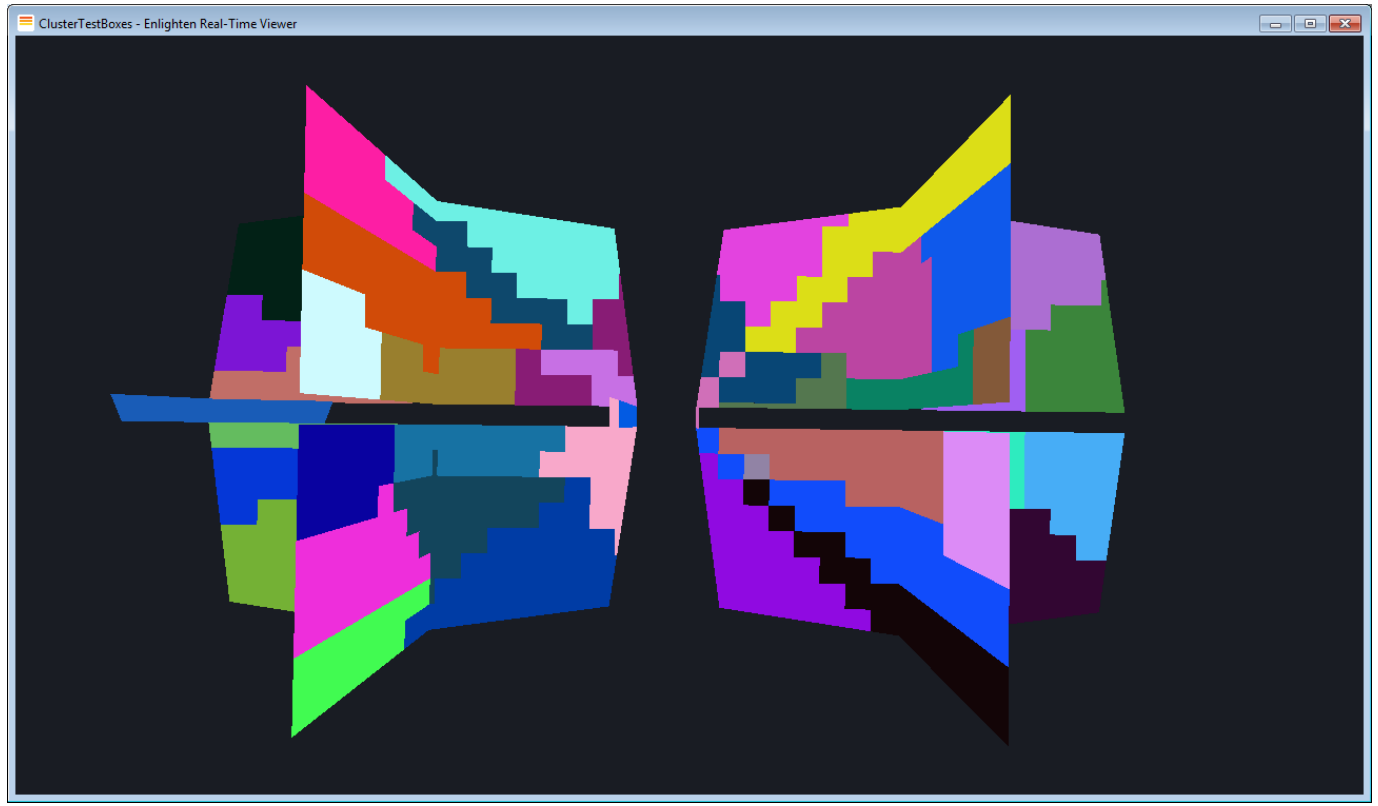

Geometrical hierarchy

Level 8

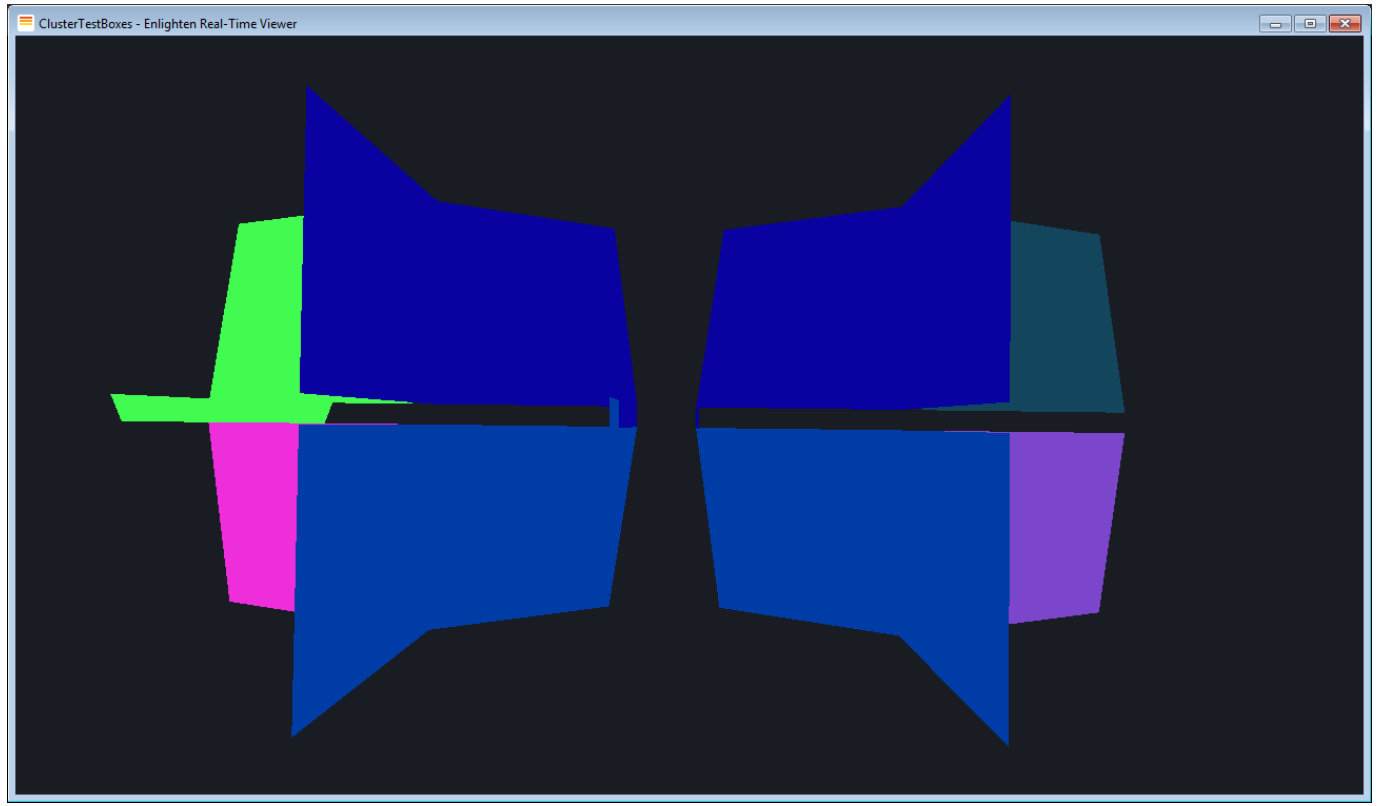

Challenge: game world construction

Common to use a “kit part” strategy to model worlds

In real life, if you have two places for a window you need two windows

In virtual world, you can put *the same window* in two different places

Game artists are not architects

Game worlds are not physically constructed – don’t care about the bits you can’t see!

Common for meshes to interpenetrate and geometry to be redundant

Need to detect and cope with this in hierarchy creation

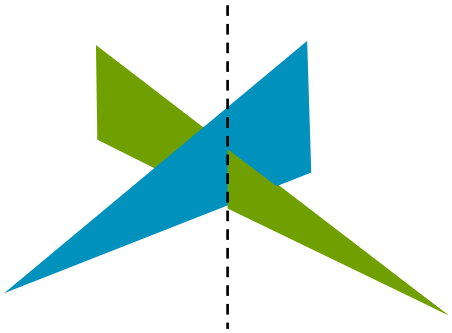

Understand shape at indirect lighting scale

Compute indirect lighting for low detail “target mesh”

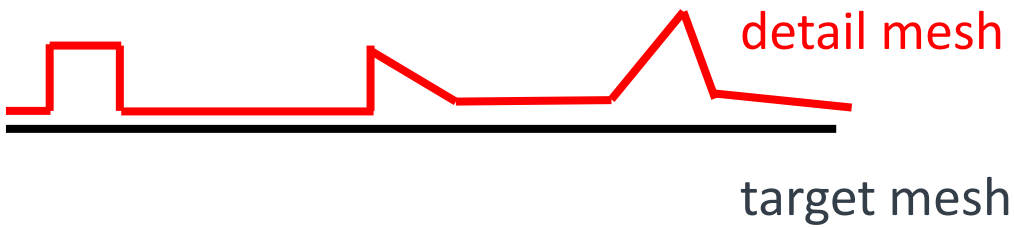

Project detailed geometry on to target mesh to compute lighting coordinates

Use optimised shader model for “relighting” arbitrary normal directions

This worked, more or less, but…

Originally, target meshes were all hand-authored

- Extremely time consuming and error-prone

Multiple years spent developing automated methods and better tools

Enabled integration as lighting solution for *Unity 5*

Put *Enlighten* in the hands of millions of developers

But… *Photometria* showed simple primitives are simple

Could we create an algorithm to represent arbitrary scenes as hierarchies of primitives?

- Planes, spheres, cylinders with local deviations

This is a Computer Vision / Machine Learning problem

- If we were starting today we would approach it quite differently

Closely related to problems in other fields (eg robotics)

Sadly, I'm no longer working on it ☹

Thank you!