KIVI Engineering Society

KIVI Chair - Big Data Science Master Class

Computer Graphics and Visualization

prof.dr. Elmar Eisemann

10 November 2016

DELFT DATA SCIENCE

[Ritschel, Eisemann, Ha, Kim, Seidel - CGF 2011]
[Kehl, de Haan – Gi4DM 2012]
[Kroes, Post, Botha -PLoS ONE 2012]

Introduction

- Computer Graphics has many applications
  - Architecture/Design
  - Engineering/Medical
  - Education
  - Movies
  - …

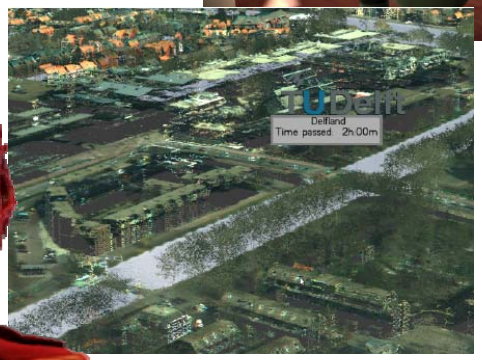

DELFT DATA SCIENCE

Sensors everywhere…

- Big brother is measuring you…

There will be a trillion (!) sensors within the next 5 years…

E.g., Google Earth already stores 21 Million Gigabytes

How can we make use of this data?

Data Development

- **Processing**
- **Analysis**
- **Interaction**
- **Visualization**
- **Guidance**

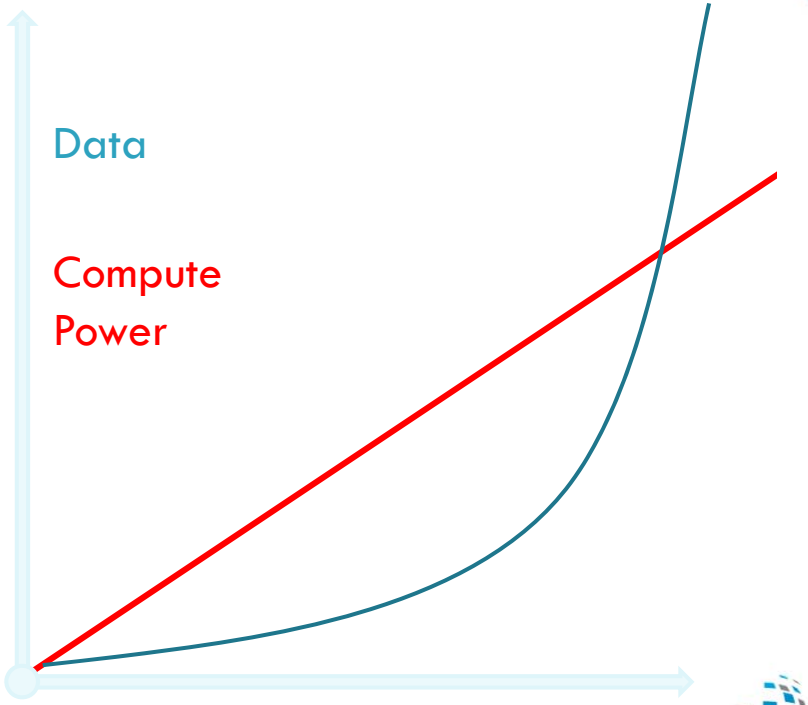

DELFT DATA SCIENCE

Modern Visualization Systems

Very complex:

- Large-scale Data Handling
- Effective Illustrations
- Responsive Interactivity
- Analysis Solutions

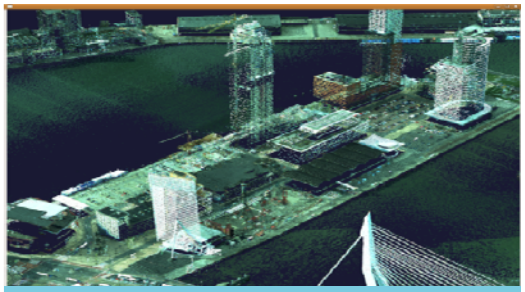

Massive 3D Pointclouds

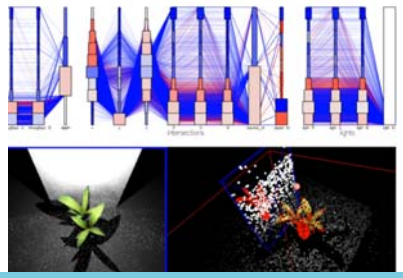

Light-Transport Analysis

Flooding Simulation

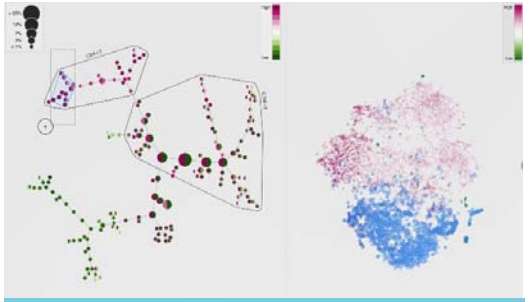

Cell-type Analysis

Volume Rendering

Multi-touch Data Mining

What topics to cover? …

>20 topics in 3h…

- Realistic Rendering
- Graphics in Data Visualization
- Aesthetic imagery
- Human perception
- Real-time processing
- Real-time large-scale data rendering
- Compression
- Multi-dimensional data
- Analysis (with Machine Learning Components)
- Basic graphics approaches
- Aeronautics
- Object Simplification
- User experience/Interaction
- Solutions used in practice
- Collision Detection
- Fabrication
- Geoscience …

DELFT DATA SCIENCE

Effective Data Visualization Requires

- Large-Scale Rendering

- Visualization and Percepti

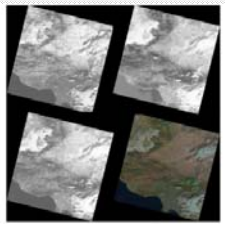

- Data Analysis

DELFT DATA SCIENCE

Large-Scale Rendering

- Ray Tracing
- Graphics Pipeline
- Specialized Methods for Different Data Types
  - Height-Field Data, Voxel Data, Data Management, Compression

Visualization and Perception

- Realistic Rendering
- Perceptual Methods
- Visualization & Interface

Data Analysis

- High-dimensional/Heterogeneous Data
- Dimensionality Reduction
- Visual Analytics

Large-Scale Rendering

- Ray Tracing
- Graphics Pipeline
- Specialized Methods for Different Data Types
  - Height-Field Data, Voxel Data, Data Management, Compression

How to produce an image?

- Computers can only calculate…

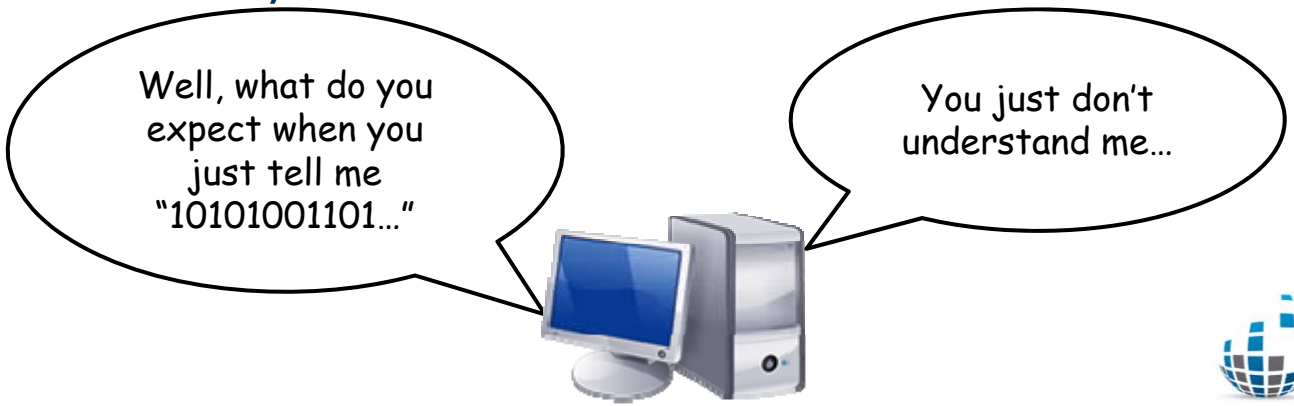

Making images with a Computer

Pixel

DELFT DATA SCIENCE

Pixels – Picture elements

- A colored pixel has typically Red Green Blue values.

- We color by numbers…

  sounds simple…

  but choosing the values can be difficult

Producing Images in the Real World

DELFT DATA SCIENCE

Producing Images in the Real World

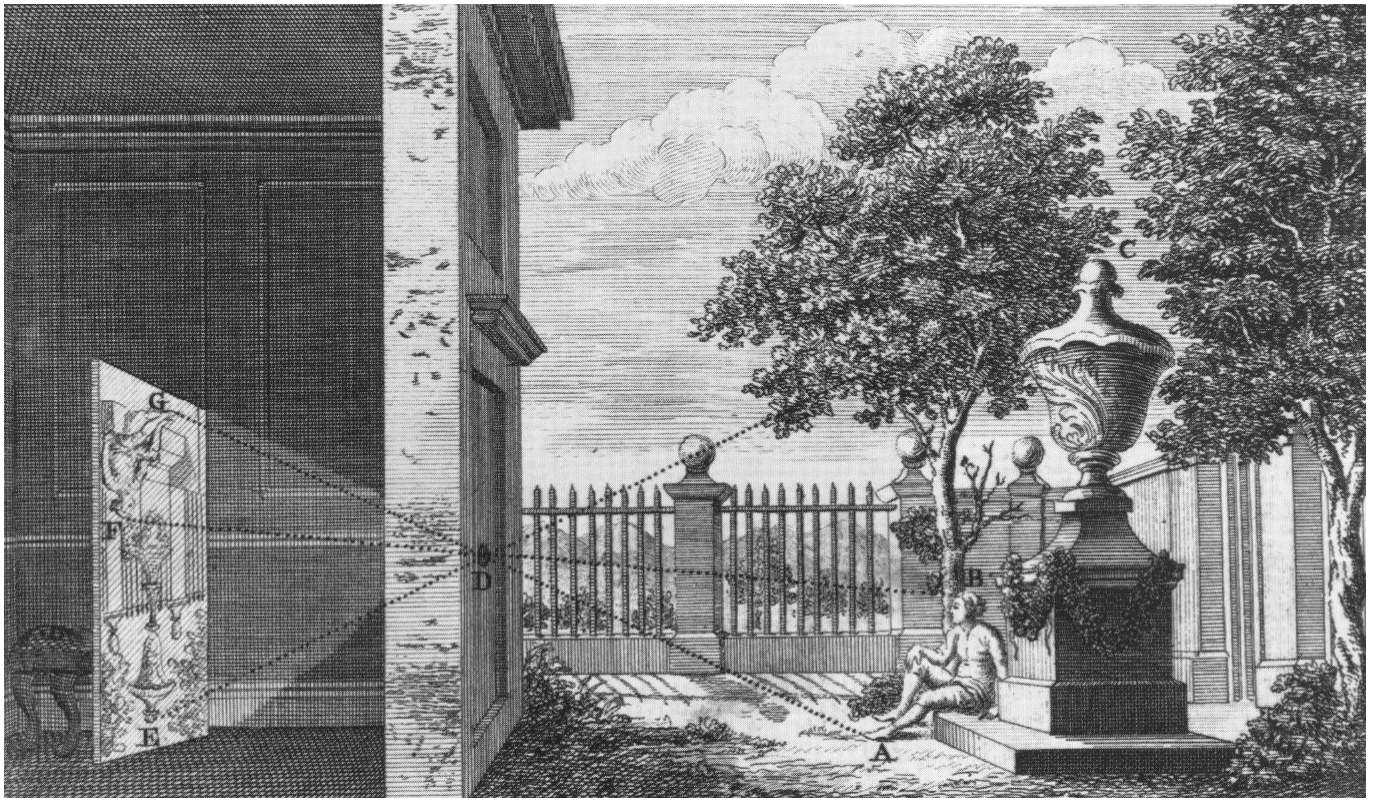

Camera obscura

DELFT DATA SCIENCE

Pinhole camera

- Box with hole
- Perfect image for “point-sized” hole

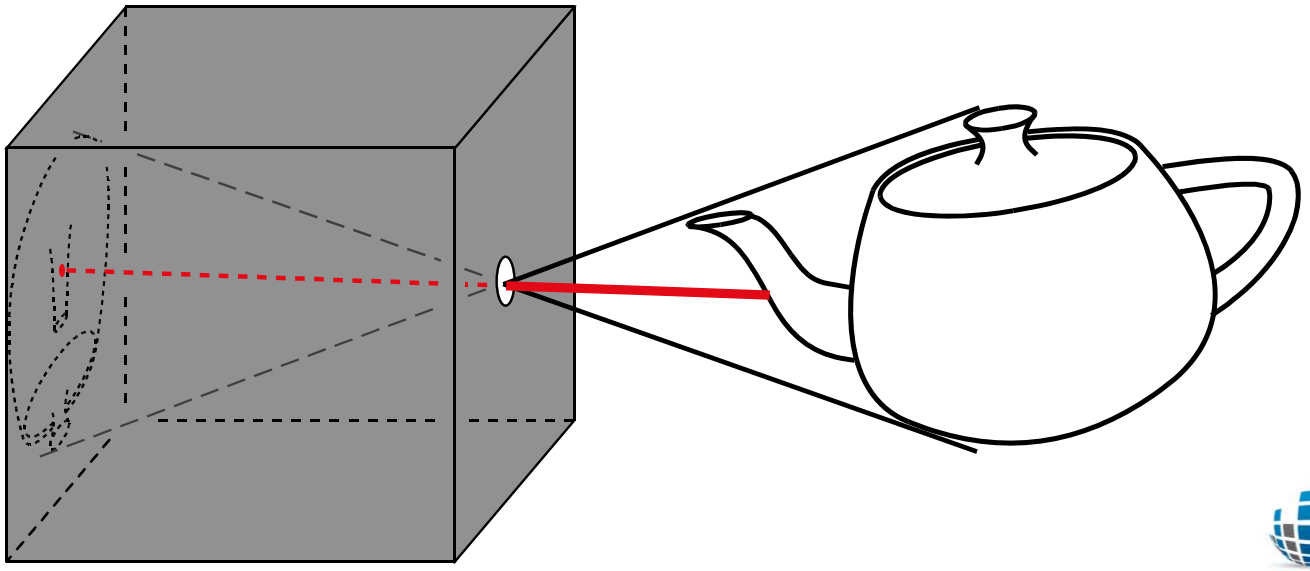

DELFT DATA SCIENCE

Virtual Camera

- Take a pixel on the image in the virtual world
- Compute ray through pixel and camera center
- Intersect ray with scene [?]

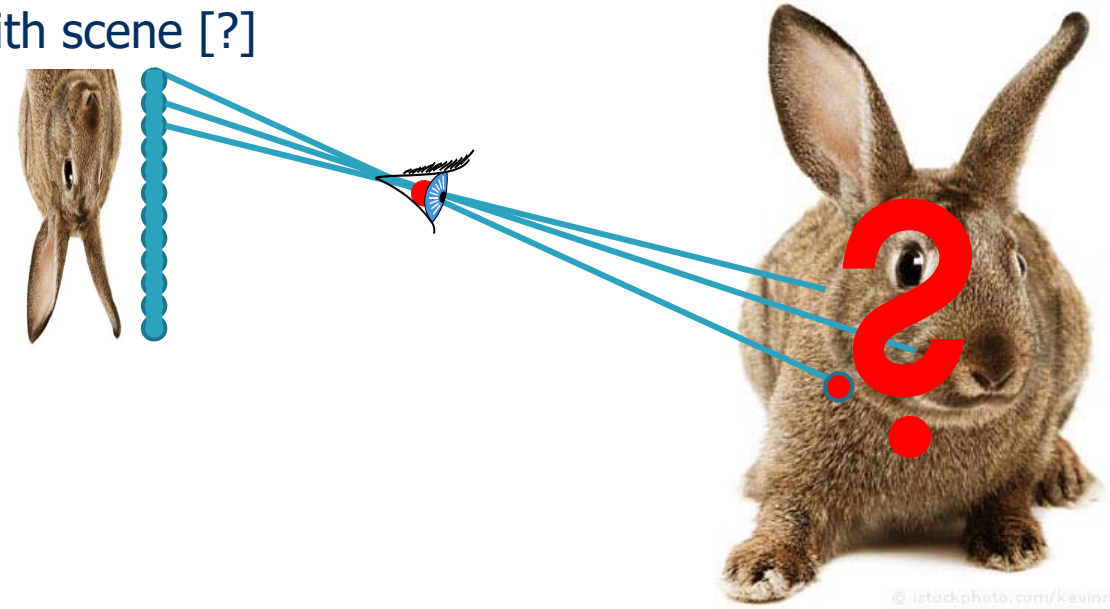

DELFT DATA SCIENCE

Models

- Representing reality

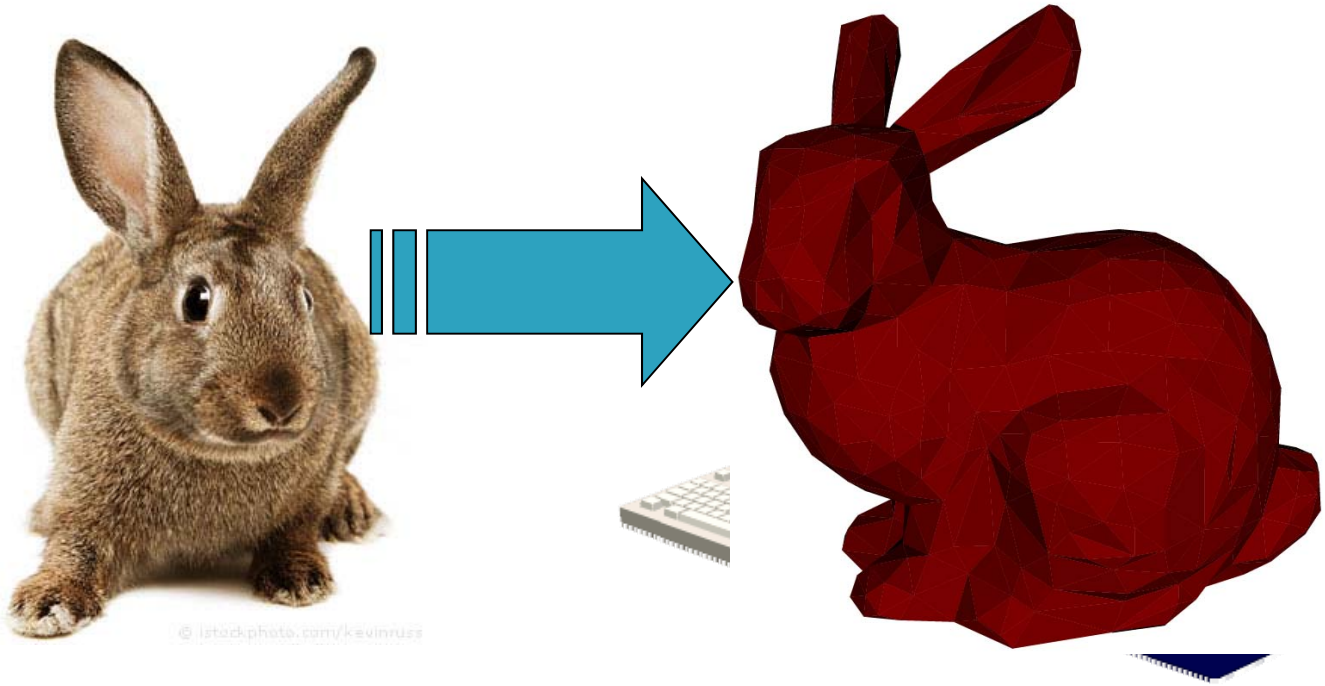

DELFT DATA SCIENCE

Models

- Models are typically lists of triangles

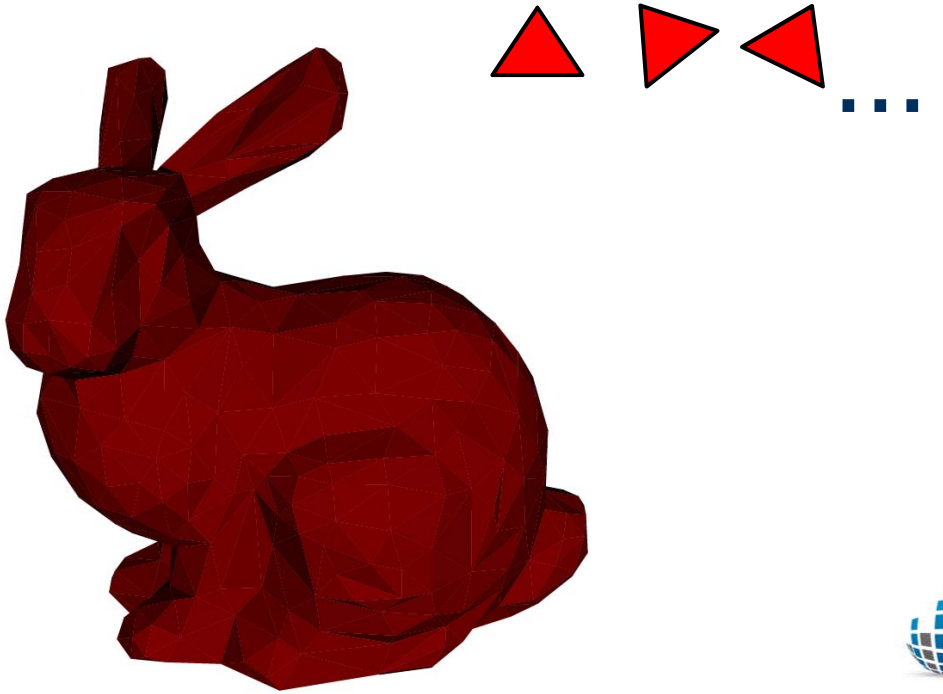

Solve Equations

- $R+tD= a+\beta\,(b-a)+\gamma\,(c-a)$

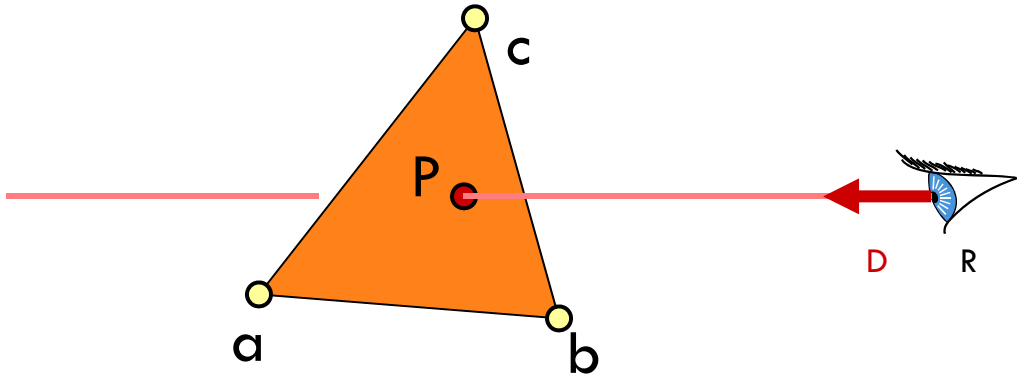

Solve Equations

- $R_x+tD_x= a_x+\beta (b_x-a_x)+\gamma (c_x-a_x)$
- $R_y+tD_y= a_y+\beta (b_y-a_y)+\gamma (c_y-a_y)$
- $R_z+tD_z= a_z+\beta (b_z-a_z)+\gamma (c_z-a_z)$

$$\begin{bmatrix} a_x-b_x & a_x-c_x & D_x \\ a_y-b_y & a_y-c_y & D_y \\ a_z-b_z & a_z-c_z & D_z \end{bmatrix} \begin{bmatrix} \beta \\ \gamma \\ t \end{bmatrix} = \begin{bmatrix} a_x-R_x \\ a_y-R_y \\ a_z-R_z \end{bmatrix}$$

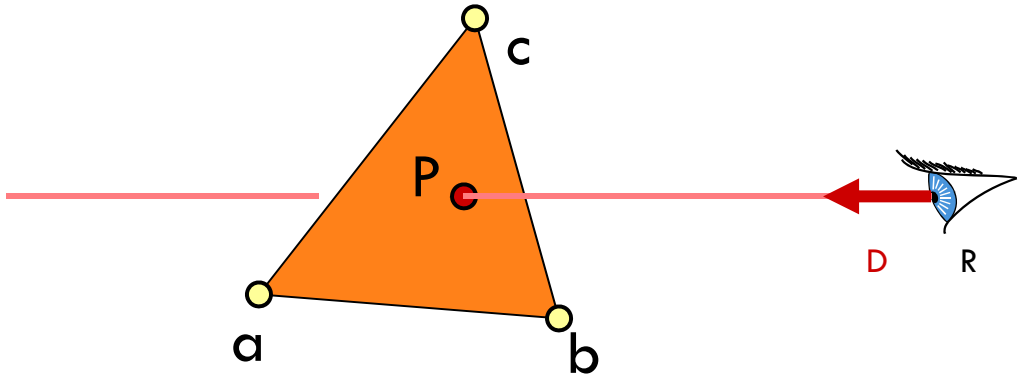

Produce Final Image

- Keep the closest intersection point

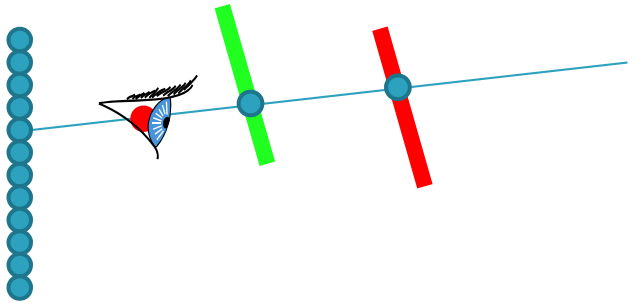

Ray Tracing - Recap

For each pixel

```
Distance=MAX
Color=0
Ray=computeRay(pixel)
For each triangle
    (CurrColor,CurrDistance)=computeIntersection(Ray)
    If (CurrDistance<Distance)
                Distance=CurrDistance
                Color=CurrColor
```

Ray Tracing - Cost

```
For each pixel
    Distance=MAX
    Color=0
    R=computeRay(pixel)
    For each triangle
        (CurrColor,CurrDistance)=testIntersection(R)
        If (CurrDistance<Distance)
                Distance=CurrDistance
                Color=CurrColor
```

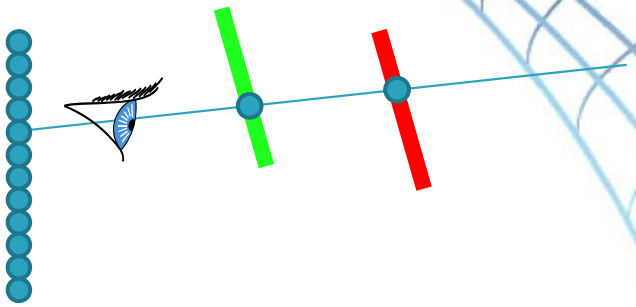

Performance Analysis

- Stupid implementation:
- Ray Tracing:

  Cost = Pixels * Triangles

e.g., 100.000 triangles and a 1000^2 screen:

Raytracing: 100.000 * 1.000.000 = 10^11

Performance Analysis

- Smart implementation:
- Ray Tracing:

  Cost = Pixels * log(Triangles)

  

  + building a structure

e.g., 100.000 triangles and a 1000^2 screen:

Raytracing: 1.000.000 * 5 + X = 5 * 10^6

Especially suitable for static data

What about real-time (30 Images/Sec)

"building a structure" is generally slow

Alternative approach:

Rasterization via the Graphics Pipeline

Virtual Camera

- Camera Plane in front of the eye

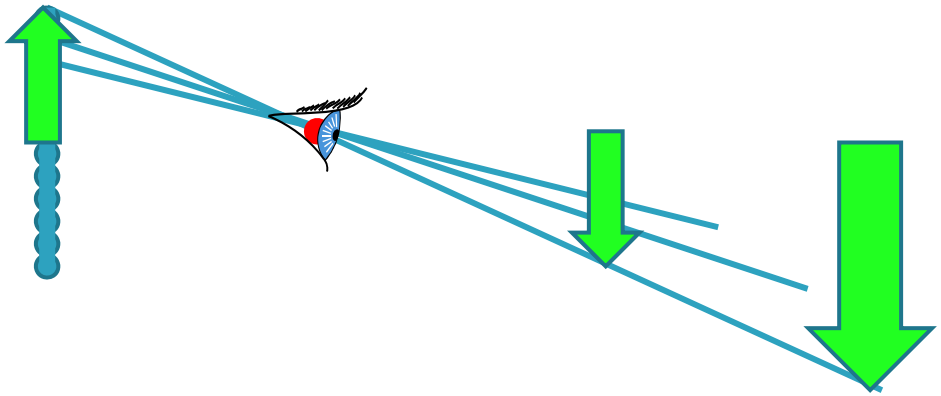

DELFT DATA SCIENCE

Simplified Graphics Pipeline

- Projection: Transform coordinates to screen

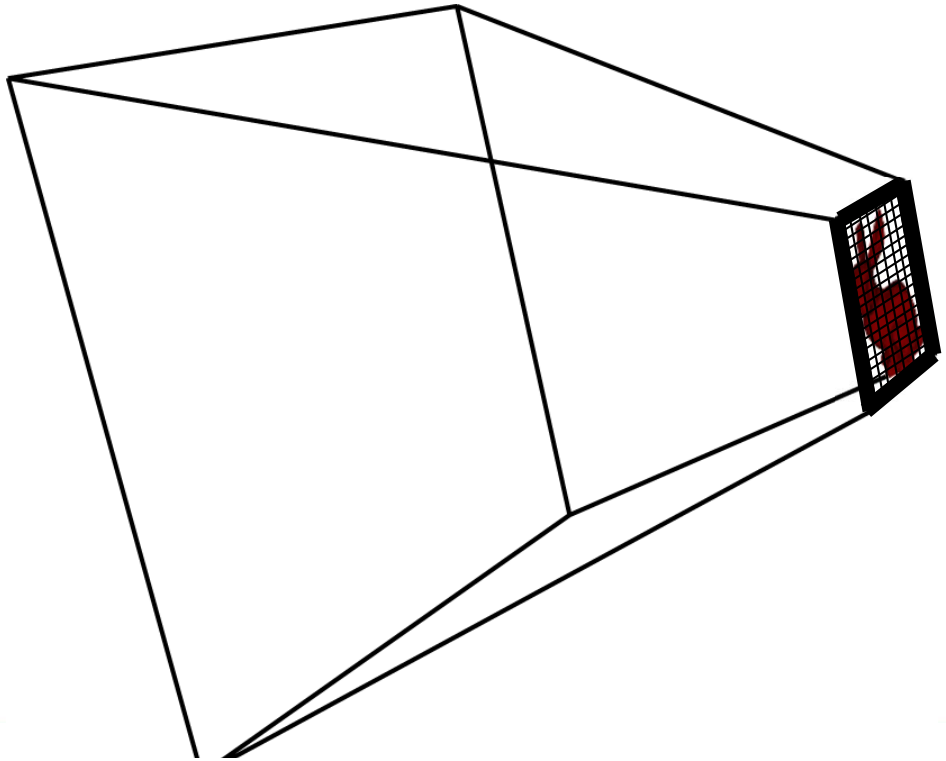

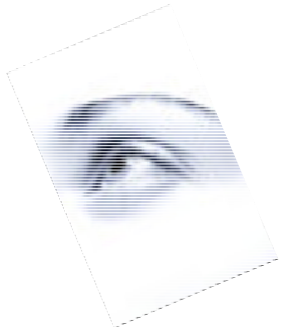

DELFT DATA SCIENCE

Simplified Graphics Pipeline

- Rasterization: Fill screen pixels

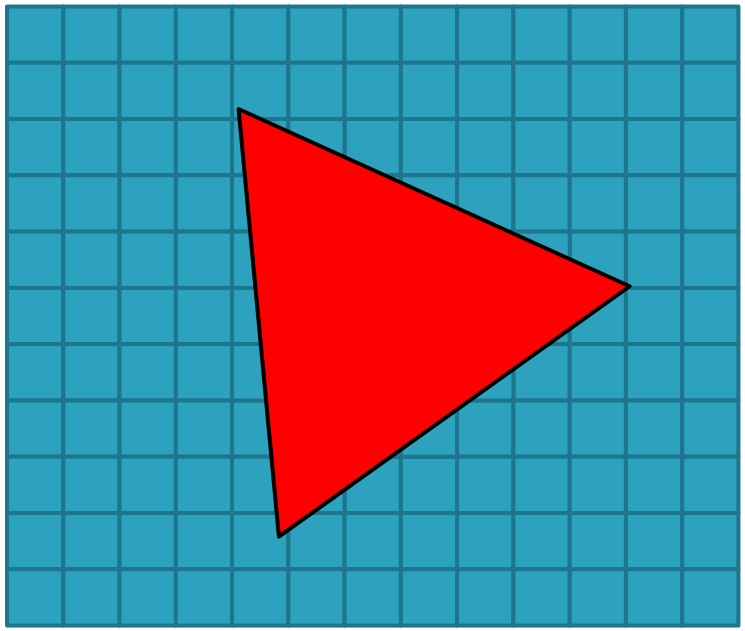

Simplified Graphics Pipeline

- Highly parallelizable ➡ GPUs

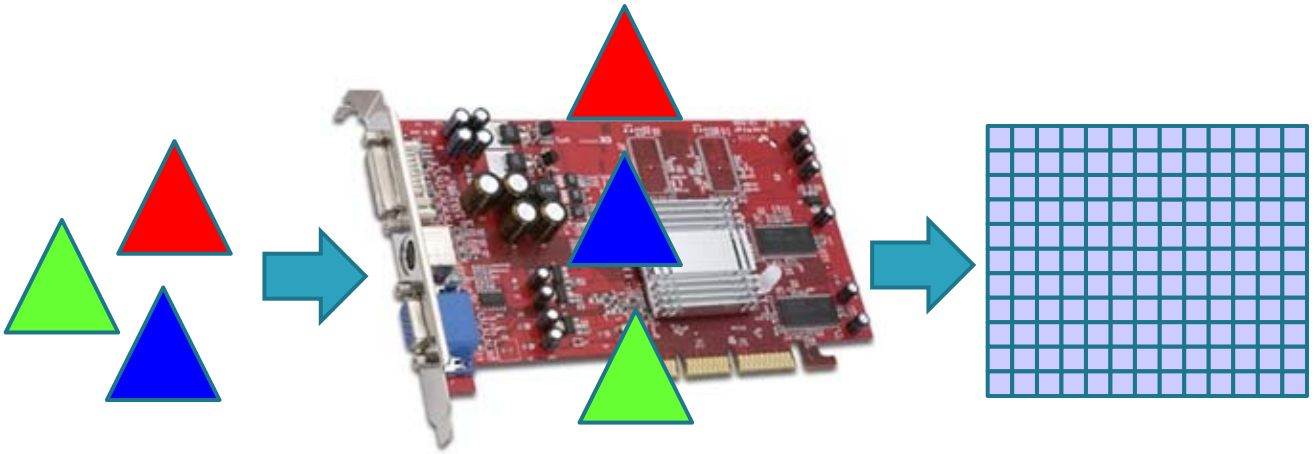

Thousands of processors working in parallel

Simplified Graphics Pipeline

- Catch: Let's look at a second triangle…

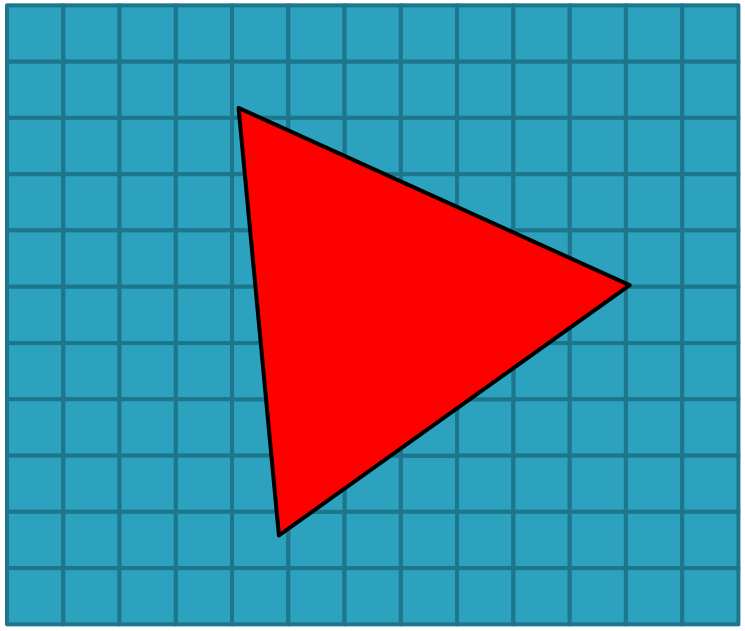

Simplified Graphics Pipeline

- Catch: Let's look at a second triangle…

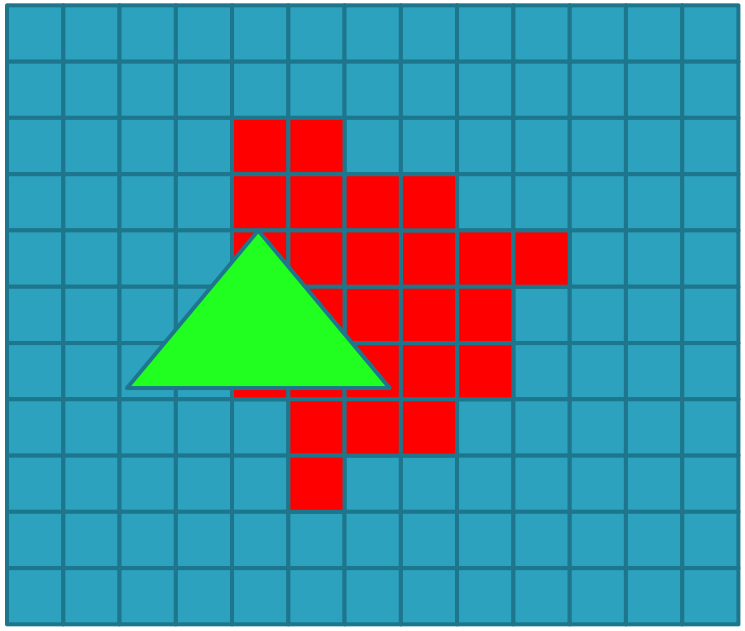

Simplified Graphics Pipeline

- Catch: Triangle drawing order changes result

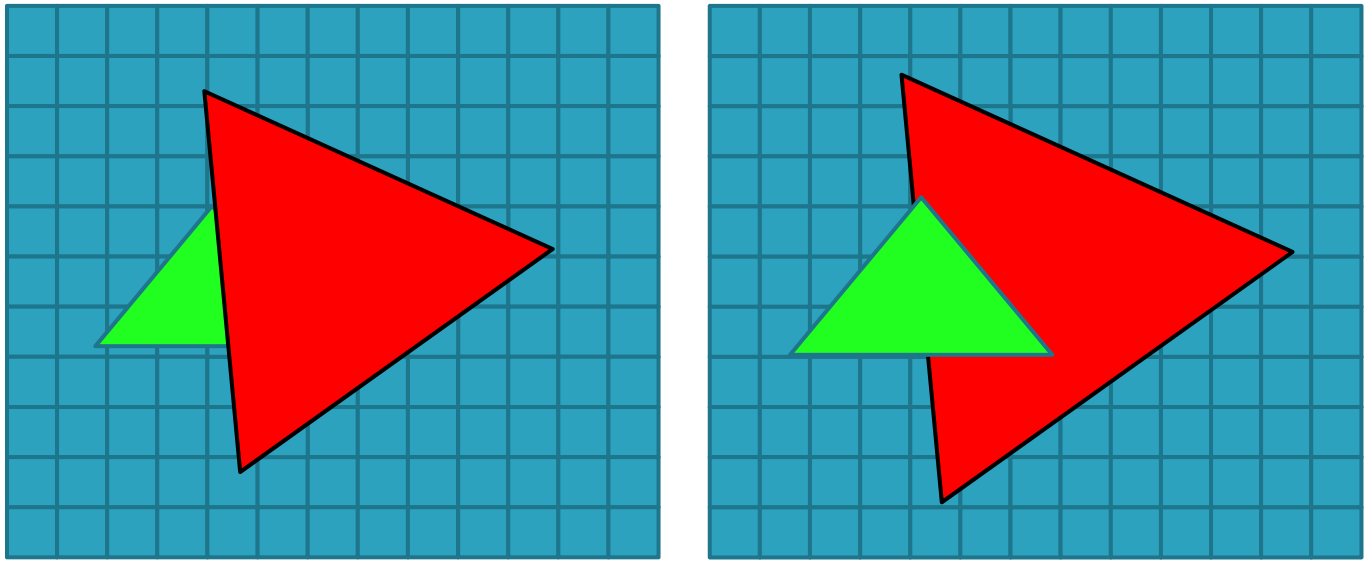

As for ray tracing: need the closest triangle

[Catmull74] , [Strasser74]

Simplified Graphics Pipeline

- Depth Buffer: Avoid sorting triangles!
- Store a color and depth in each pixel

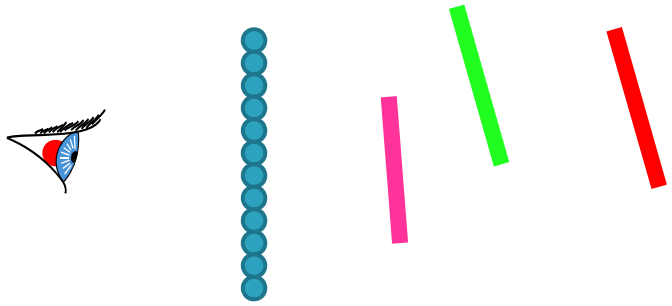

[Catmull74] , [Strasser74]

Simplified Graphics Pipeline

- Depth Buffer: Avoid sorting triangles!
- Store a color and depth in each pixel

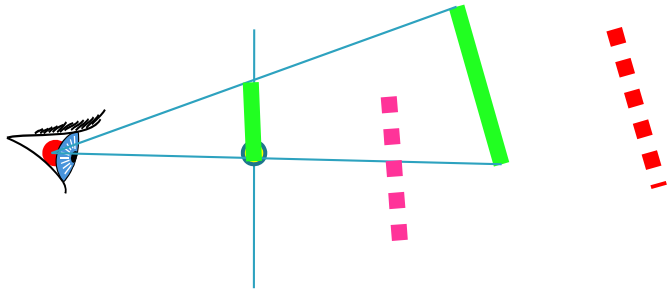

[Catmull74] , [Strasser74]

Simplified Graphics Pipeline

- Depth Buffer: Avoid sorting triangles!
- Store a color and depth in each pixel

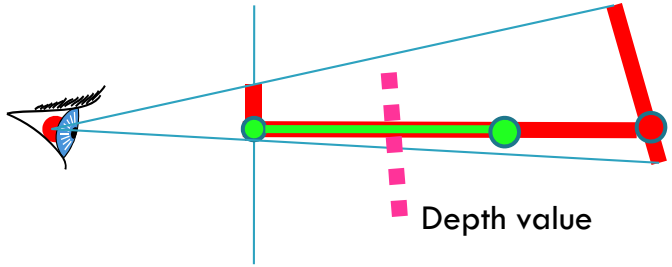

[Catmull74] , [Strasser74]

Simplified Graphics Pipeline

- Depth Buffer: Avoid sorting triangles!
- Store a color and depth in each pixel

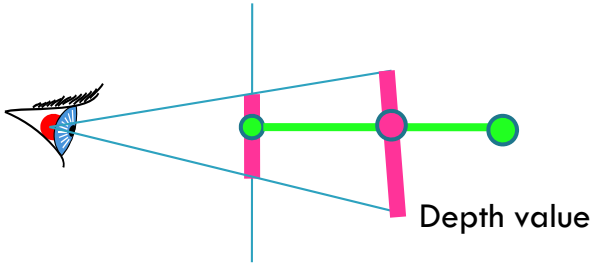

[Catmull74] , [Strasser74]

Simplified Graphics Pipeline

- Depth Buffer: Avoid sorting triangles!
- Store a color and depth in each pixel

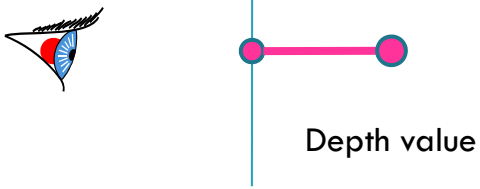

Depth Buffering [1974]

But was applied much later… Why?

Memory requirements

320x200 pixel -> 200 KB of memory !

2000x1000 pixel -> 6 MB of memory !

1974 : $314,573 /MB

1986 : $ 300 /MB

1993 : $28 /MB (Nvidia)

Cost of Rasterization

Algorithm:

For each triangle

    projTri=projectTriangle(triangle)

    fillPixels(projTri)

Cost = Triangles + “drawn pixels”

Performance Analysis

Ray Tracing:

Cost = Pixels * log(Triangles) + structure

vs.

Rasterization:

Cost = Triangles + "drawn pixels"

e.g., 100.000 triangles and a 1000^2 screen:

Raytracing: X+5 * 1.000.000

Rasterization: 100.000 + "drawn pixels"

Raytracing/Rasterization time: ~50

Performance Analysis

Ray Tracing:

Cost = Pixels * log(Triangles) + structure

vs.

Rasterization:

Cost = Triangles + "drawn pixels"

e.g., 100.000.000 triangles and a 1000^2 screen:

Raytracing: X+8 * 1.000.000

Rasterization: 100.000.000 + "drawn pixels"

Raytracing/Rasterization time: ~0.1

… but Rasterization can be made smarter too…

What complexity do we work with?

- Today's Games:
  - 200.000 triangles
- Today's Movies
  - more than 1 Billion

Movie Rendering Costs

>1000 hours!

Big Hero Six – copyright Disney

What complexity do we work with?

- Today's Games:
  - 200.000 triangles

Designed for optimality

Often real-world data

- Visualization Systems:
  - Millions of Data points
  - Potentially summarizing much more information

- Today's Movies
  - more than 1 Billion

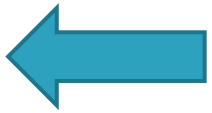

Designed for artistic look

For large-scale data display, we need to mix rendering strategies

DELFT DATA SCIENCE

Questions?

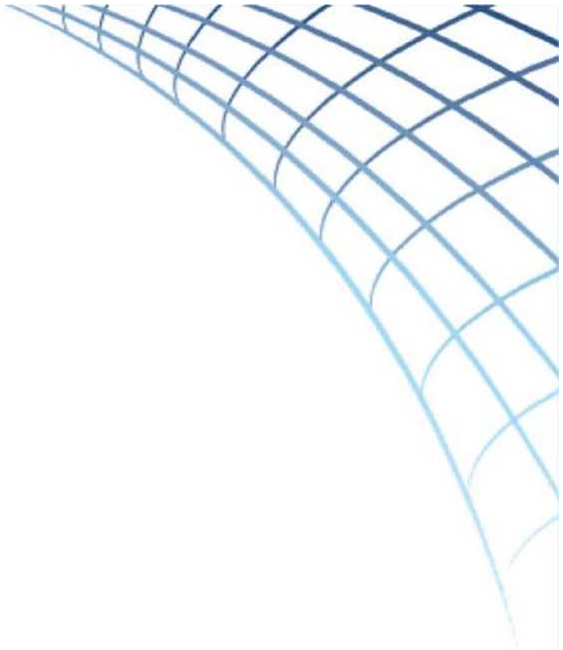

Scanned Data Sets

- AHN2 dataset with satellite imagery ~8 TB of data

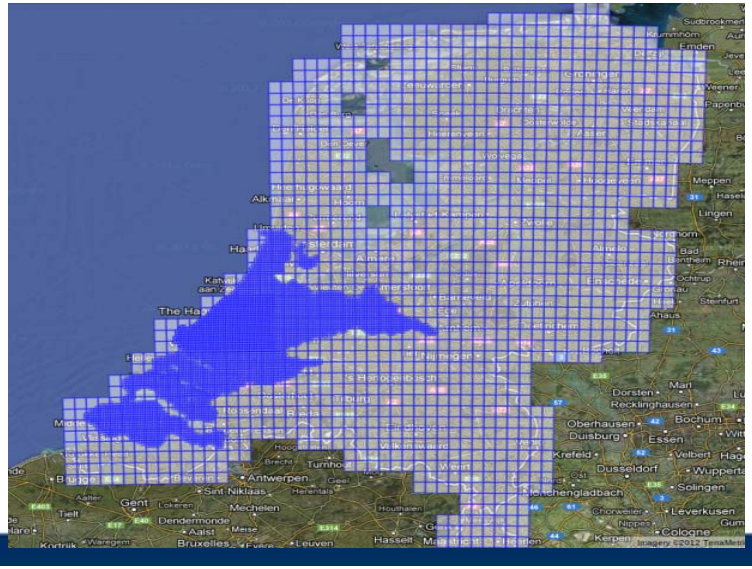

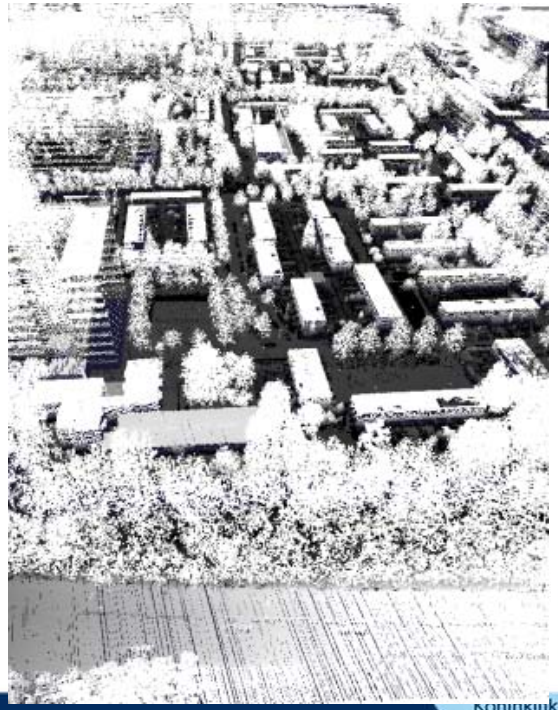

DELFT DATA SCIENCE

Koninklijk Instituut Van Ingenieurs

Scanned Data Sets

DELFT DATA SCIENCE

Koninklijk Instituut Van Ingenieurs

Height-field Representation

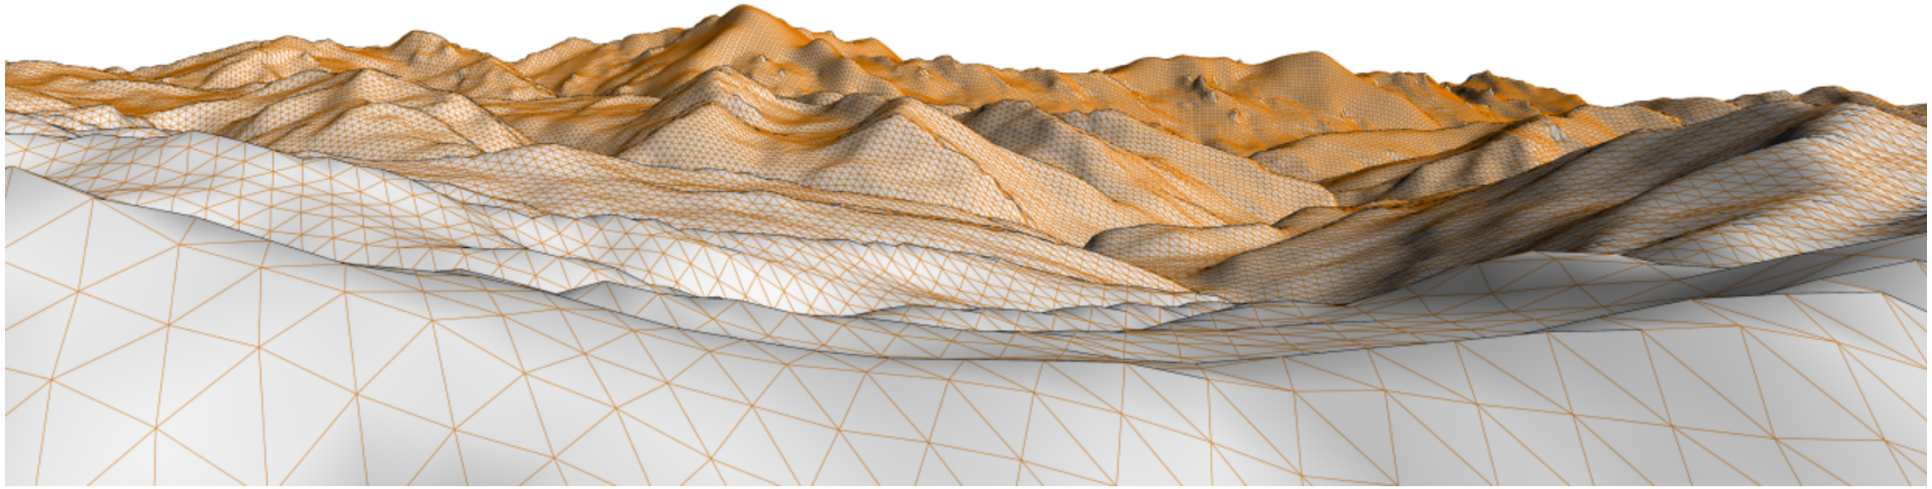

Height-field Representation

- Storage in form of an 2D image with height per pixel

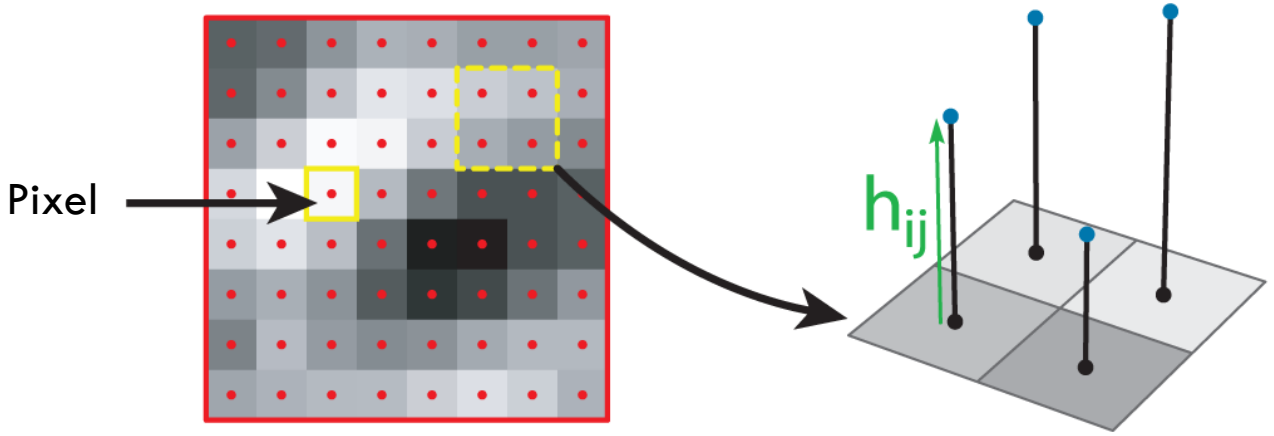

DELFT DATA SCIENCE

Height-Field Rendering

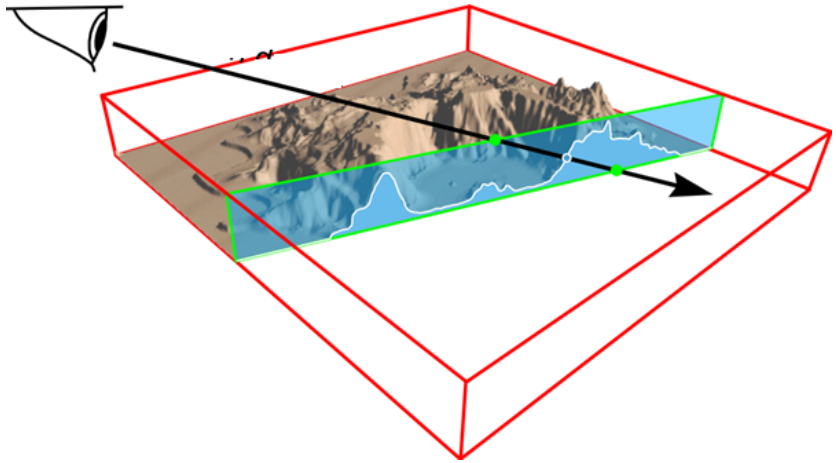

Traditional Height-Field Rendering

- Many steps necessary and limited precision

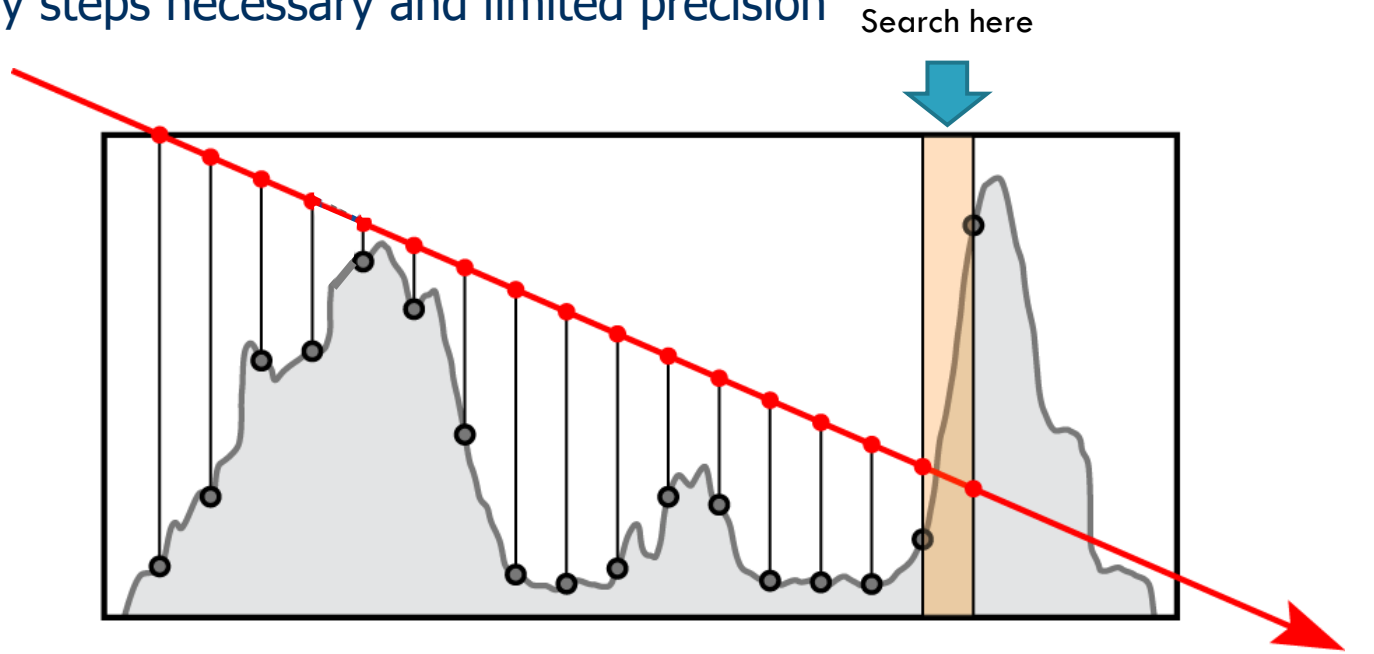

Intersection Algorithm

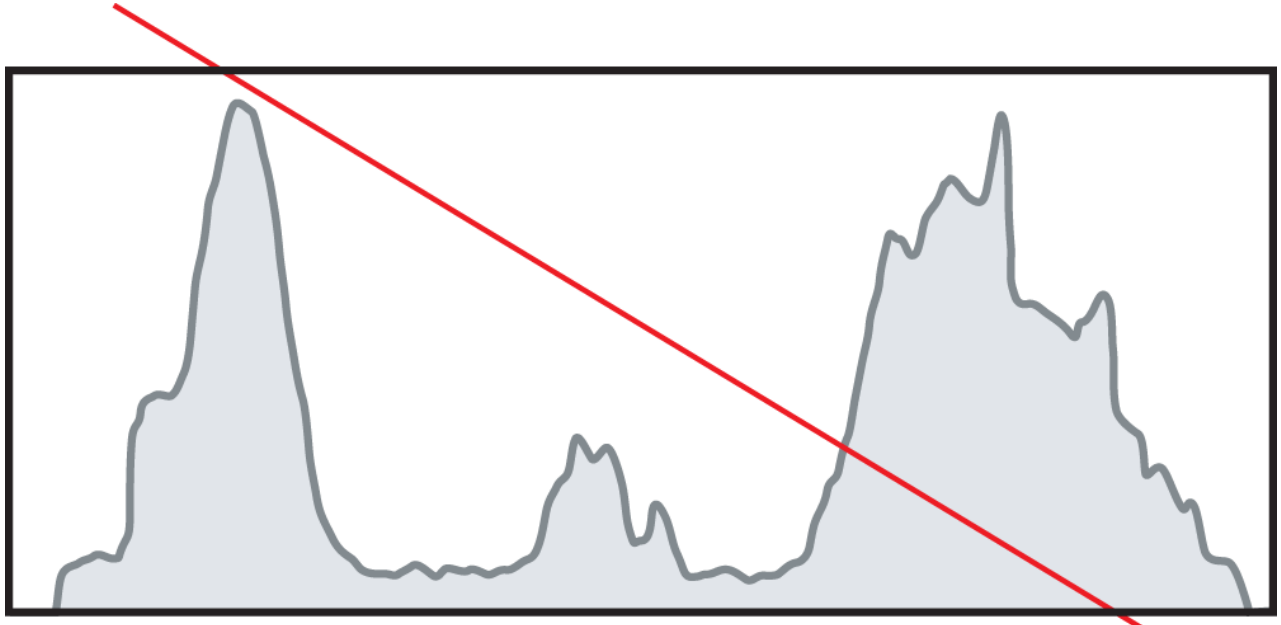

Intersection Algorithm

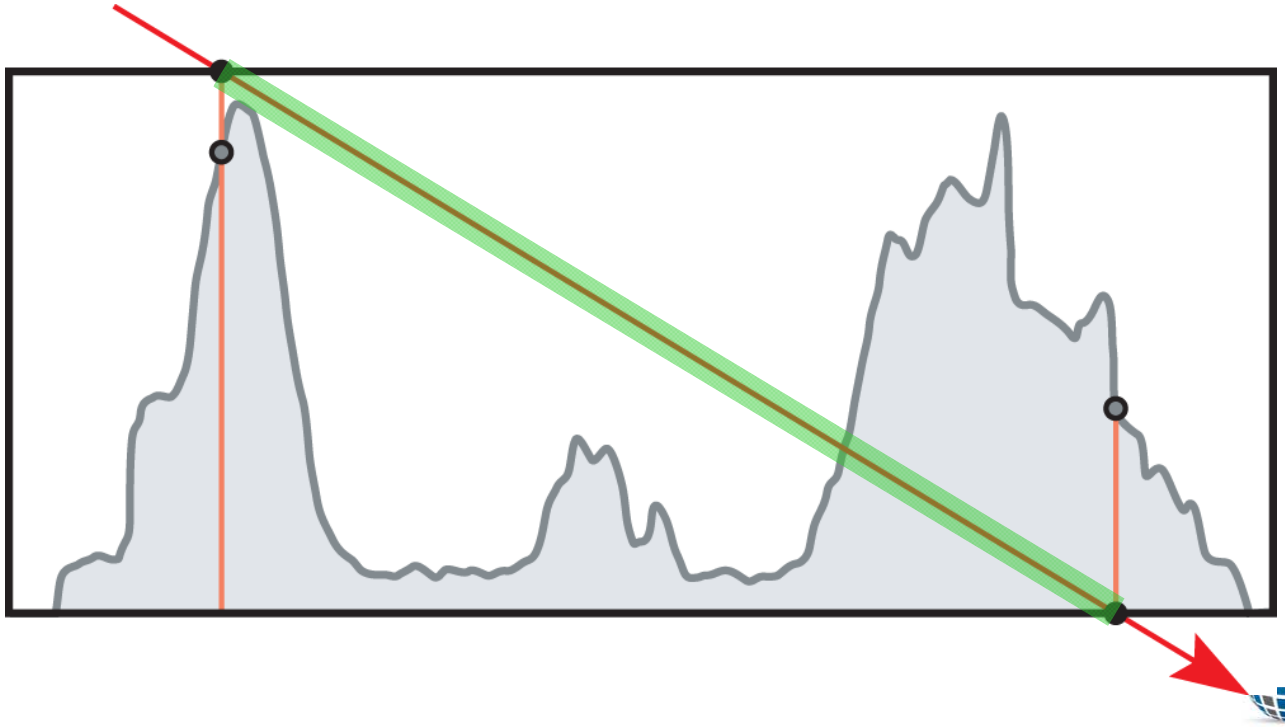

Intersection Algorithm

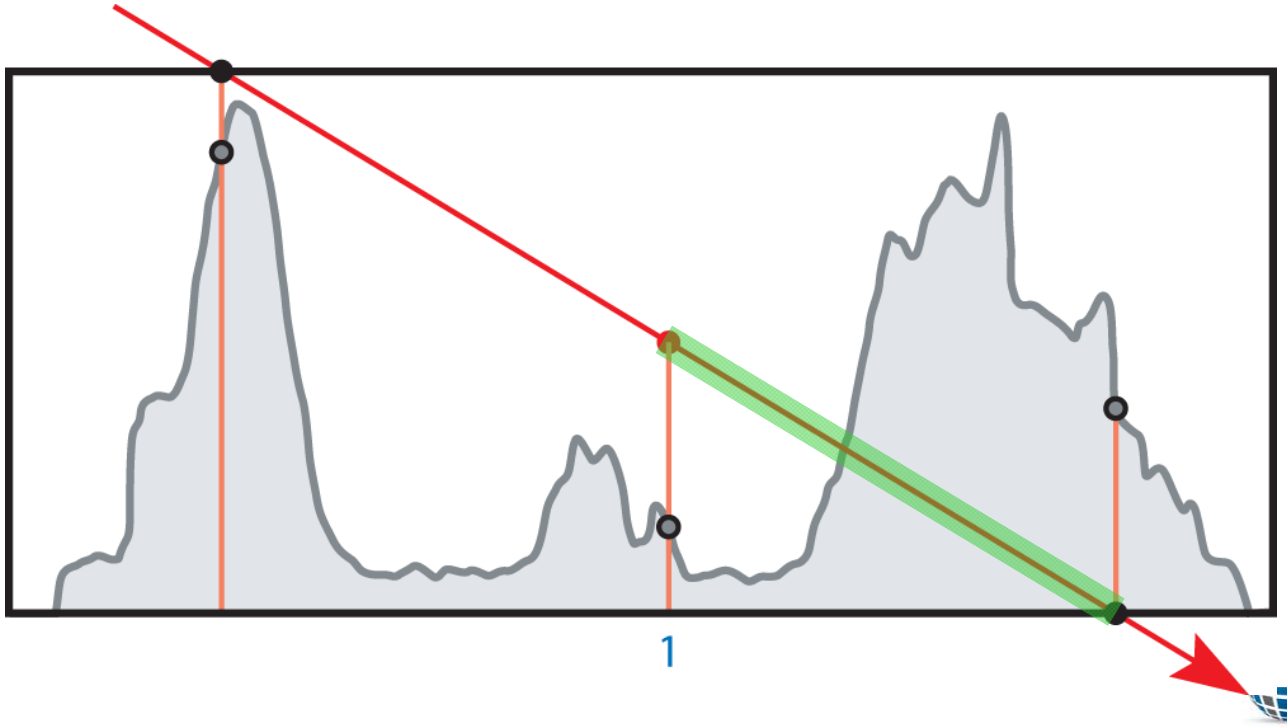

Intersection Algorithm

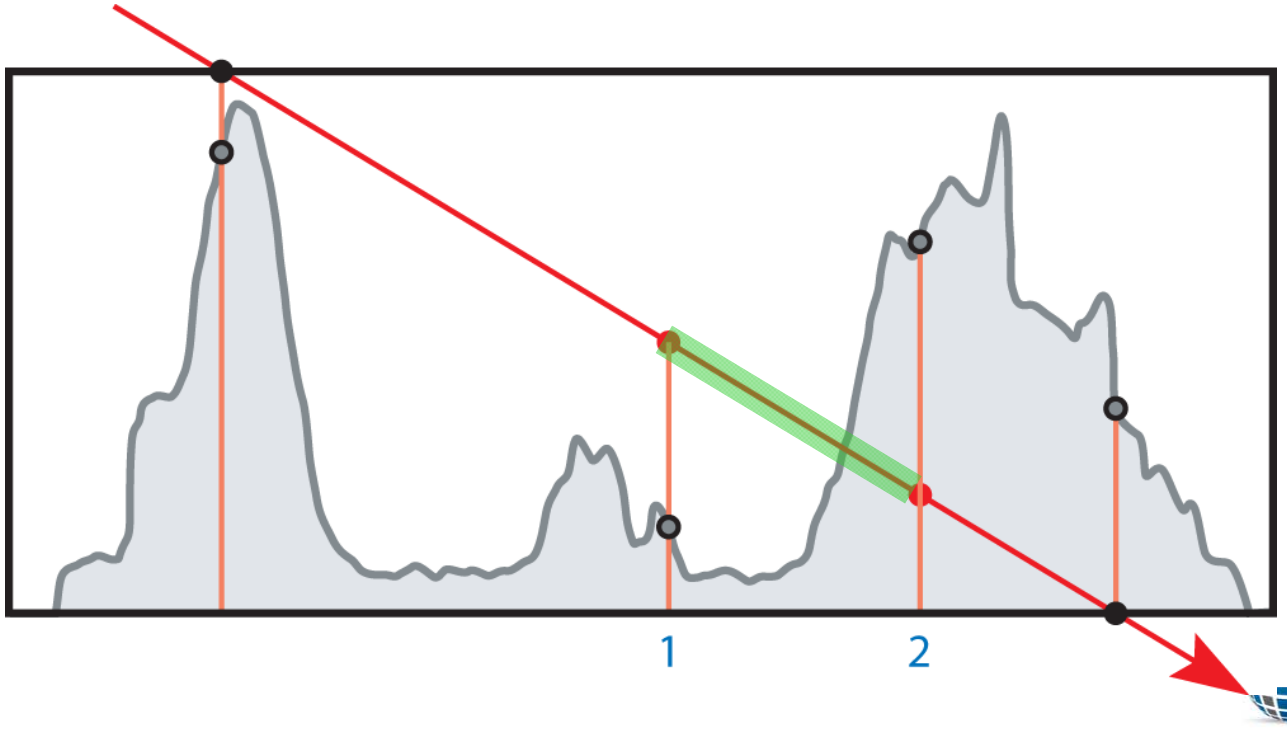

DELFT DATA SCIENCE

Intersection Algorithm

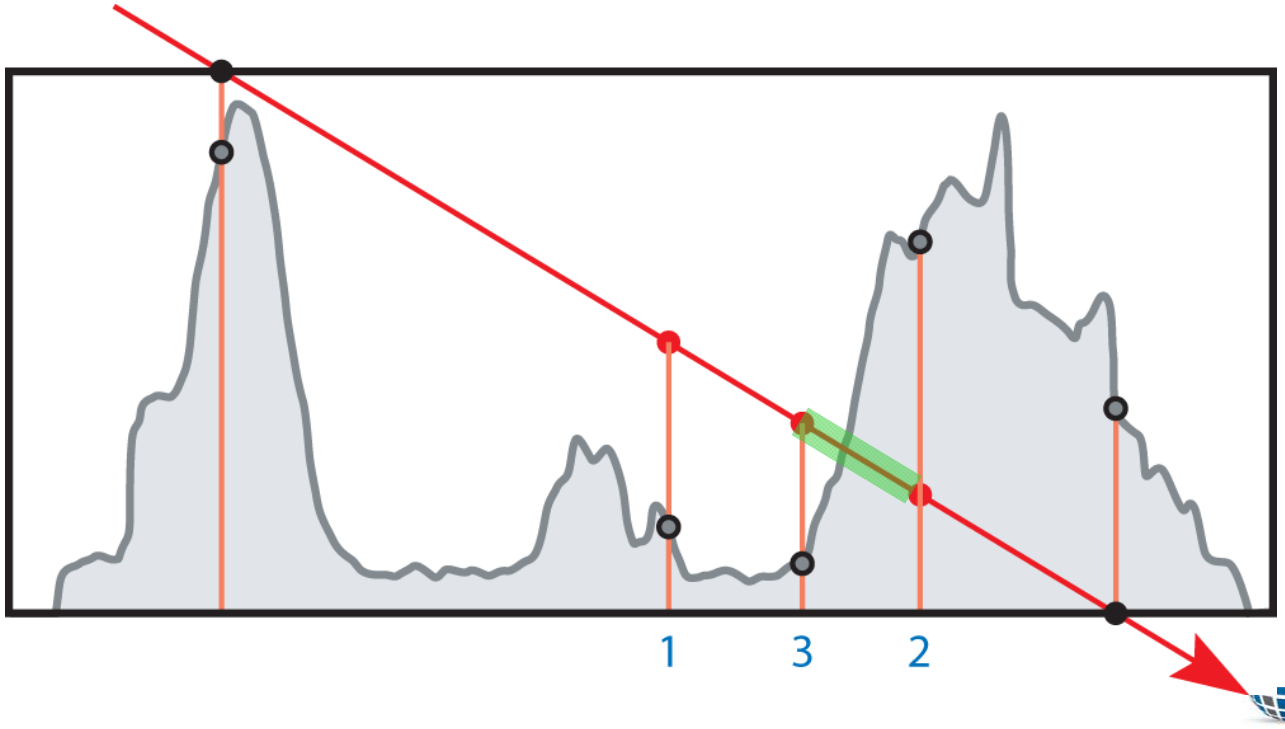

DELFT DATA SCIENCE

Intersection Algorithm

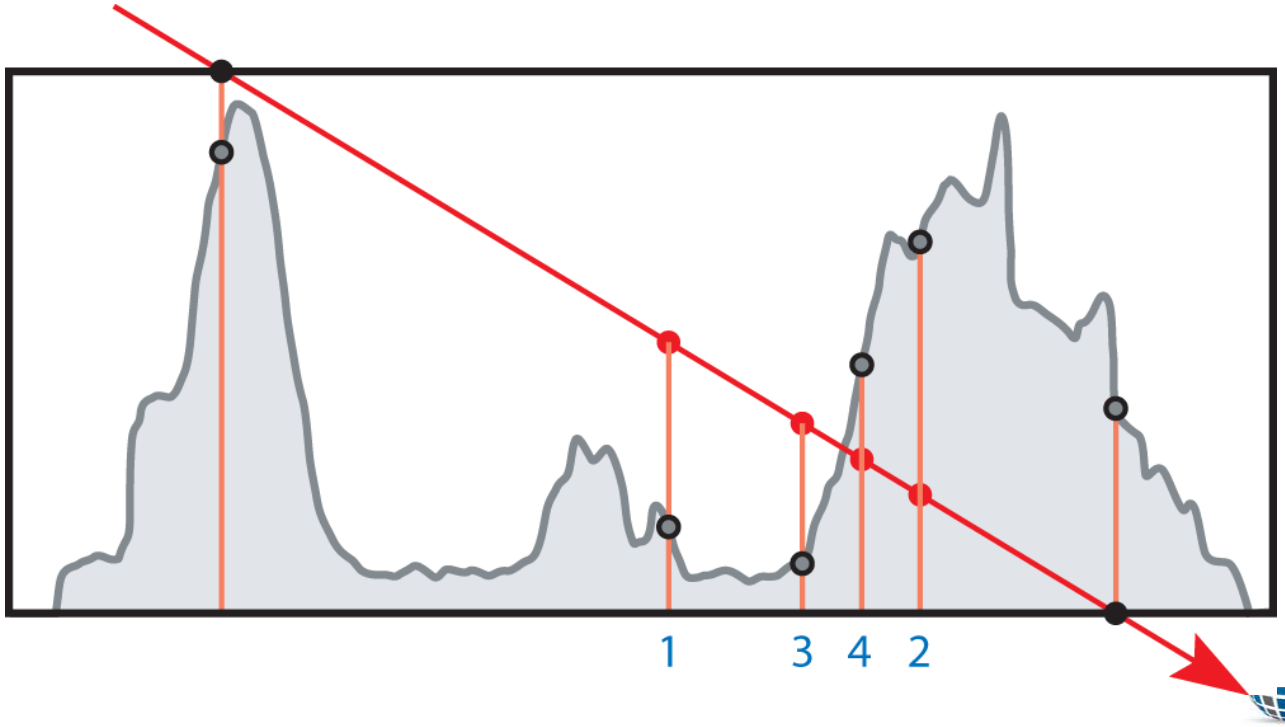

DELFT DATA SCIENCE

Intersection Algorithm

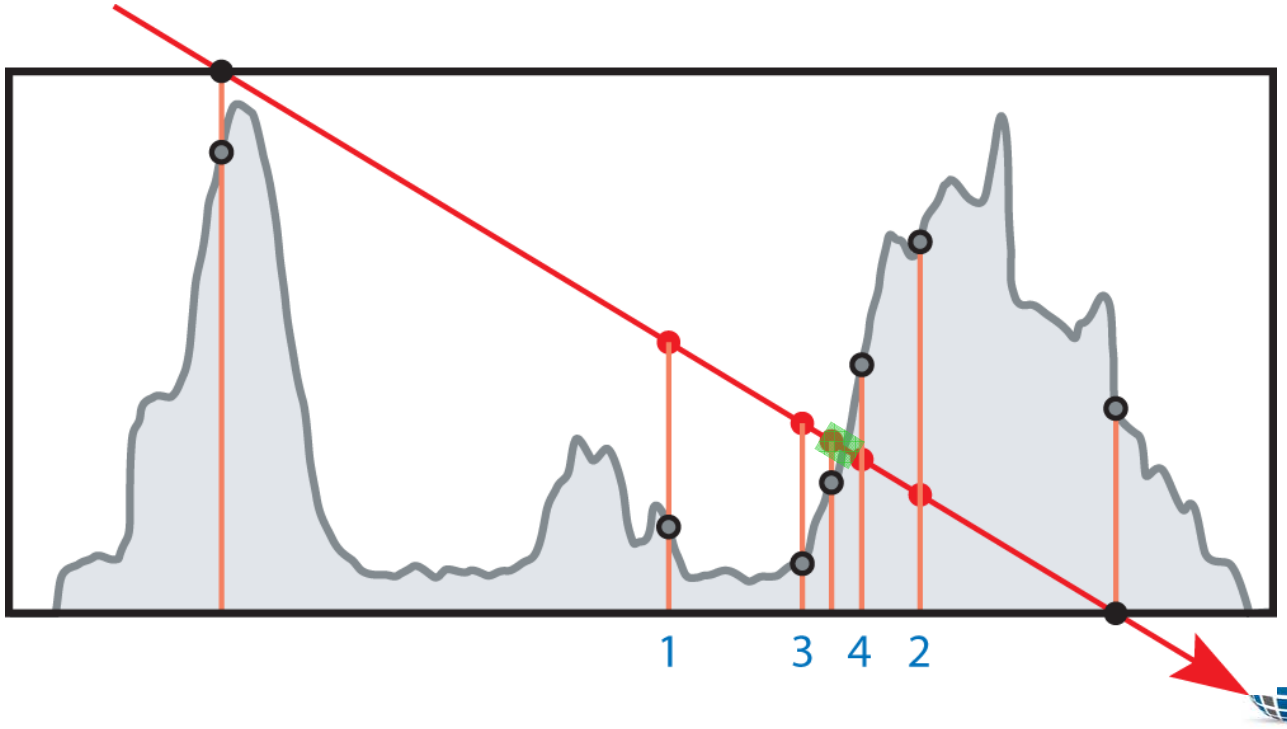

DELFT DATA SCIENCE

Intersection Algorithm

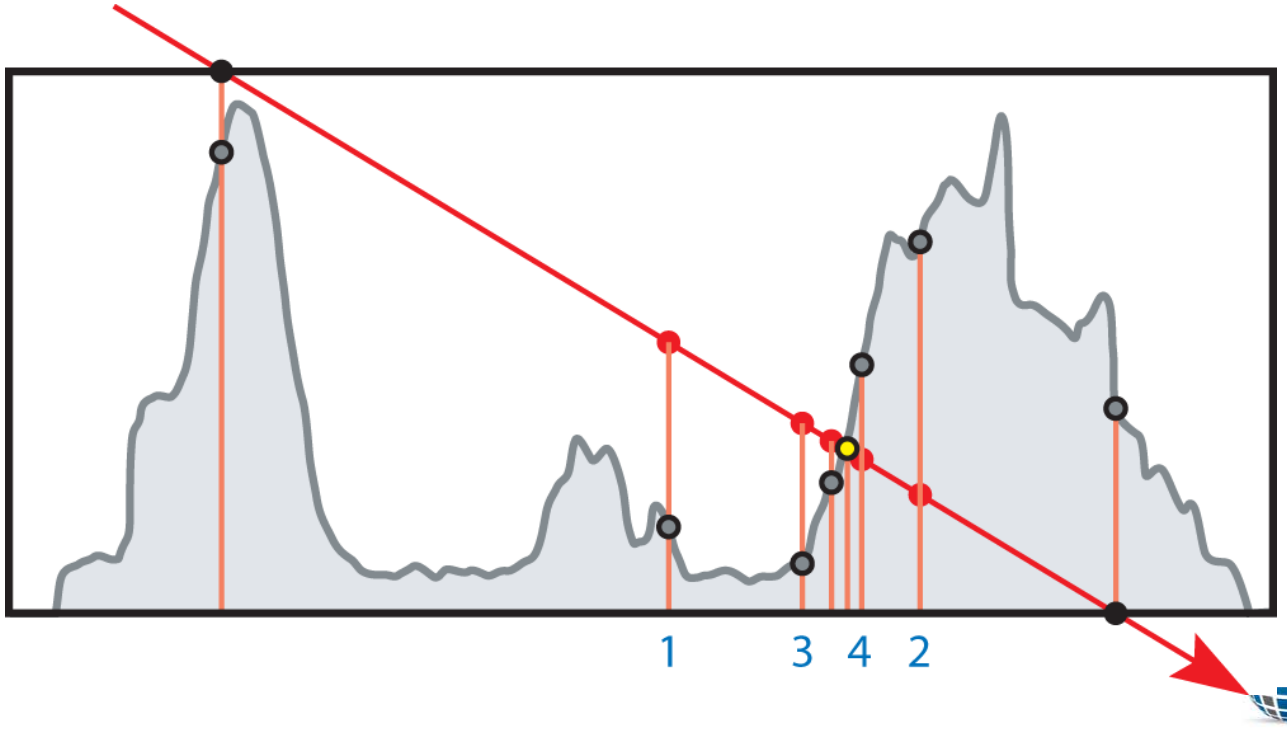

DELFT DATA SCIENCE

Intersection Algorithm

- Observation: Only works for one intersection!

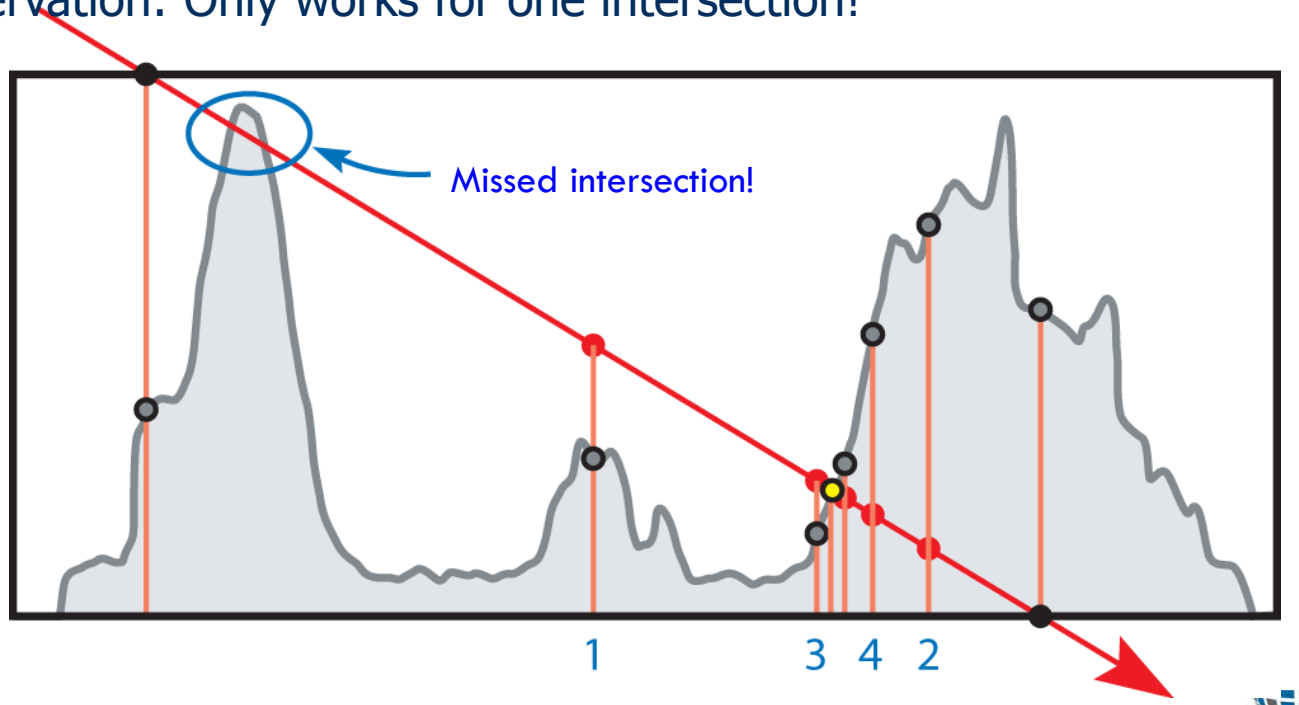

DELFT DATA SCIENCE

Traditional Height-Field Rendering

- Combine uniform steps with bisection method

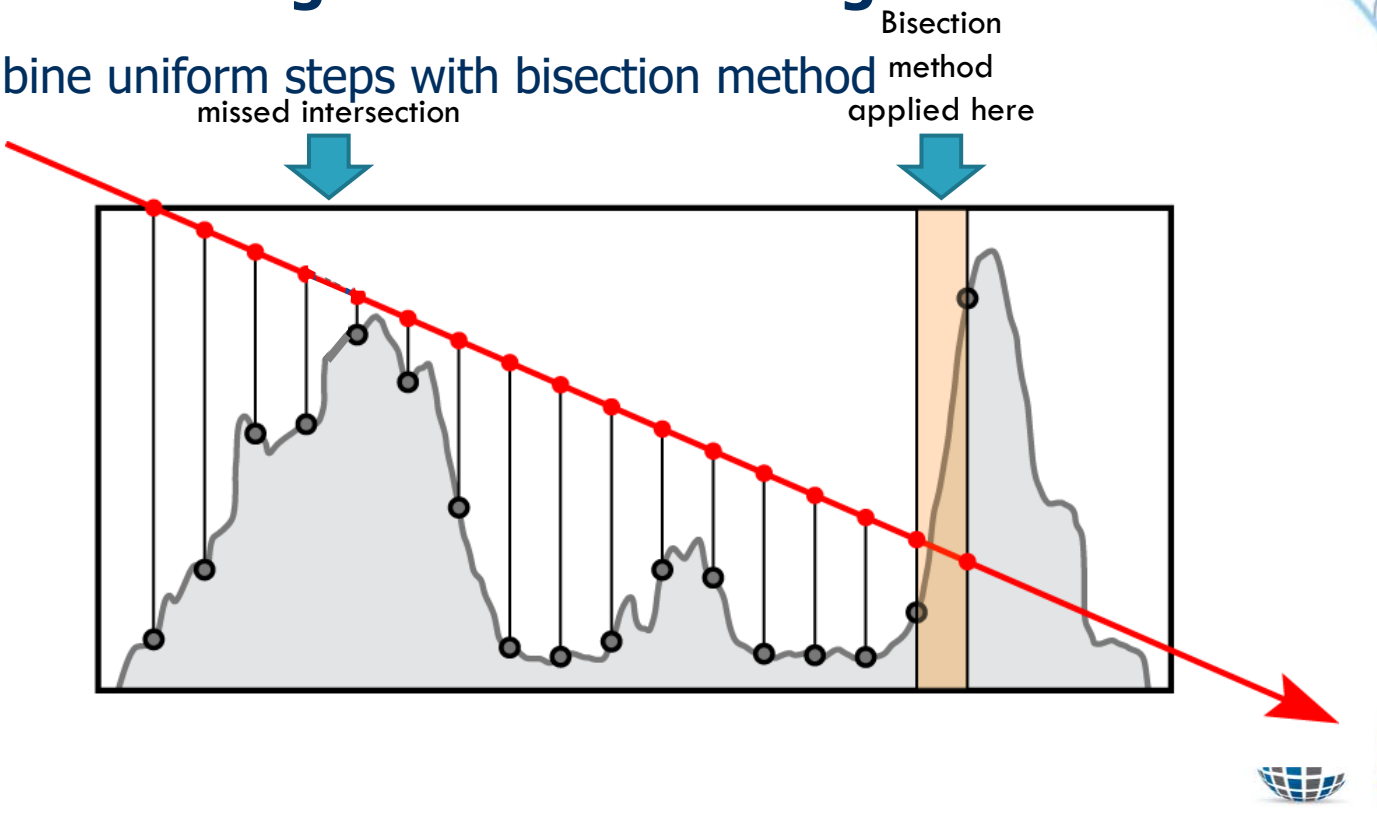

Previous Work

- Advanced Accelerations
  [Policarpo 05,07], [Tatarchuk 06]

100

steps

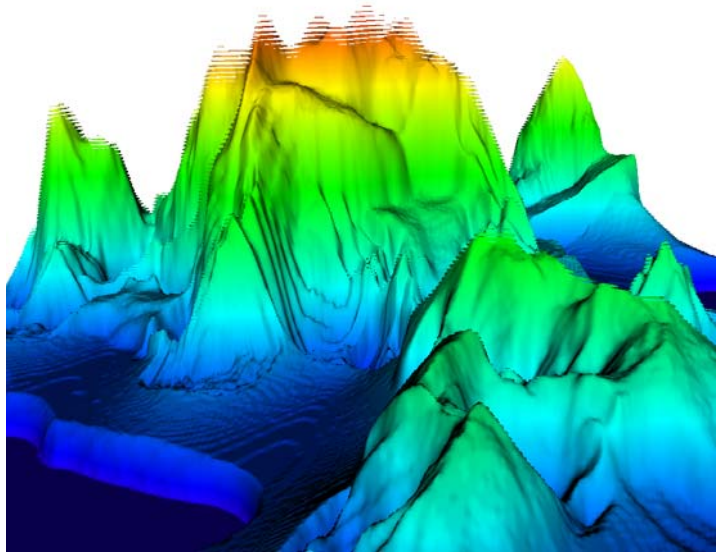

DELFT DATA SCIENCE

Previous Work

- Advanced Accelerations
  [Policarpo 05,07], [Tatarchuk 06]

500

steps

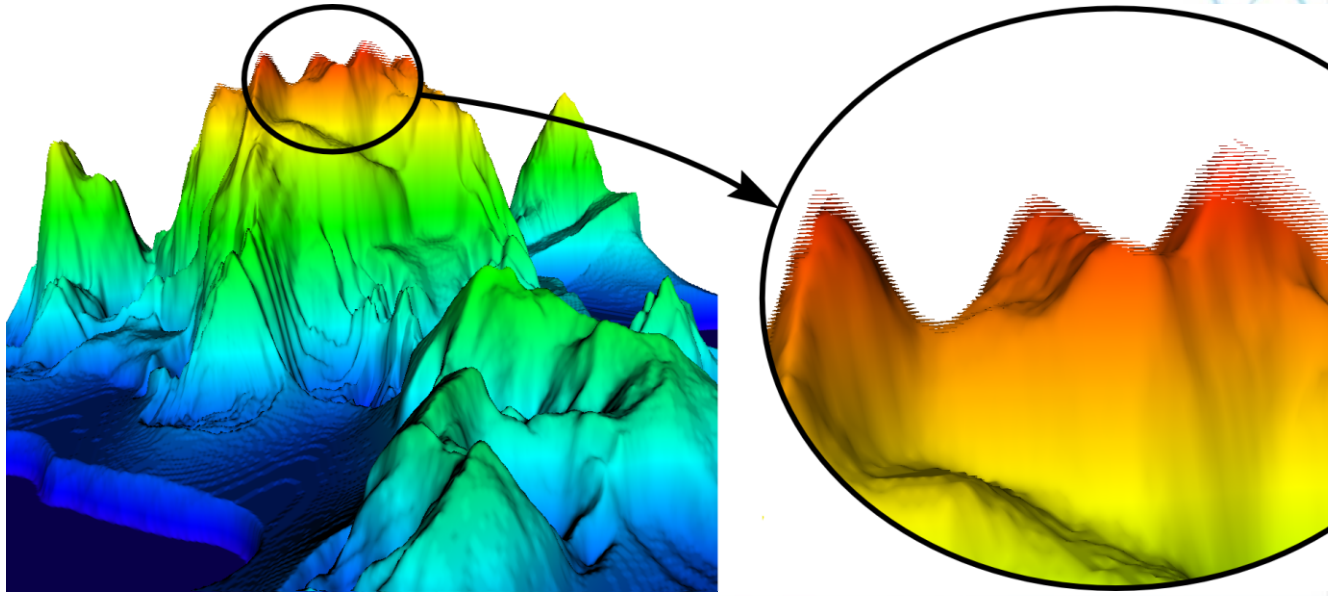

Marching Shapes – Non-collision distances

- Store distance to surface in 3D volume

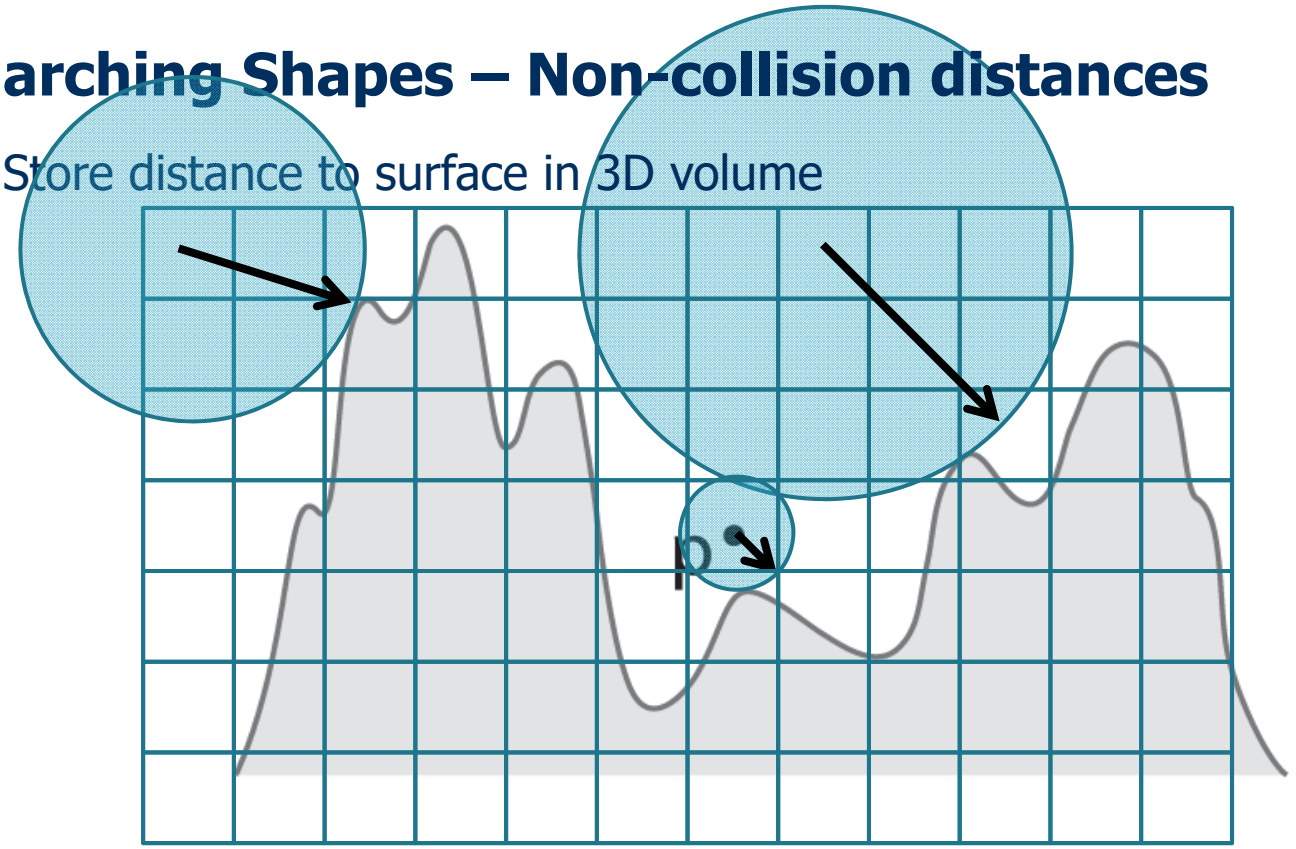

Marching Shapes – Non-collision distances

- Problems: Requires 3D storage
  No converge on surface

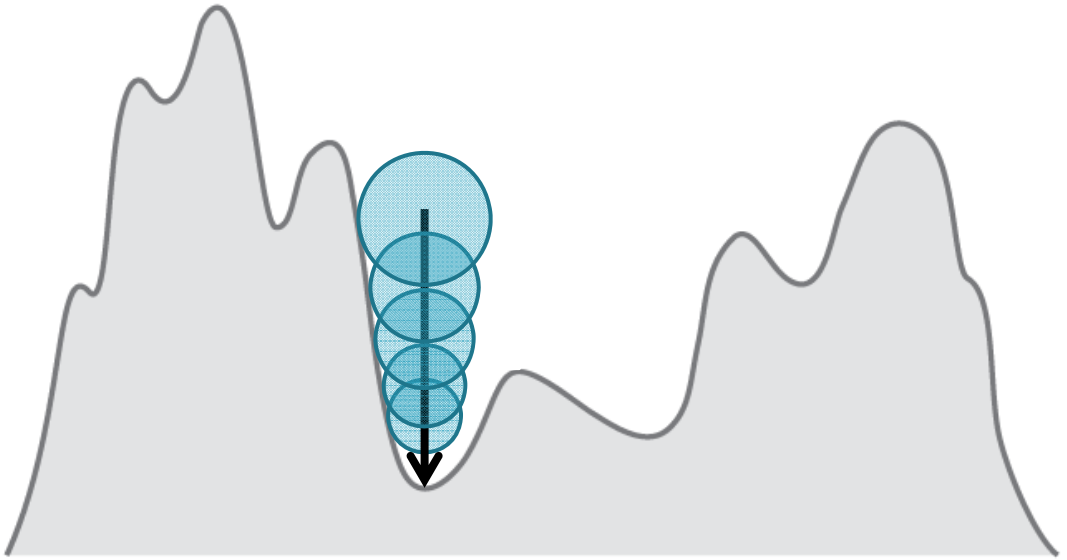

DELFT DATA SCIENCE

Marching Shapes

- Store distance to surface in 3D volume

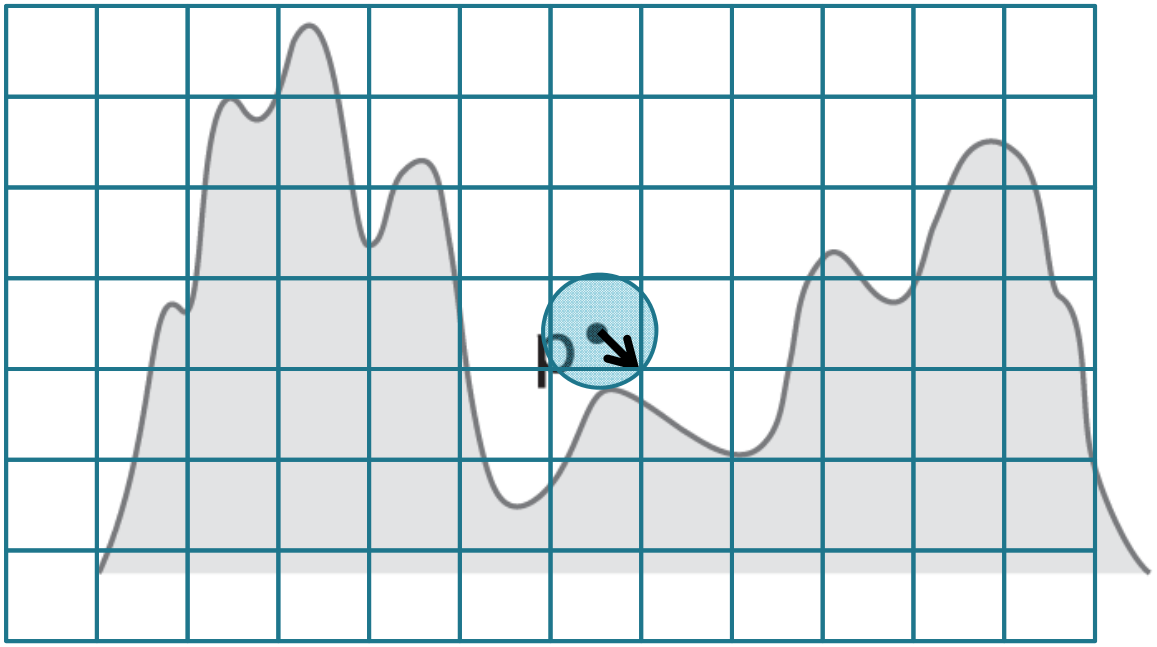

[Baboud, Eisemann, Seidel – IEEE TVCG 2012]

Safety Shapes

- Store distance to 2nd surface in 3D volume

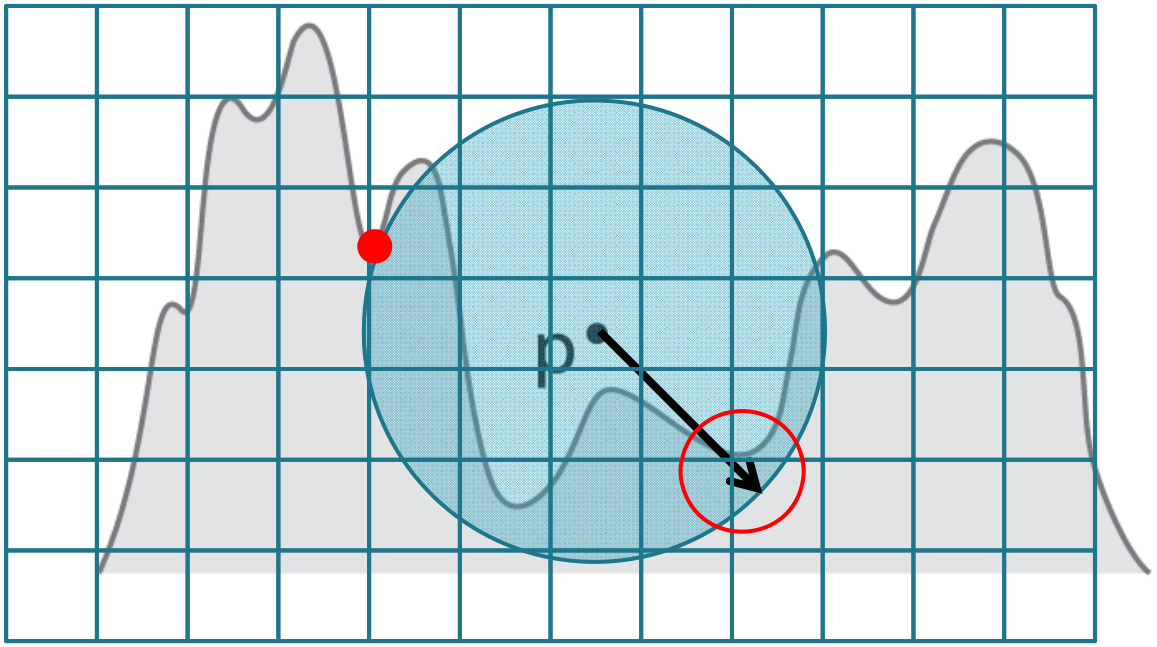

Start bisection method

[Baboud, Eisemann, Seidel – IEEE TVCG 2012]

Reduce Memory Cost

- Store 2D square in the height-field texture (only one additional value)

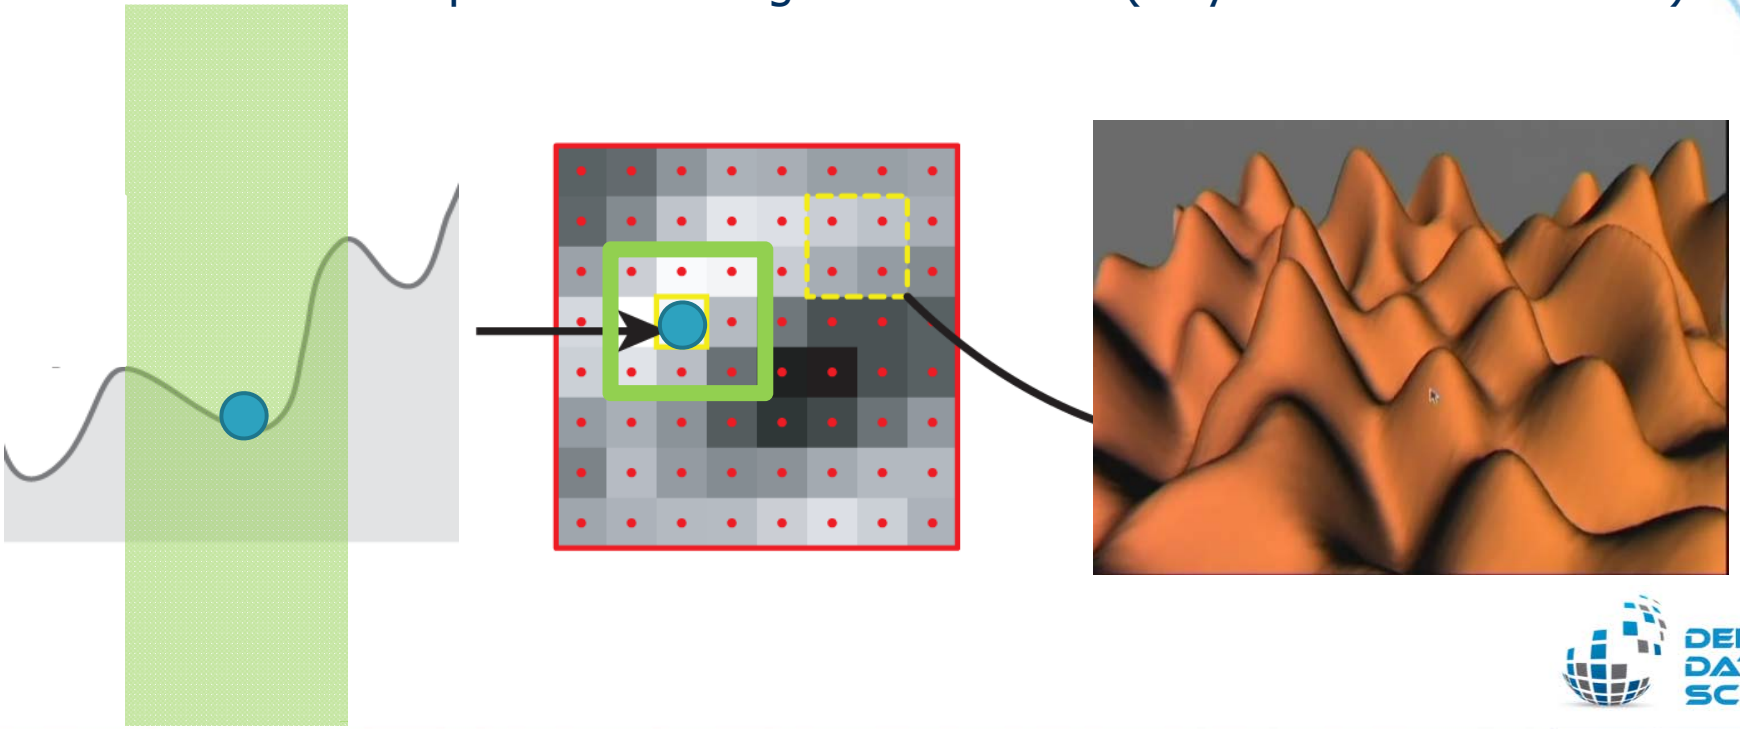

[Baboud, Eisemann, Seidel – IEEE TVCG 2012]

Safety Shapes – Bells and Whistles

- Fast Computation
- Novel Interpolation Method
- Improved Safety Shapes (Directional, Cylicone etc.)
- Highest Numerical Stability

[Baboud, Eisemann, Seidel – IEEE TVCG 2012]

Precomputed Safety Shapes

- Fast and accurate height-field rendering
- Arbitrary rays
- Low precomputation (3-4 orders of magnitude)

- Pure 2.5D
- All data has to reside in memory

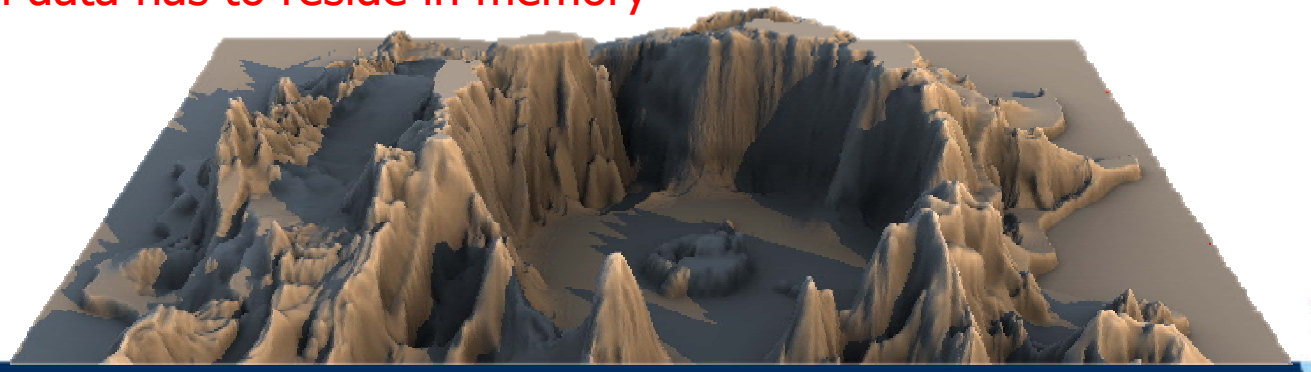

Questions?

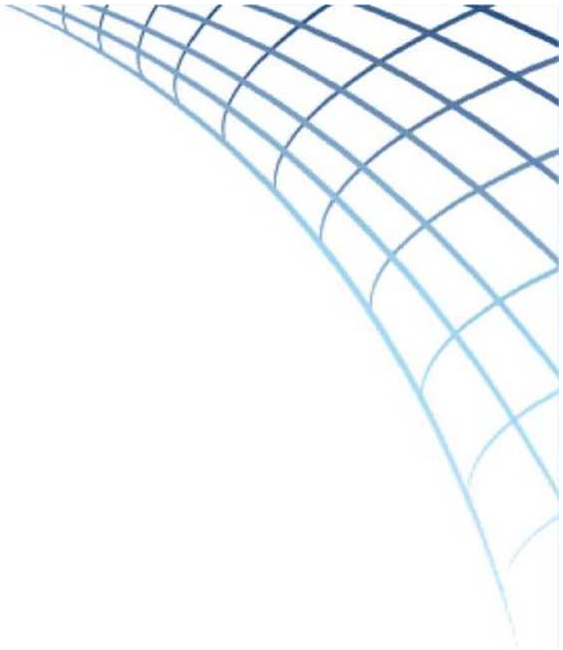

Large-scale Out-of-core Rendering

- Create Hierarchy (tile area and create multiple resolution height fields)
- Choose resolution according to viewer
- Cull nodes outside the frustum

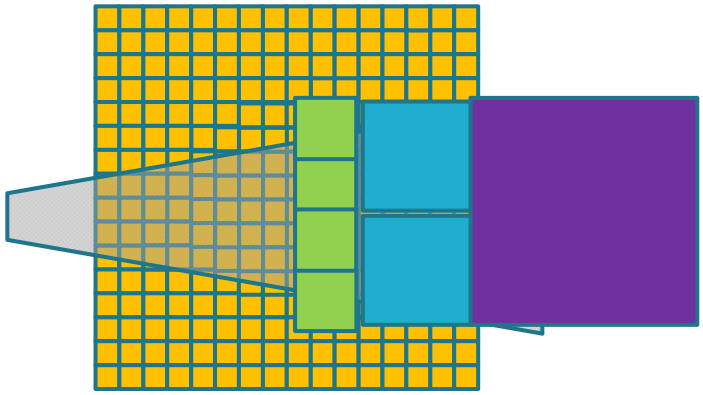

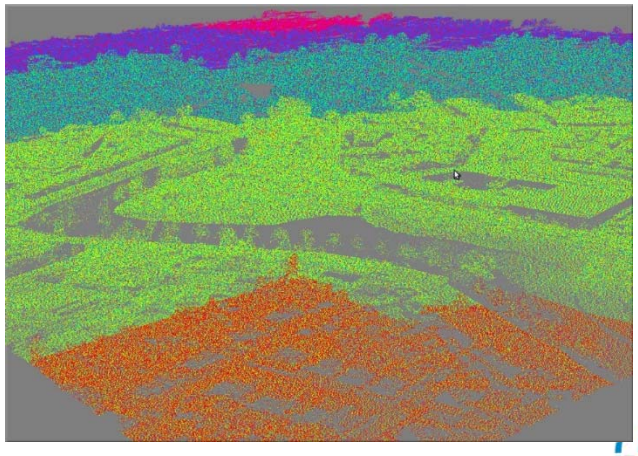

DELFT DATA SCIENCE

Large-Scale Data

- Integration of street-view scans into the AHN2 dataset

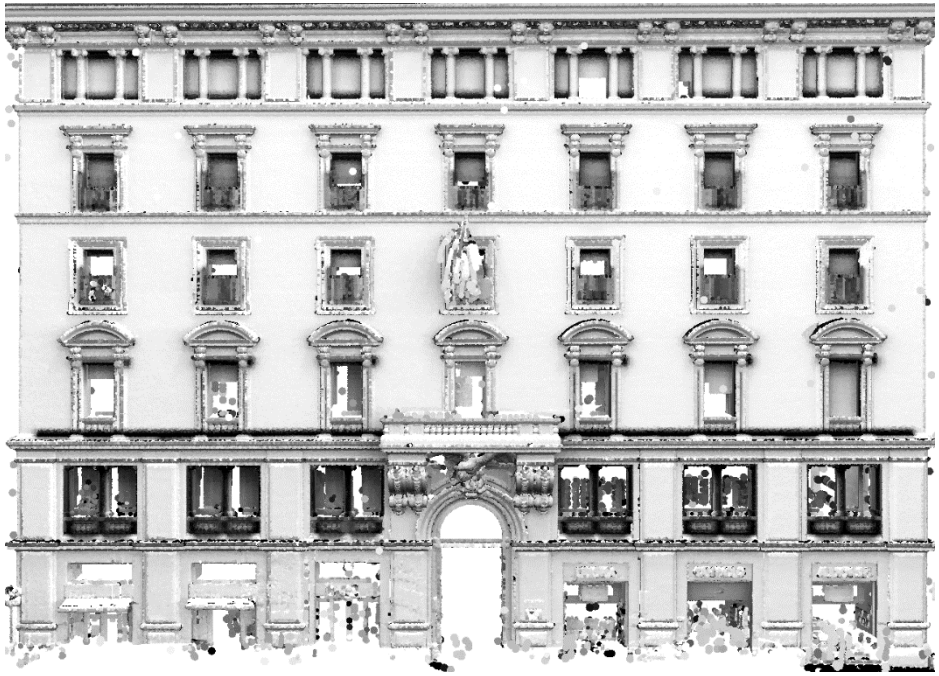

DELFT DATA SCIENCE

Add structure…

- Voxel Representations

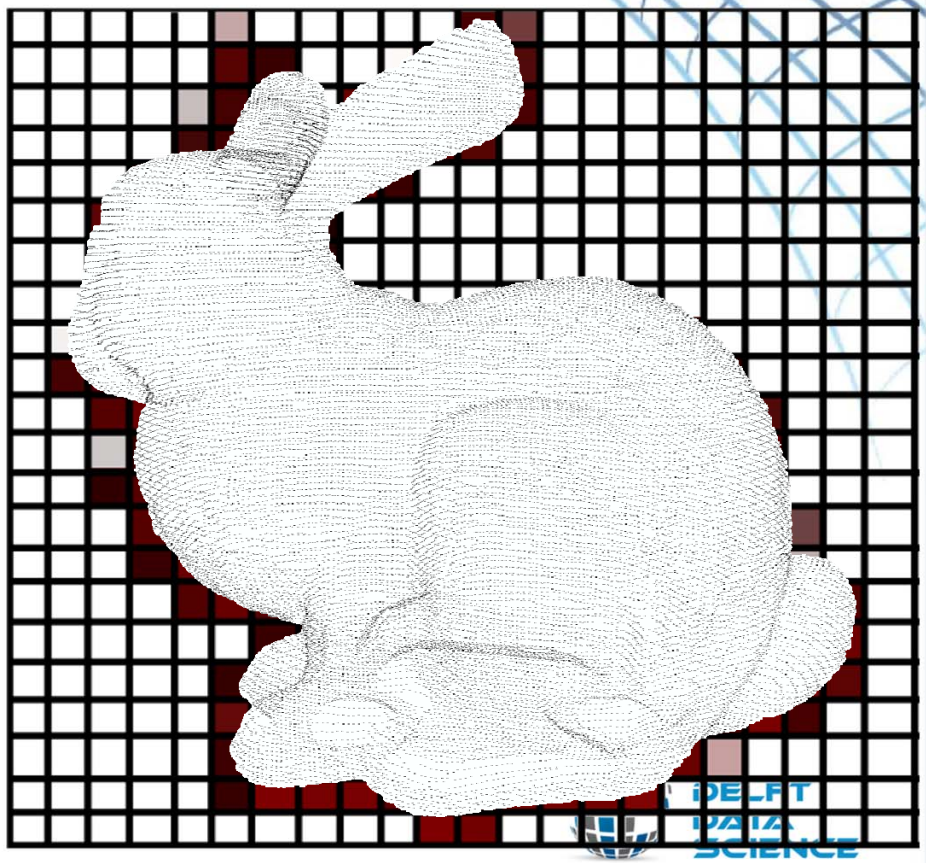

Voxels - Excursion

- Have many applications
- Allow for random-access queries
- Fast algorithms to voxelize objects in real-time
  [Eisemann & Decoret i3D2006], [Eisemann & Decoret GI2008]
  Back then: 1024^3 binary voxelization at 100 fps for 300K triangles (NvidiaG80)

[Eisemann & Decoret i3D2006], [Eisemann & Decoret GI2008]

Voxels

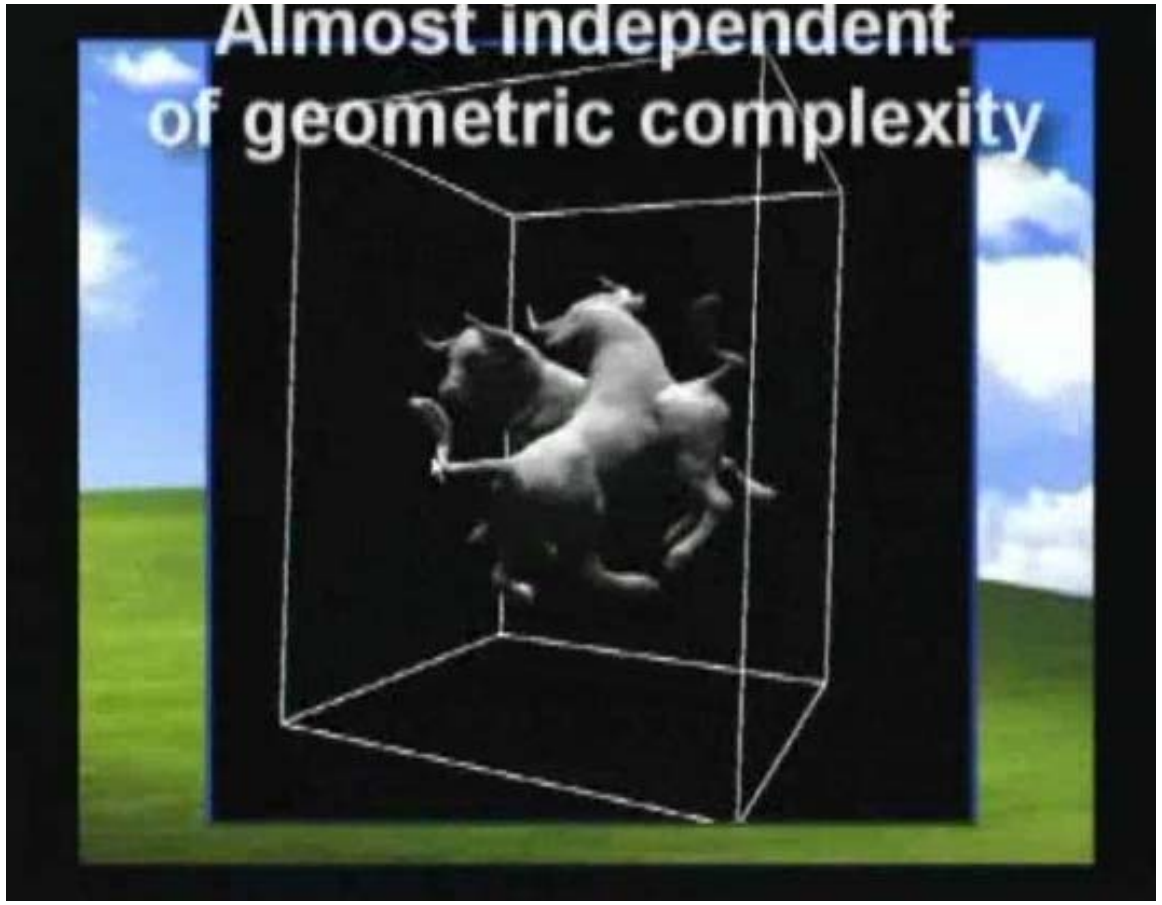

DELFT DATA SCIENCE

[Eisemann & Decoret i3D2006], [Eisemann & Decoret GI2008]

Voxels

Voxels

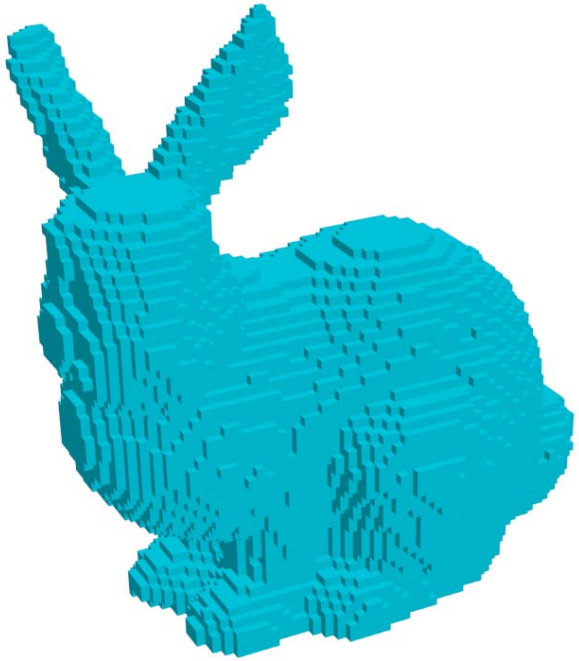

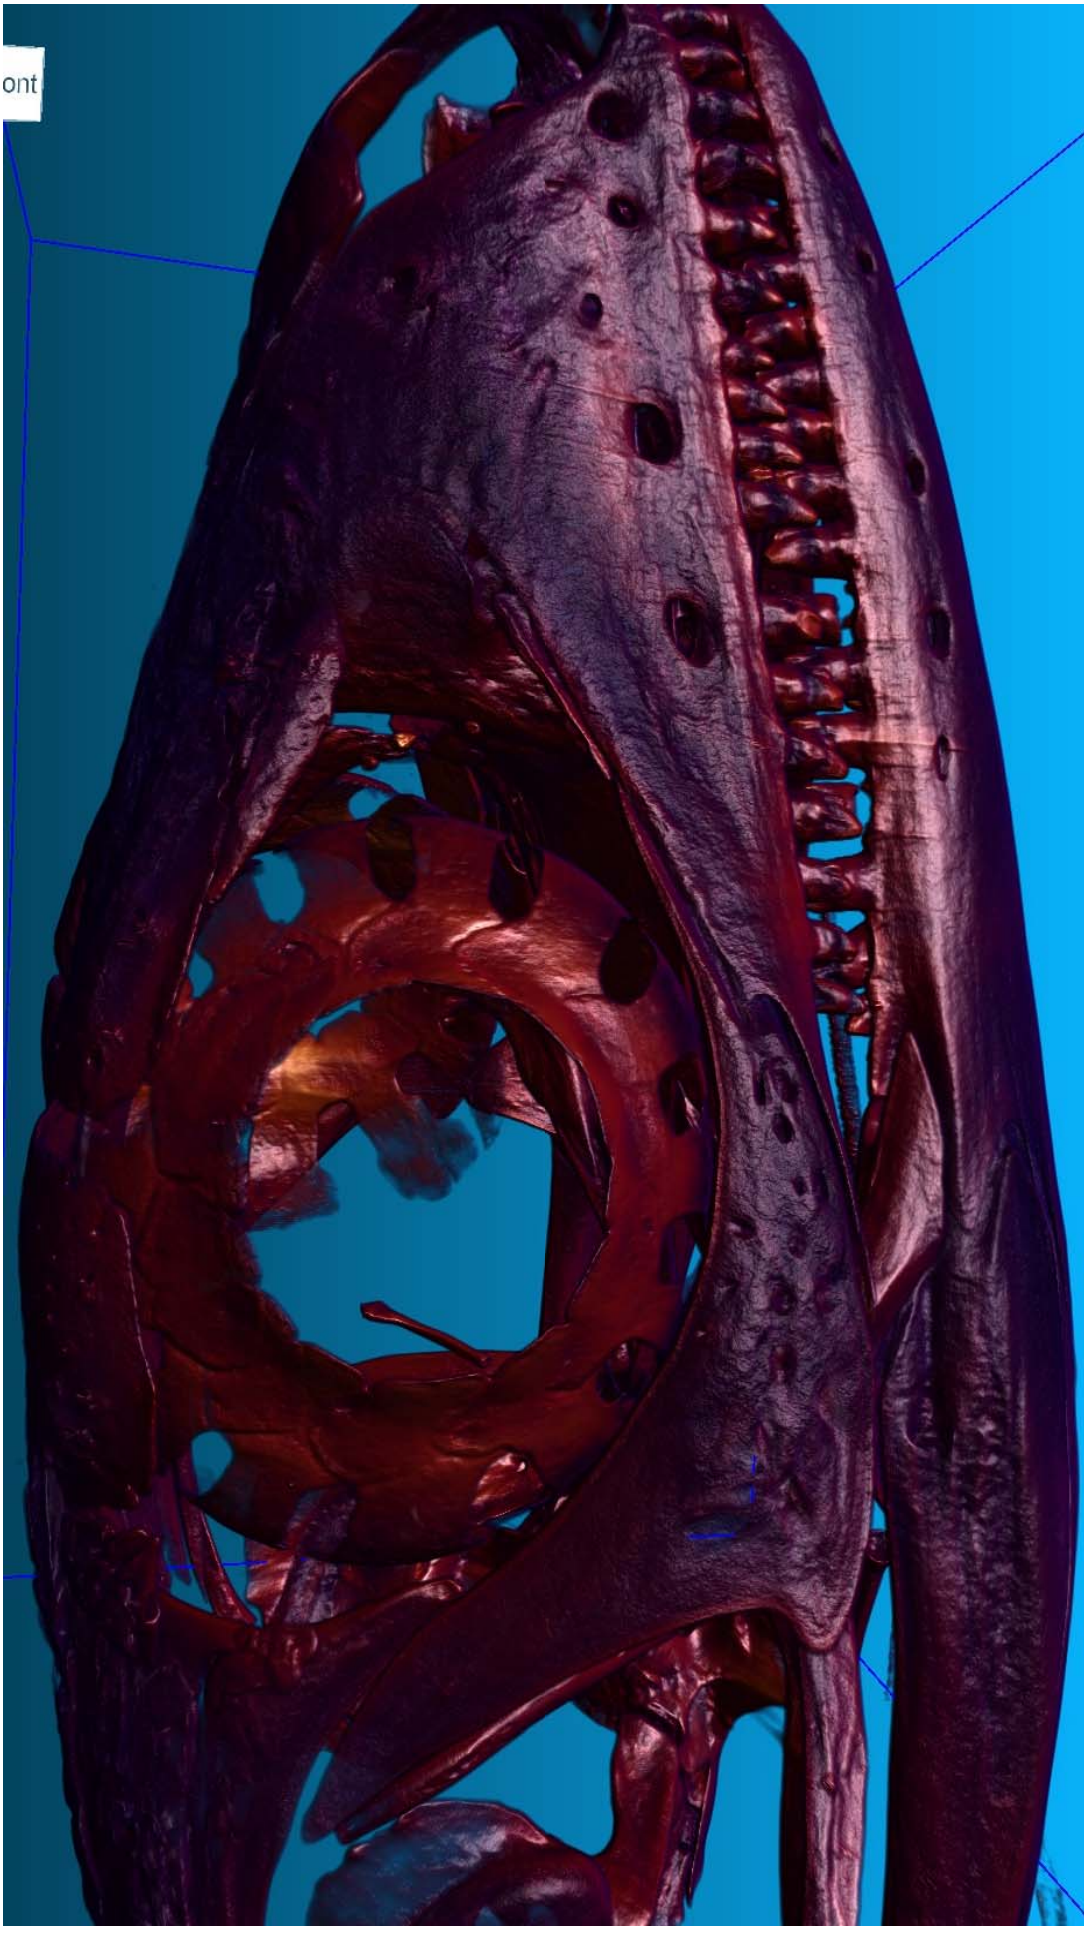
ont

Voxel Representations

- **Main problem:**
  - Memory is a key issue!
    - E.g. 2048^3 x RGBA = **32 GB**

DELFT DATA SCIENCE

How to deal with large-scale scenes?

Level of detail, visibility tests, compresssion…

Today: only a glimpse of what is done…

Ray-tracing methods [Wald et al. 2004]

- During ray traversal record missing data and upload

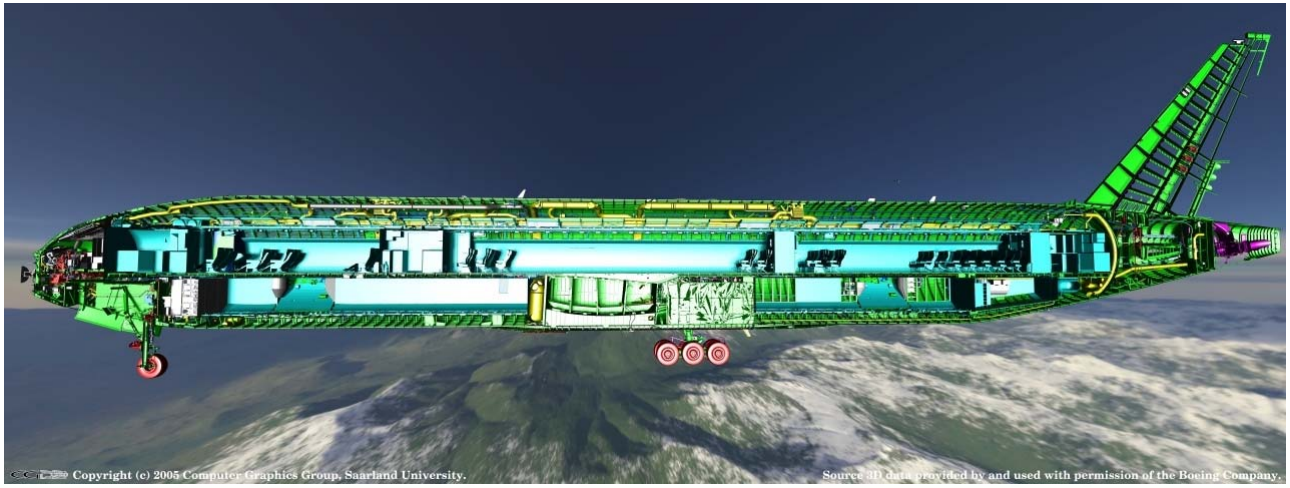

DELFT DATA SCIENCE

Modern GPU instances

- Hybrid CPU/GPU approach [Gobetti et al. 2008]
- Full GPU approach [Crassin et al. 2009]

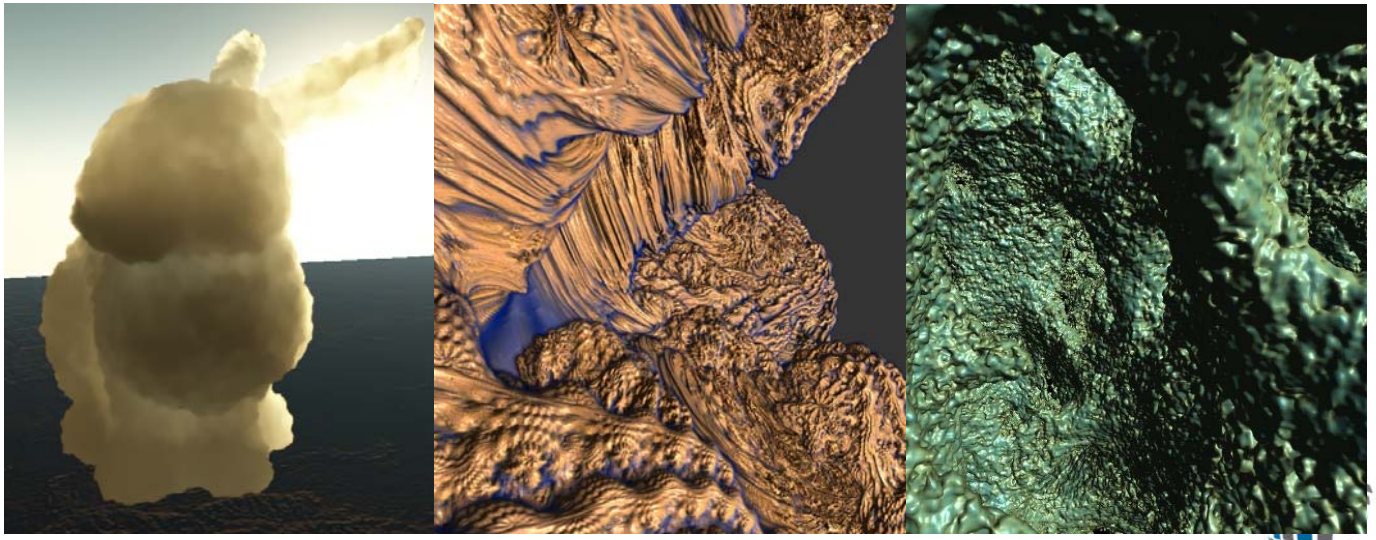

Ways to deal with large data

- Constant valued areas
  (also empty)

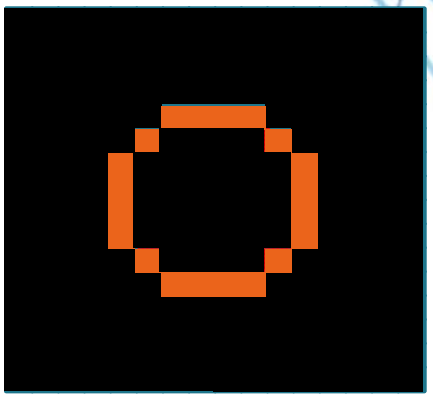

Ways to deal with large data

- Constant valued areas
  (also empty)

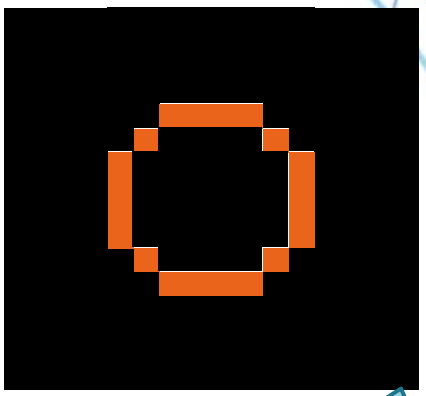

- View-dependence
  - Visibility

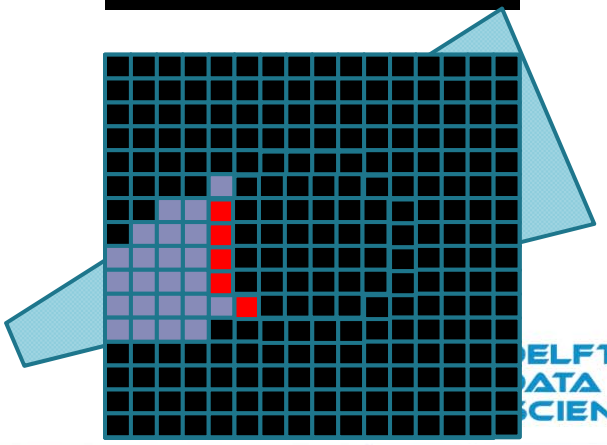

Ways to deal with large data

- Constant valued areas
  (also empty)

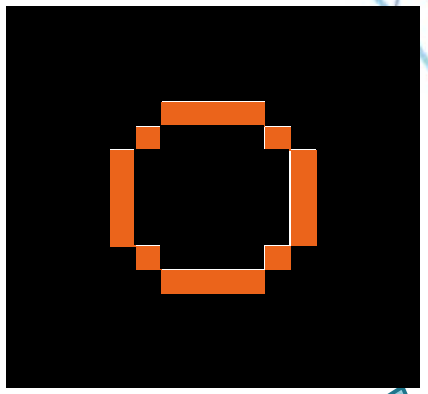

- View-dependence
  - Visibility

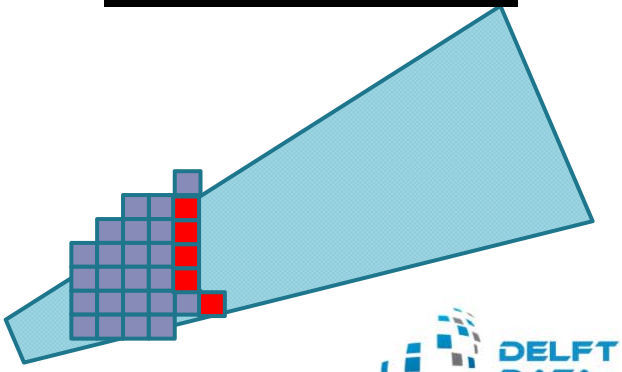

DELFT DATA SCIENCE

Ways to deal with large data

- Constant valued areas
  (also empty)

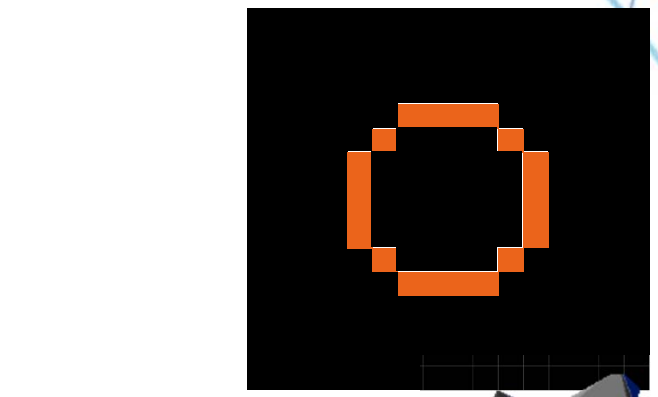

- View-dependence
  - Visibility
  - Level of Detail

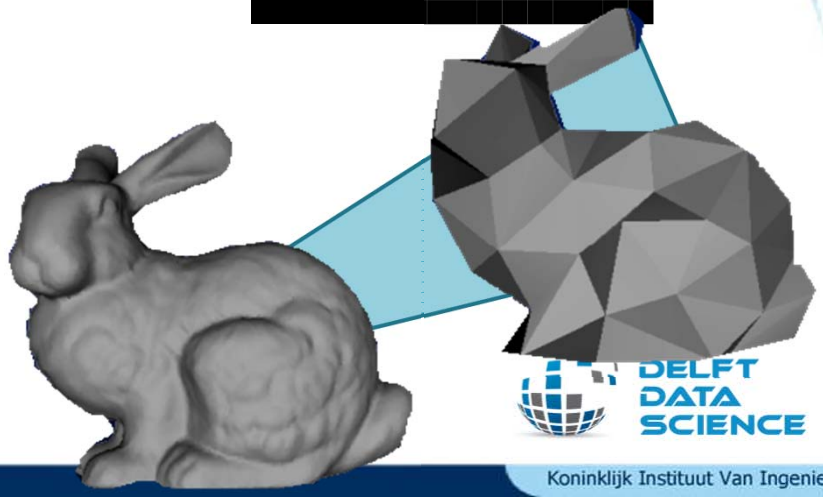

Ways to deal with large data

- Constant valued areas
  (also empty)

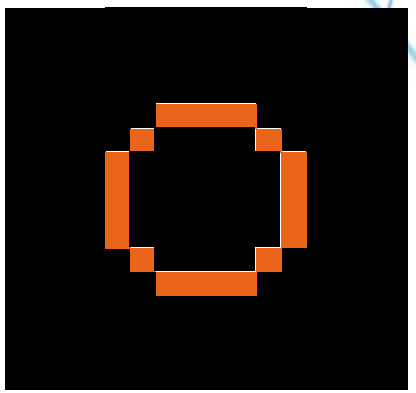

- View-dependence
  - Visibility
  - Level of Detail

DELFT DATA SCIENCE

[Crassin, Neyret, Lefebvre, Eisemann – I3D 2009]
[Crassin, Neyret, Lefebvre, Eisemann – GPUPro2010]

Free-Viewpoint Details

Ways to deal with large data

- How to exploit these observations in practice?

  Need:
  - Data structure allowing for easy adaptation
  - Memory management

[Crassin, Lefebvre, Neyret, Eisemann - i3D 2009]
[Crassin, Neyret, Lefebvre, Eisemann - GPUPro2010]

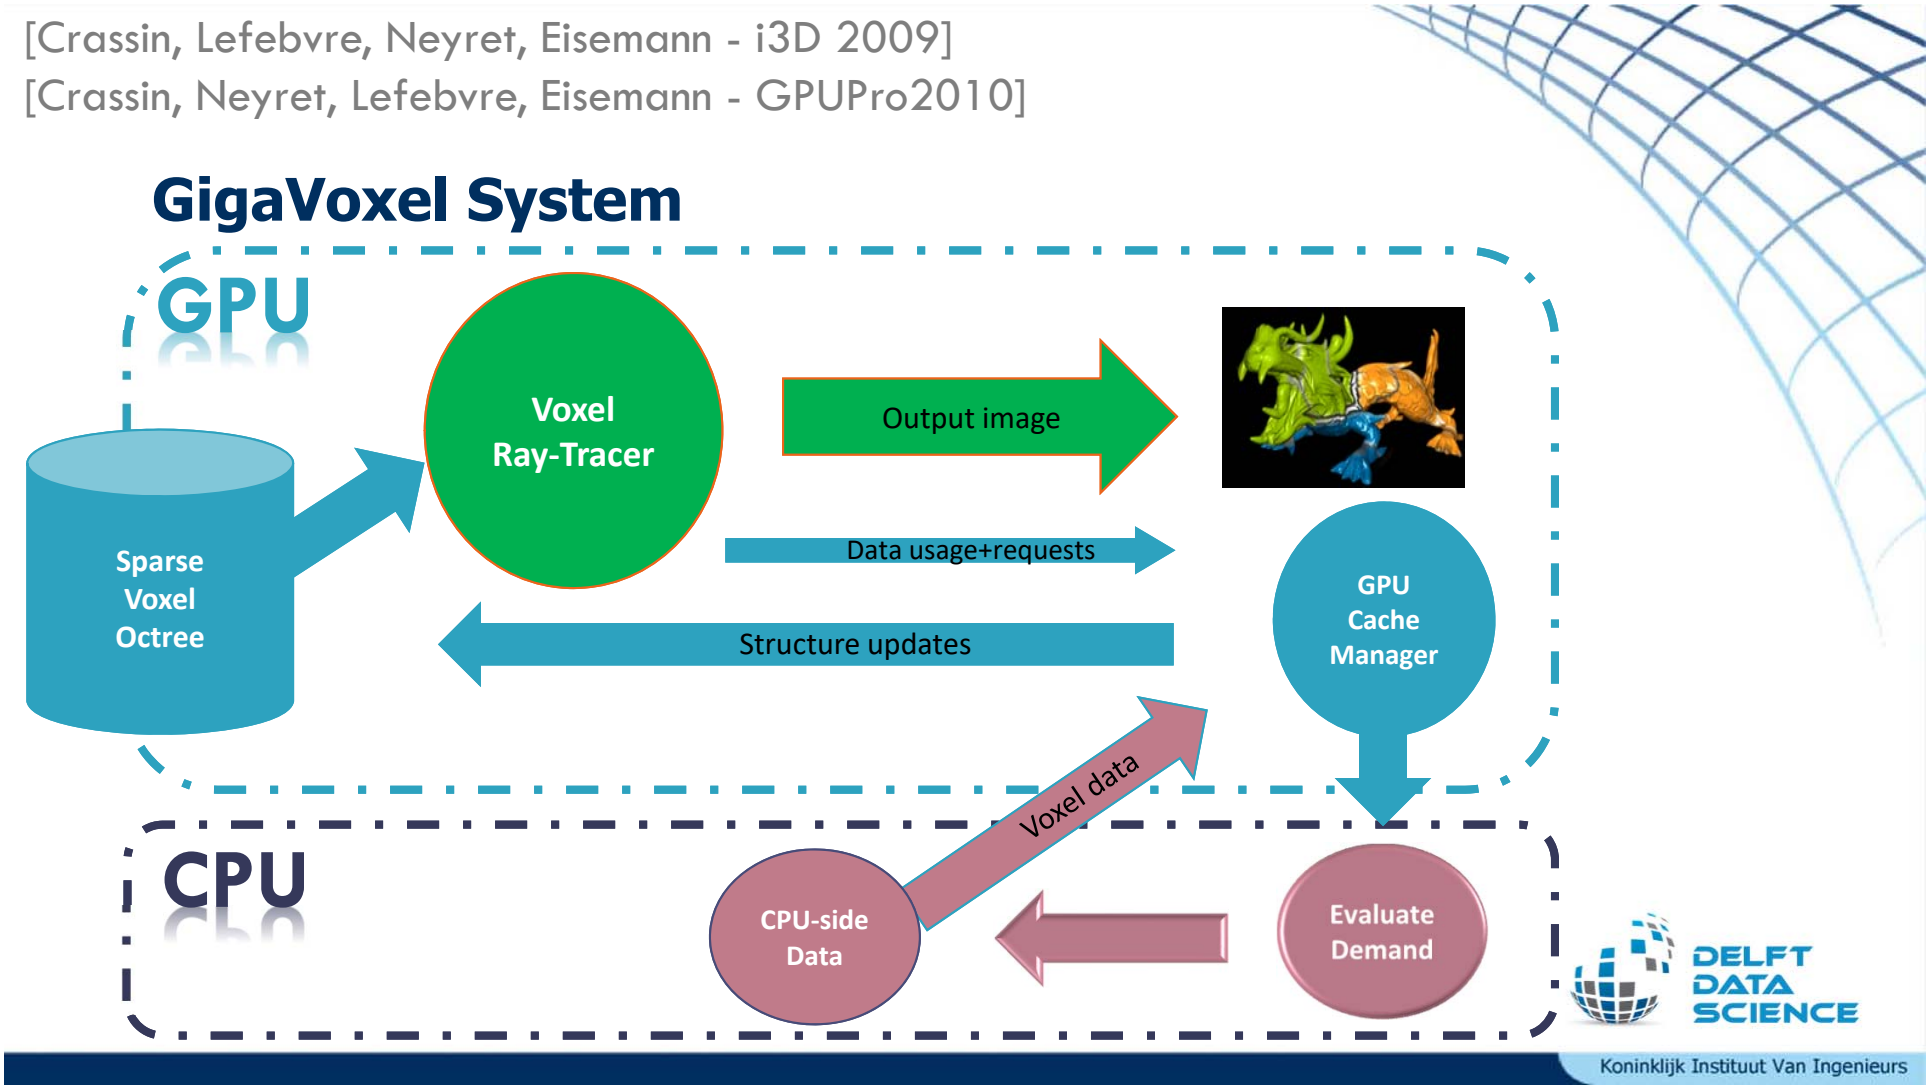

Octree of Small Volumes (Bricks)

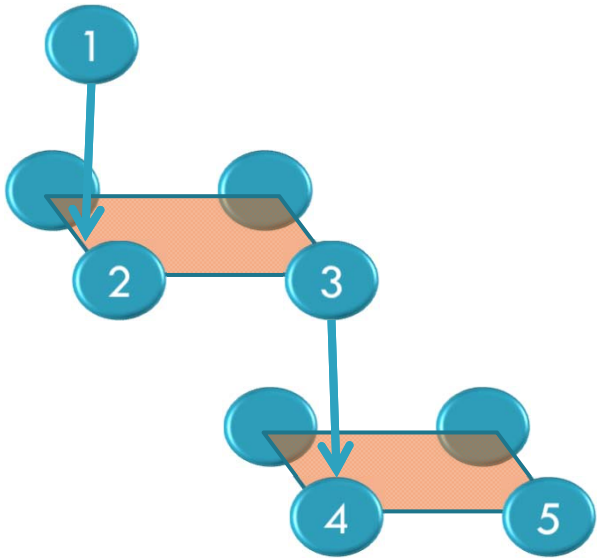

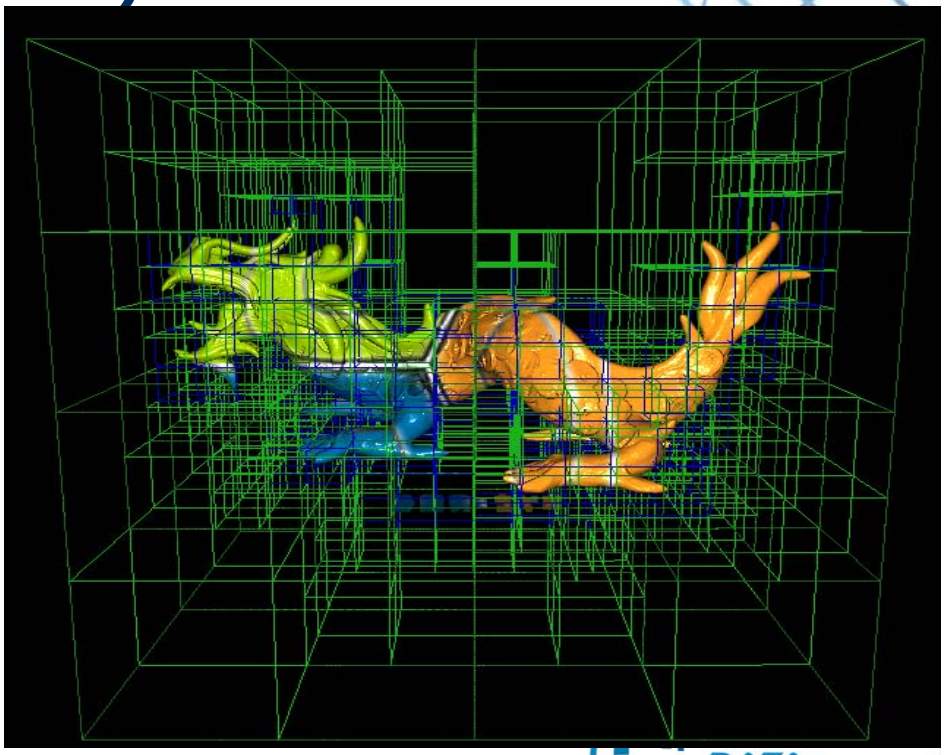

Octree of Voxel Bricks

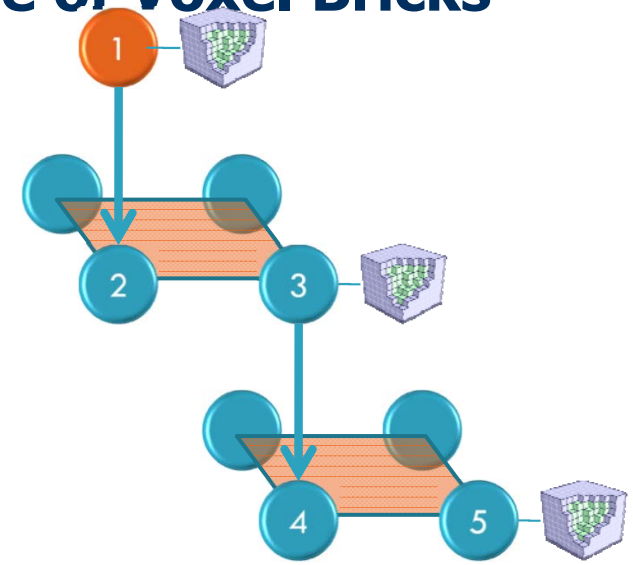

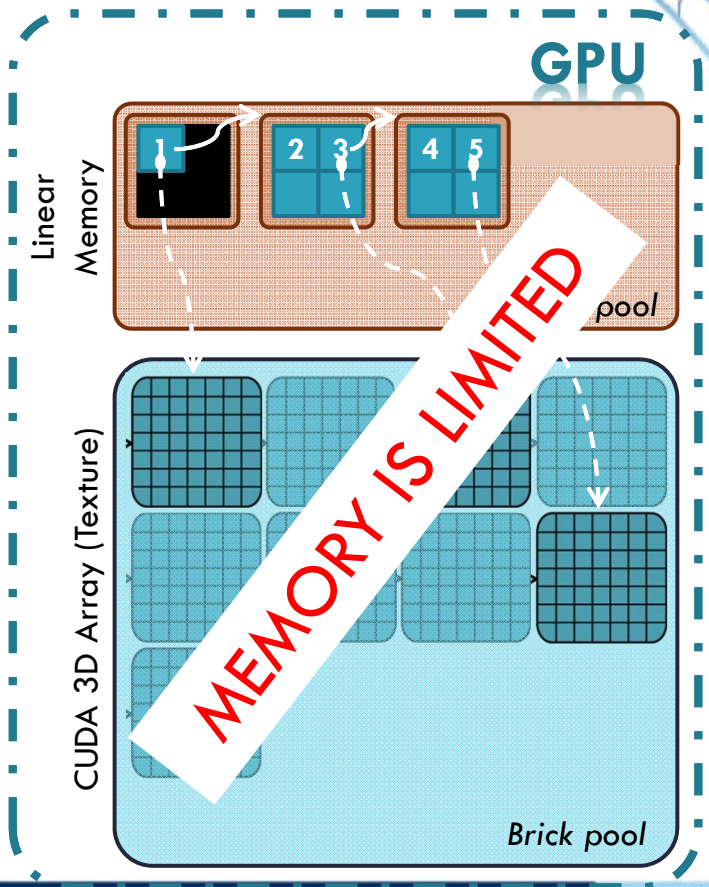

Volume Ray-Casting

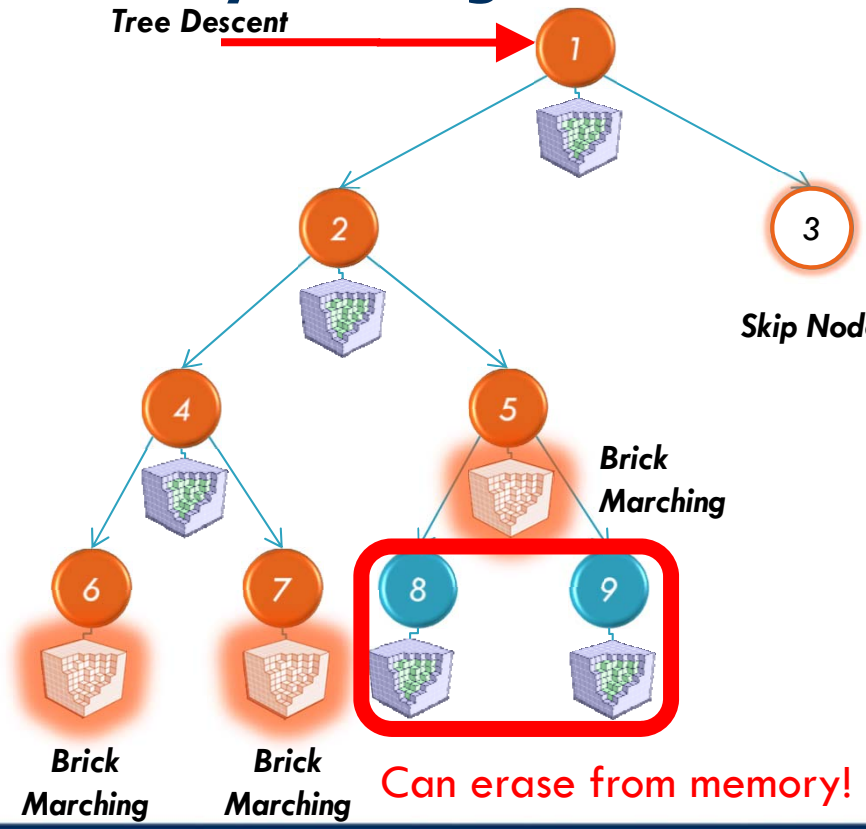

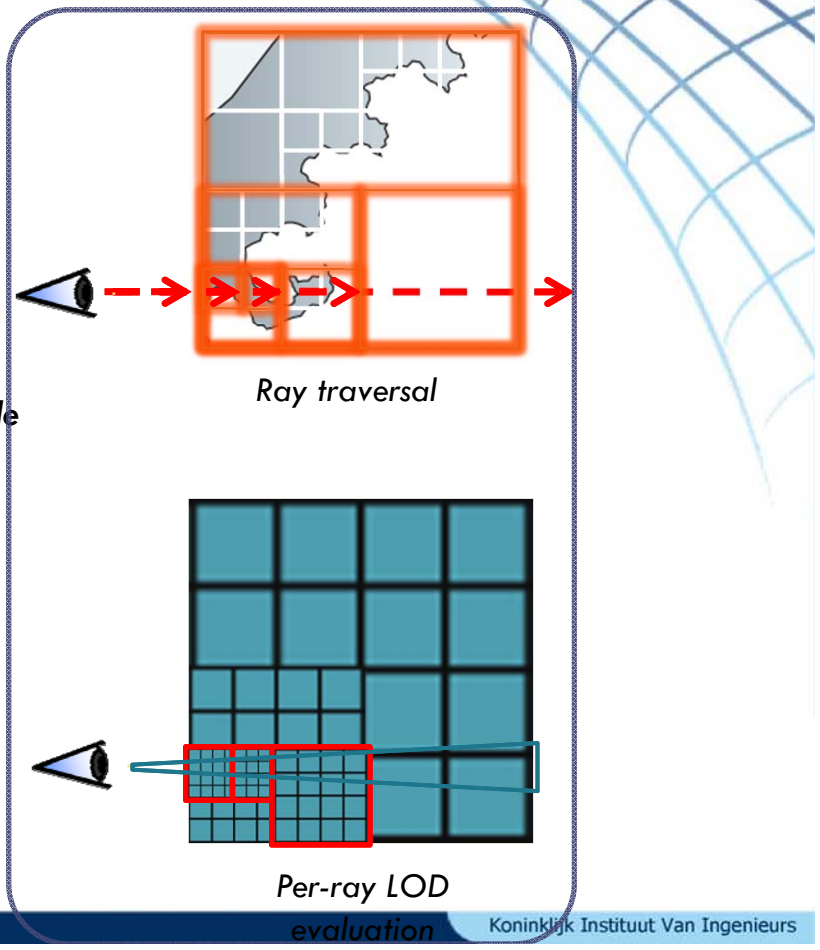

Octree of Small Volumes (Bricks)

- All elements have the same size in memory!
- Makes exchanges easy

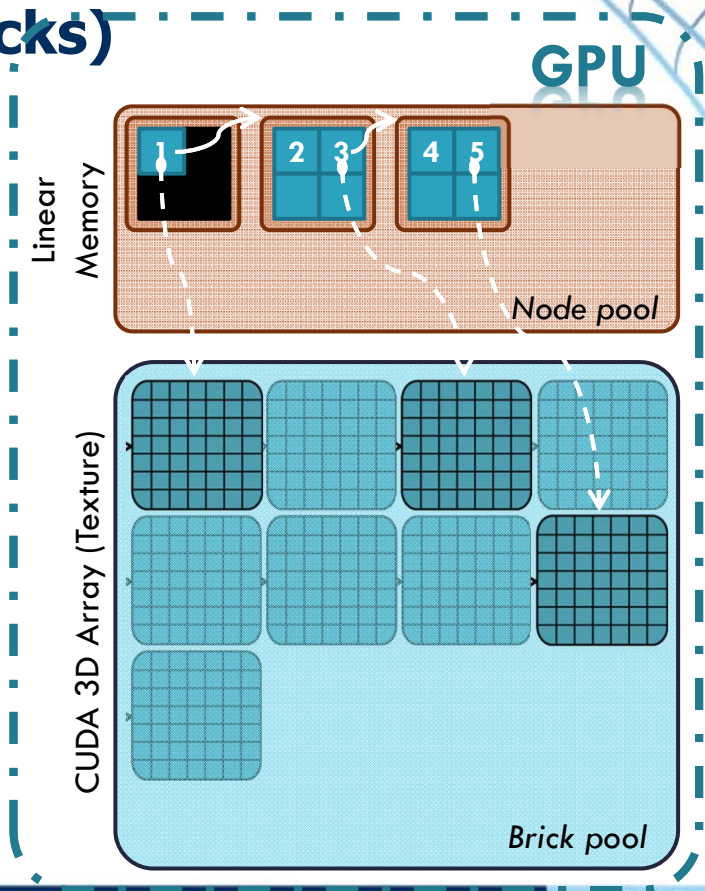

“Rendering Queries” – Ray Queries

- GPU Least-Recently-Used Cache
  - Track element usage
  - Maintain list with least-used element in front

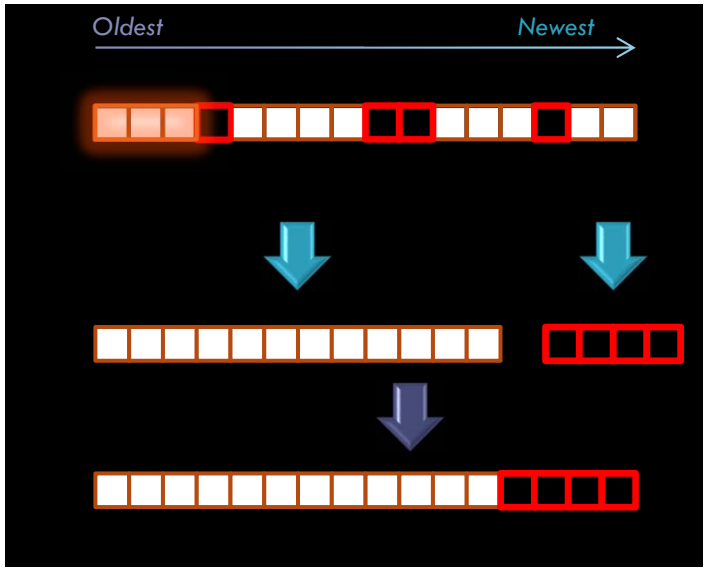

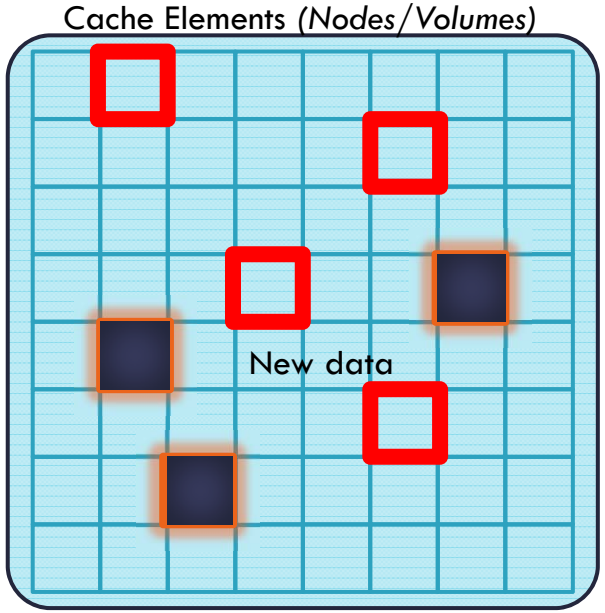

Node/Brick Pool

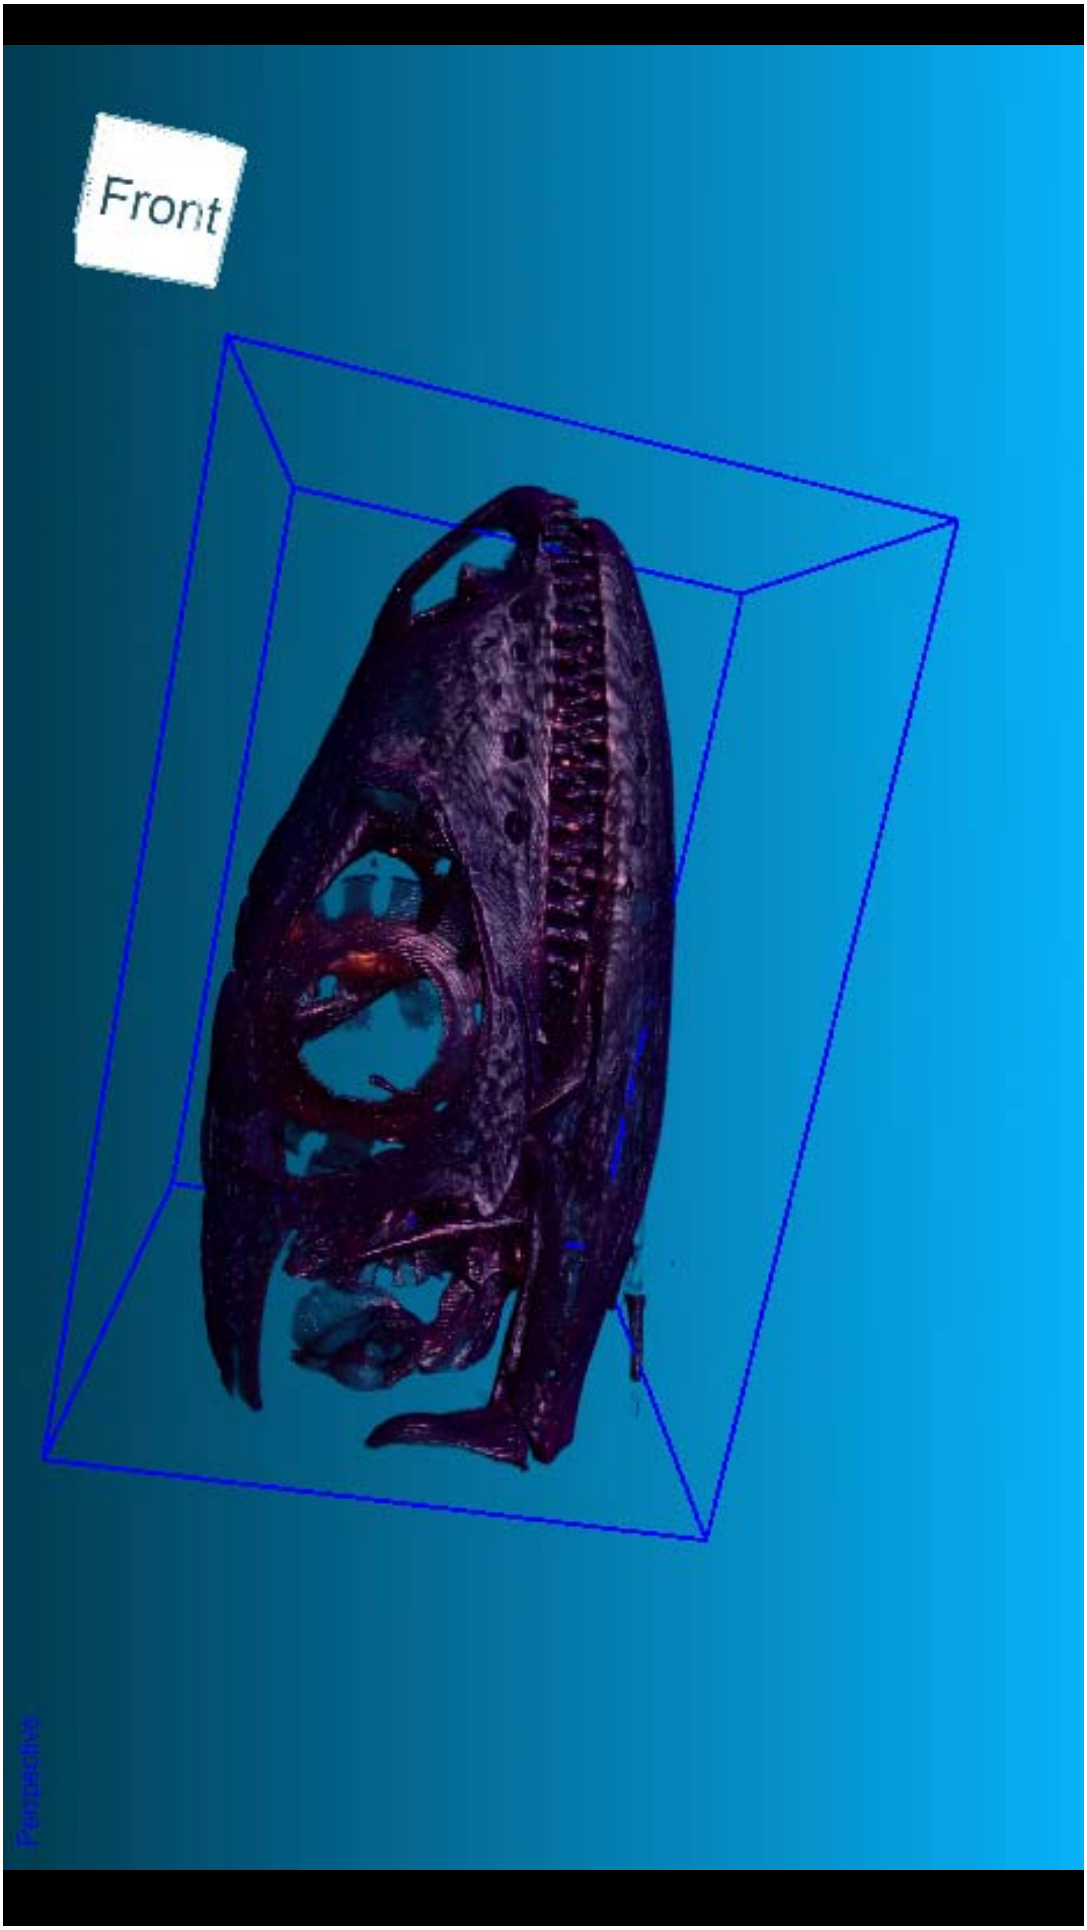
Front
Paraview

[Crassin, Lefebvre, Neyret, Eisemann - i3D 2009]
[Crassin, Neyret, Lefebvre, Eisemann - GPUPro2010]

GigaVoxels

- Interactive Out-of-core Exploration of Volume Data

- Fully GPU-oriented
  - Rendering Algorithm
  - Data Representation

Major remaining challenges:

- Bandwidth to GPU
- Storage (volumes grow quickly)

Questions?

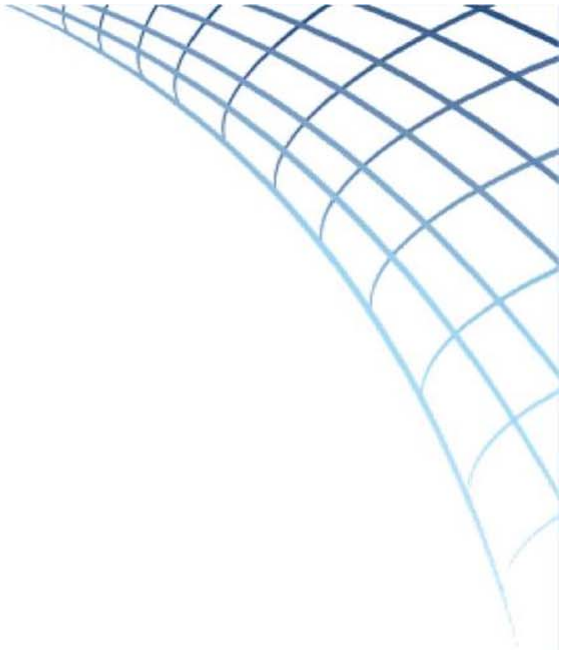

Large-Scale Voxels

Major remaining challenges for less performant systems:

- Bandwidth to GPU
- Storage (volumes grow quickly)

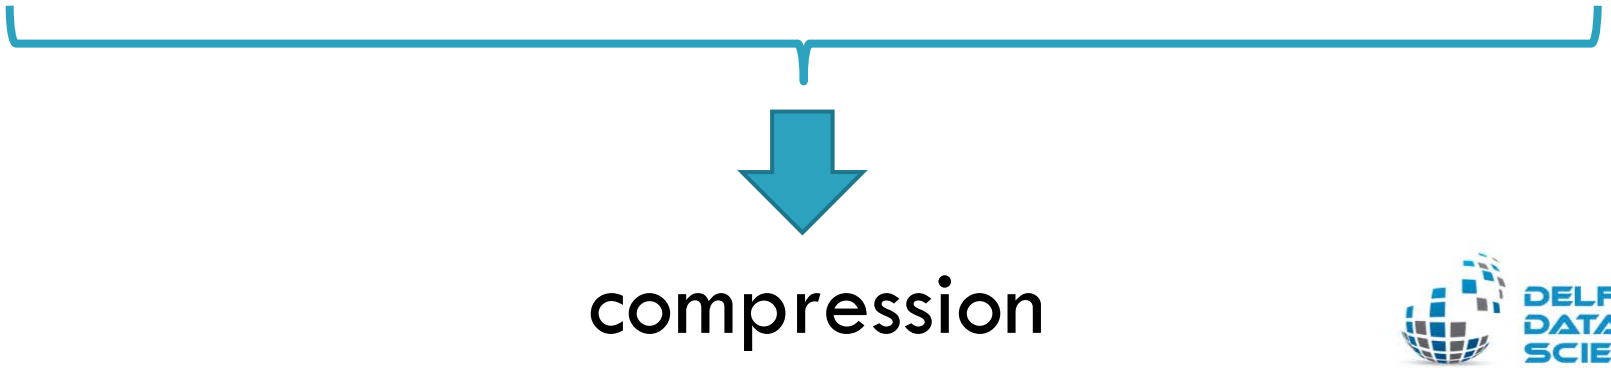

DELFT DATA SCIENCE

[Dado, Kol, Thiery, Bauszat, Eisemann – Eurographics 2016]

2^51 voxels = 2,251,799,813,685,248 Voxels

Colored DAGs

[Kaempe et al. – SIGGRAPH 2013]

Sparse Voxel Directional Acyclic Graphs

Binary Data Example

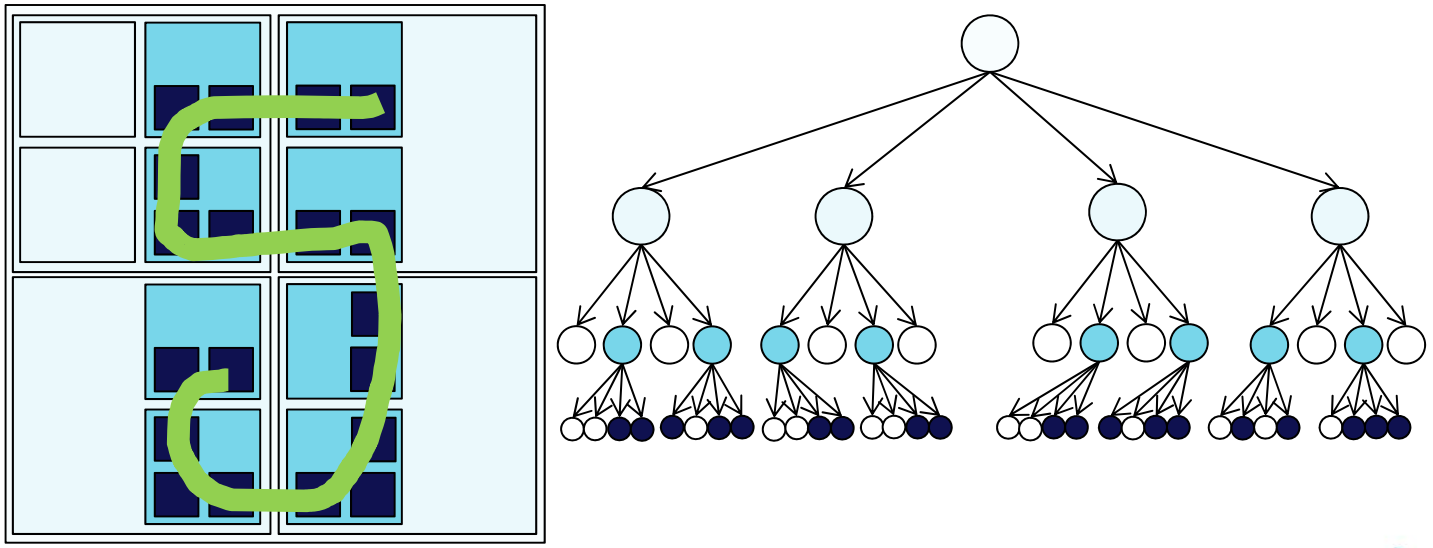

[Kaempe et al. – SIGGRAPH 2013]

Sparse Voxel Directional Acyclic Graphs

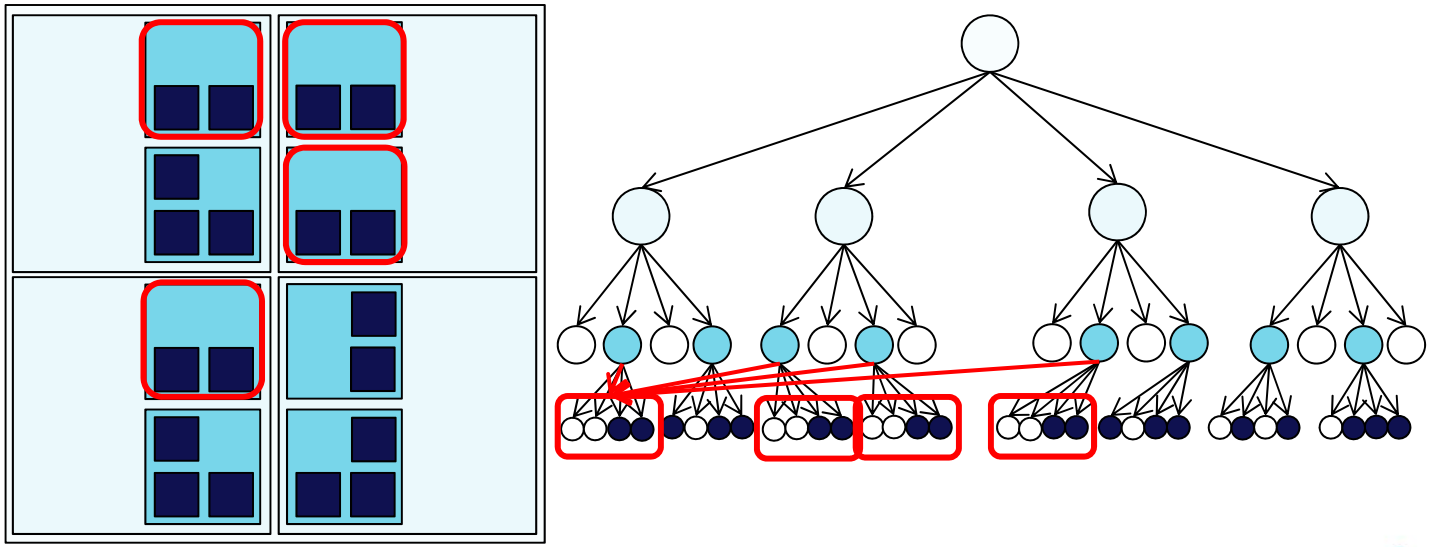

[Kaempe et al. – SIGGRAPH 2013]

Sparse Voxel Directional Acyclic Graphs

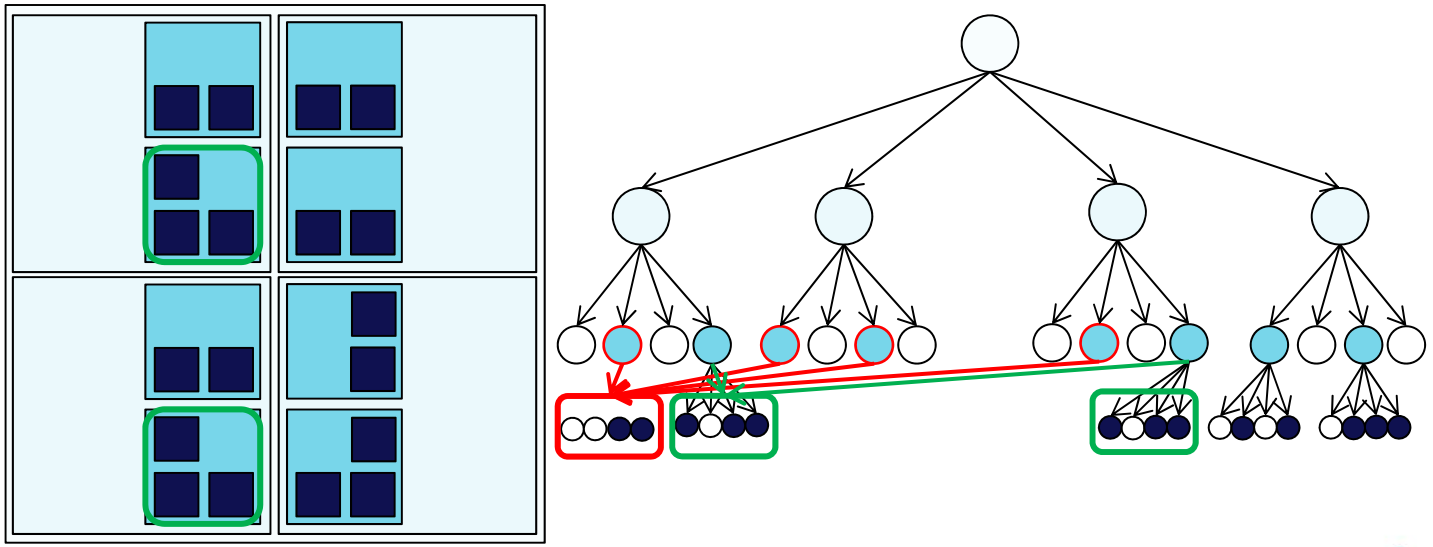

[Kaempe et al. – SIGGRAPH 2013]

Sparse Voxel Directional Acyclic Graphs

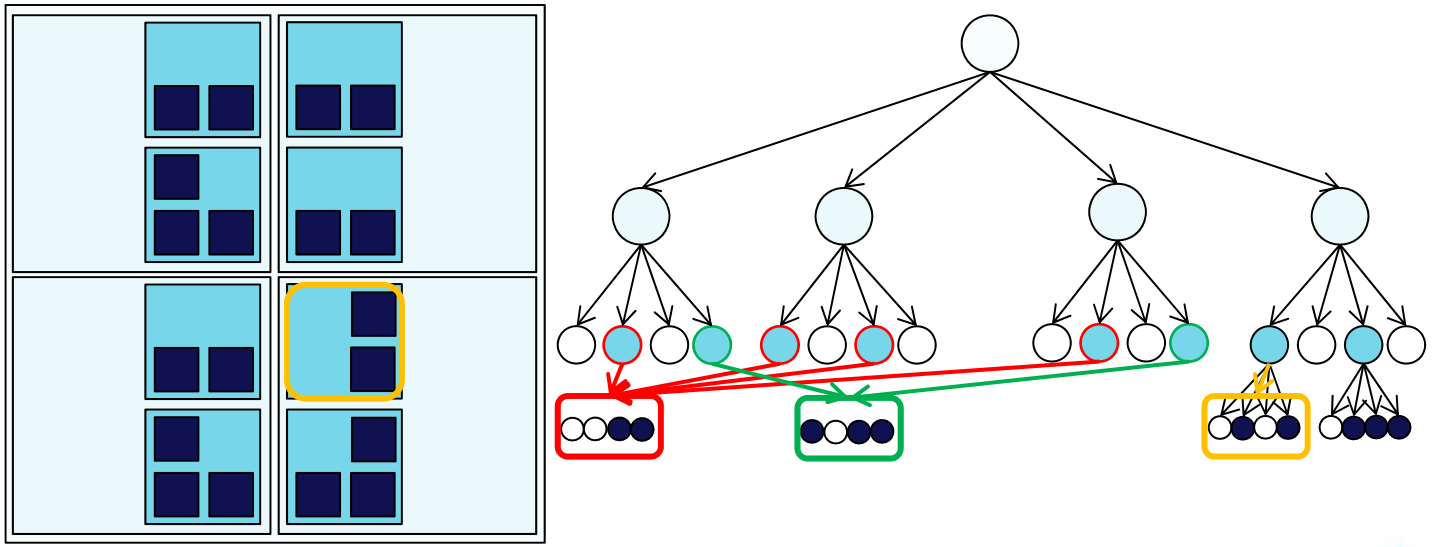

[Kaempe et al. – SIGGRAPH 2013]

Sparse Voxel Directional Acyclic Graphs

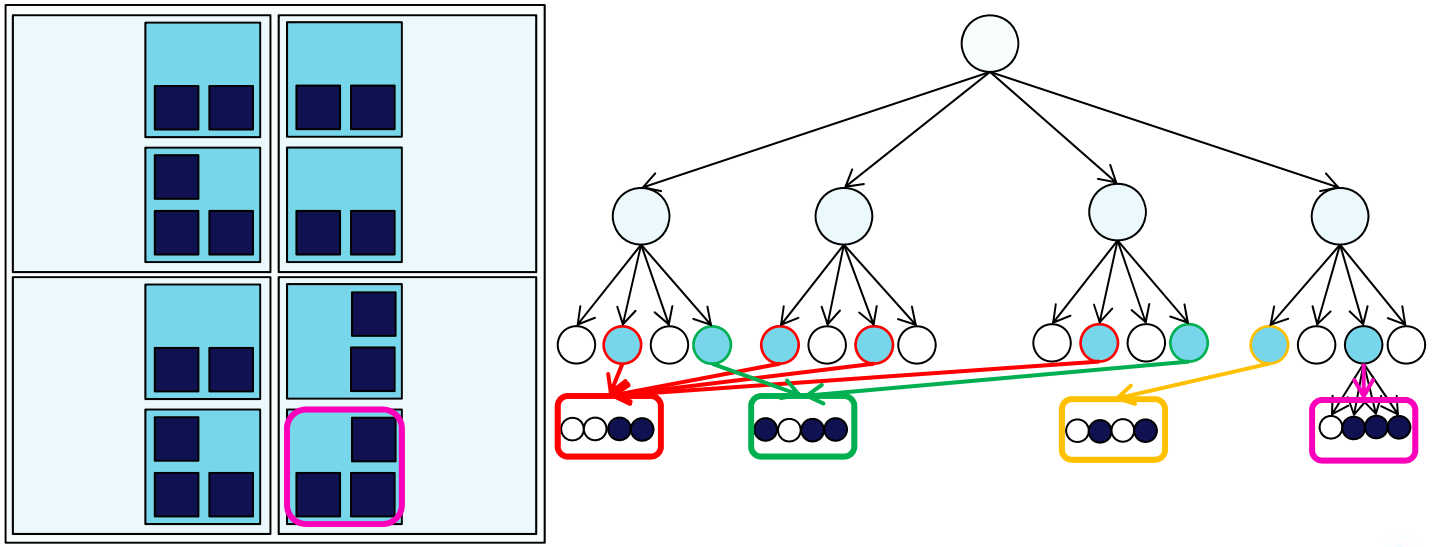

[Kaempe et al. – SIGGRAPH 2013]

Sparse Voxel Directional Acyclic Graphs

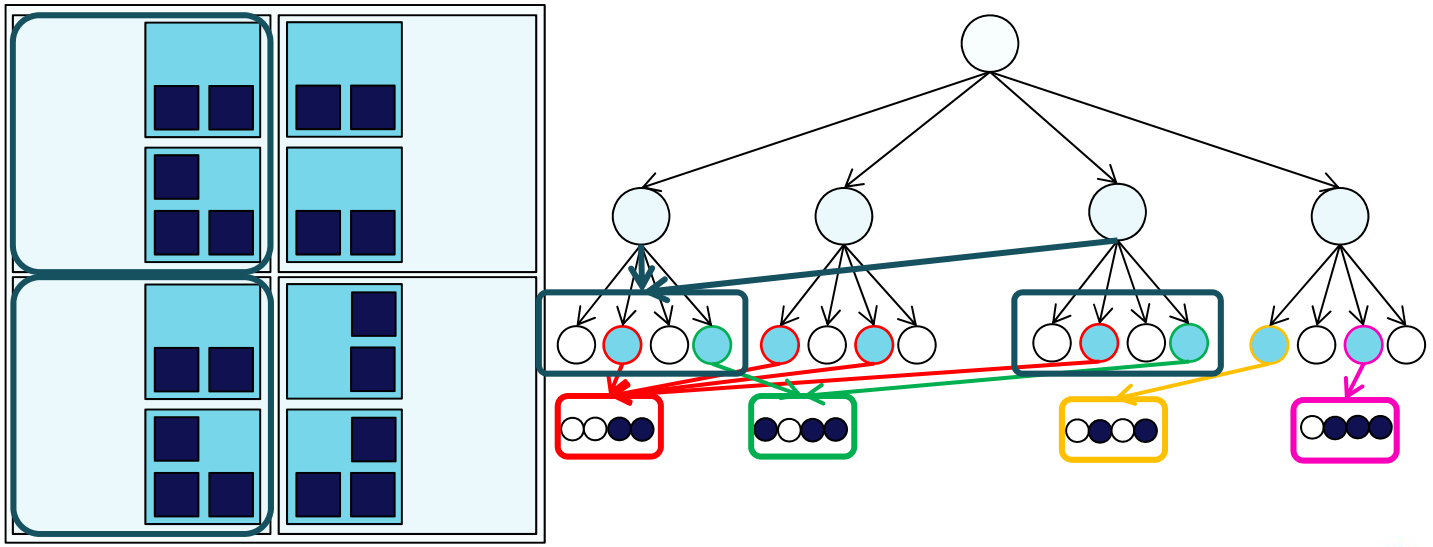

[Dado, Kol, Thiery, Bauszat, Eisemann – Eurographics 2016]

Sparse Voxel Directional Acyclic Graphs

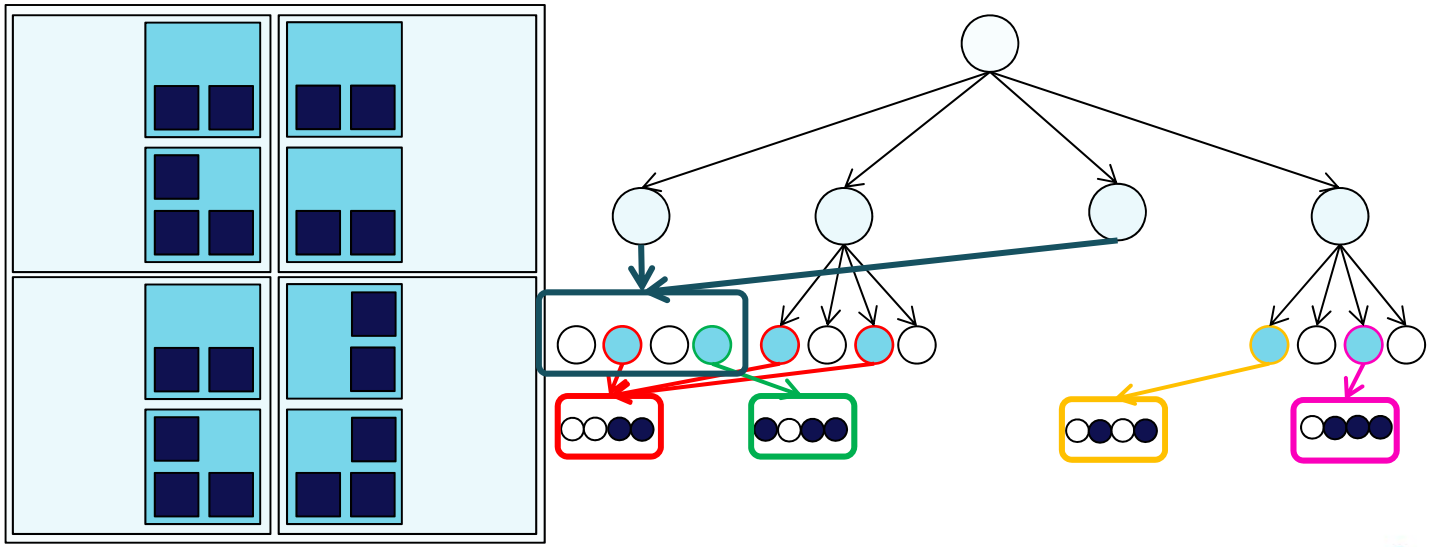

Recent improvements:

More efficient compression
Extension to voxel attributes

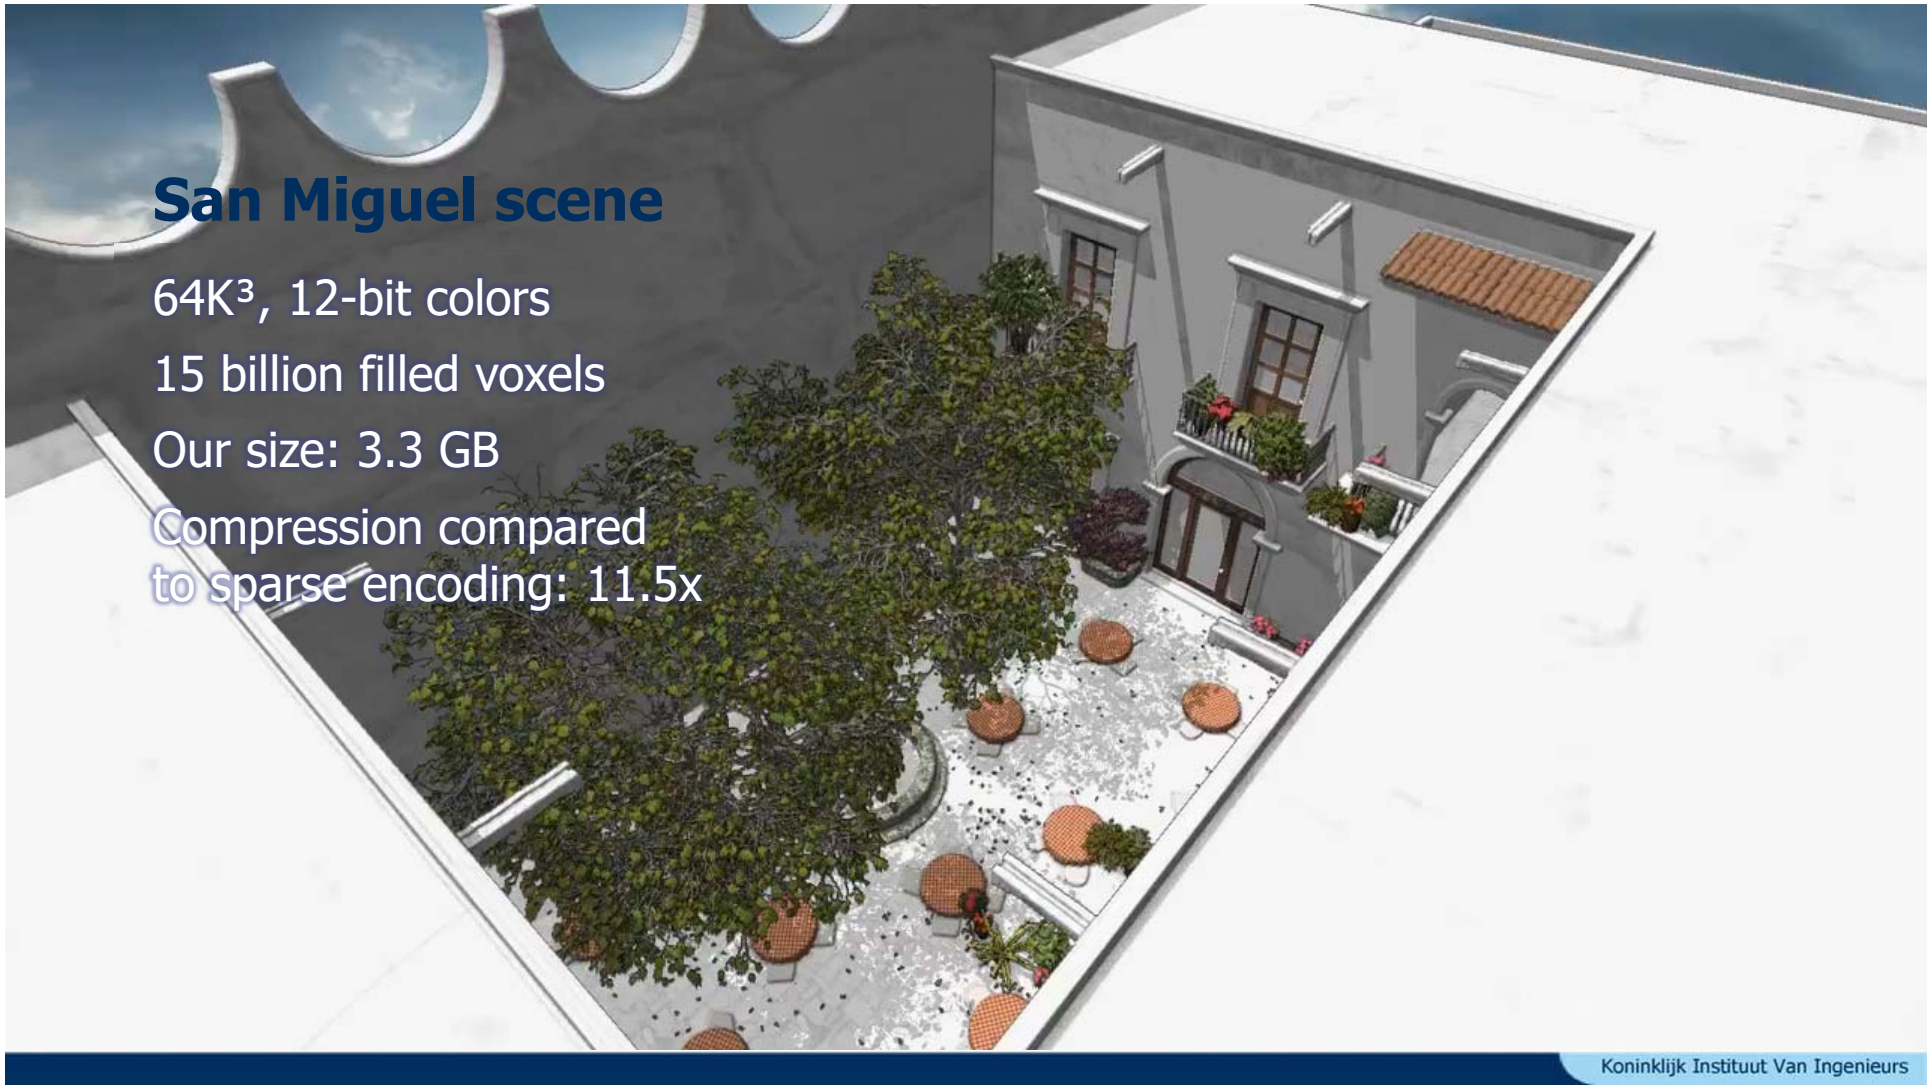
San Miguel scene
64K³, 12-bit colors
15 billion filled voxels
Our size: 3.3 GB
Compression compared
to sparse encoding: 11.5x
Koninklijk Instituut Van Ingenieurs


Arena scene
64K³, 12-bit colors
3.3 billion filled voxels
Our size: 1.0 GB
Compression compared
to sparse encoding: 9.3x
Koninklijk Instituut Van Ingenieurs

[Dado, Kol, Thiery, Bauszat, Eisemann – Eurographics 2016]

Comparison

Cannot be displayed

PSVOs: [Schnabel and Klein 2006]

ESVOs: [Laine and Karras 2011]

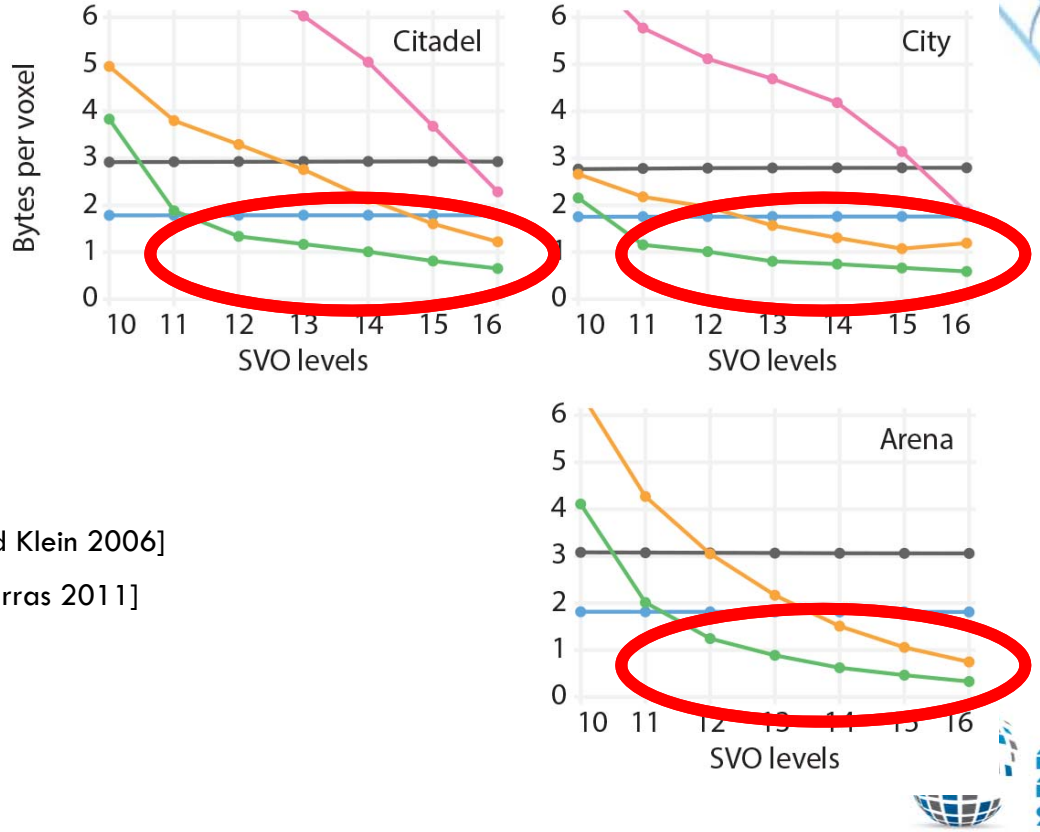

DELFT DATA SCIENCE

Large-Scale Data Rendering

3Di

DELFT DATA SCIENCE

Koninklijk Instituut Van Ingenieurs

Flooding Scenario

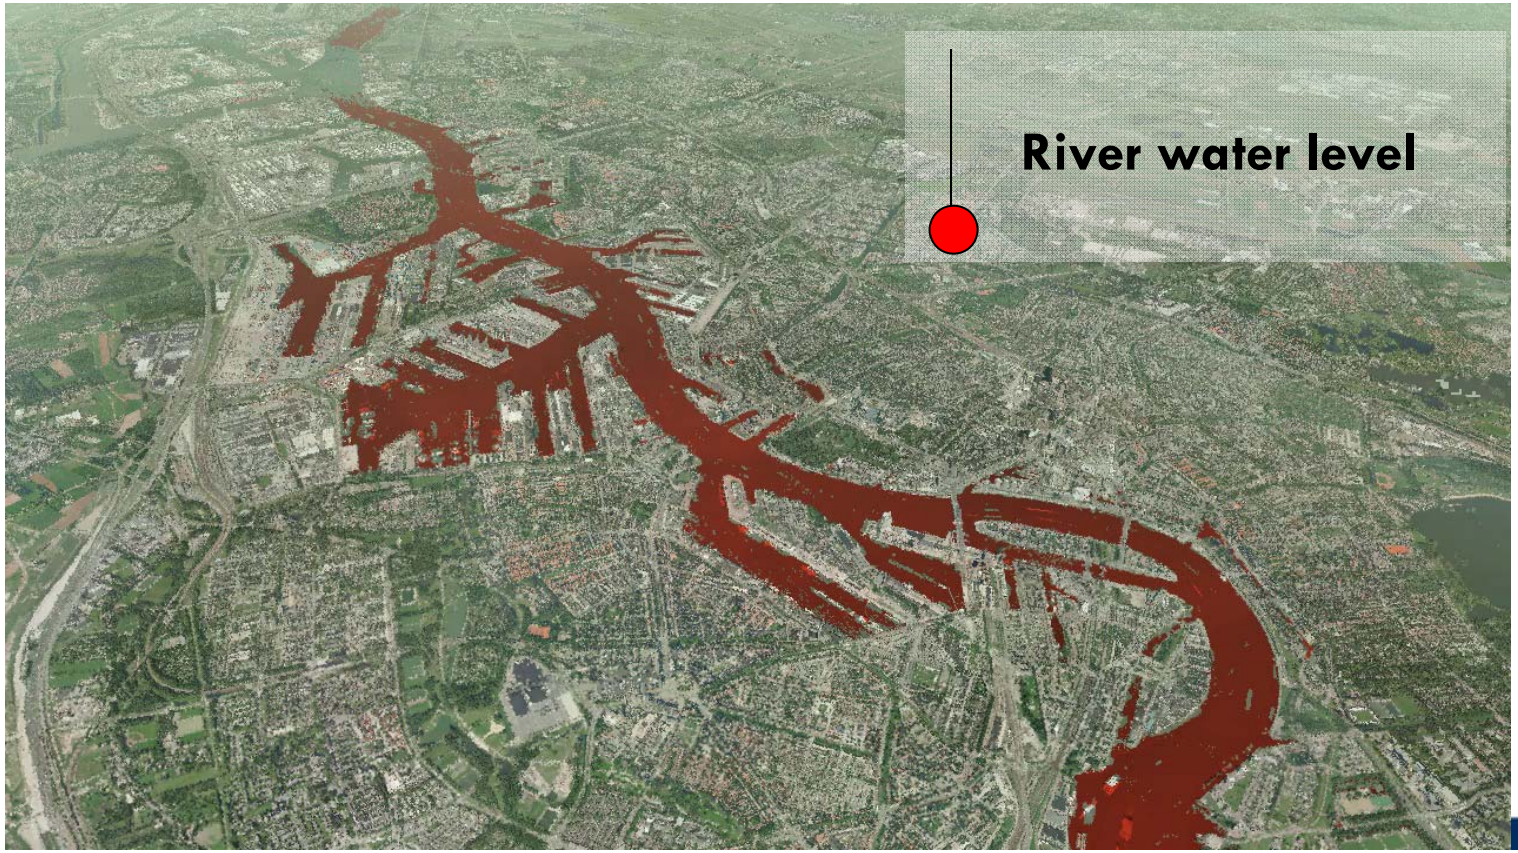

3Di

DELFT DATA SCIENCE

Koninklijk Instituut Van Ingenieurs

[Leskens, Kehl, Tutenel, de Haan, Stelling, Eisemann - Mitigation and Adaptation 2015]

Evaluation of Evacuation Scenarios

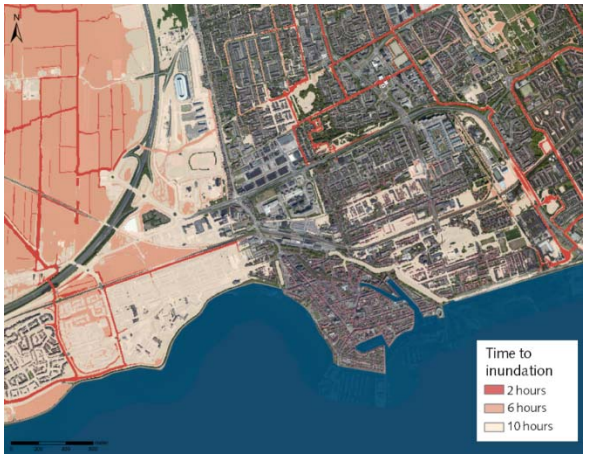

- High benefit of visualizations
  > more and more equal discussion between parties

DELFT DATA SCIENCE

[Leskens, Kehl, Tutenel, Kol, de Haan, Stelling, Eisemann – Science Env. Pol./Mit.Glob.Change 2015]

Benefits of Realistic 3D Visualization

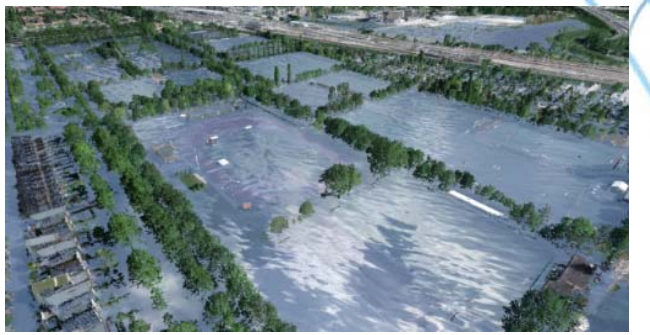

- Accessible to non-experts
- Useful for decision making
  - More involved discussions from all parties
  - Better estimation of damages/dangers
  - Better understanding of the context

DELFT DATA SCIENCE

Exhibitions

- Reproduction of 1953 flooding exhibited in the
  - Noodwatermuseum Zeeland
  - Delft Science Center

DELFT DATA SCIENCE

Large-Scale Rendering

- Ray Tracing
- Graphics Pipeline
- Specialized Methods for Different Data Types
  - Height-Field Data, Voxel Data, Data Management, Compression

Effective Data Visualization Requires

- Large-Scale Rendering

- Visualization and Percepti

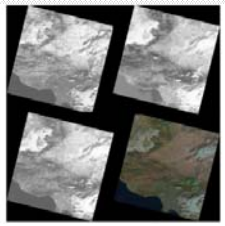

- Data Analysis

DELFT DATA SCIENCE

Questions?

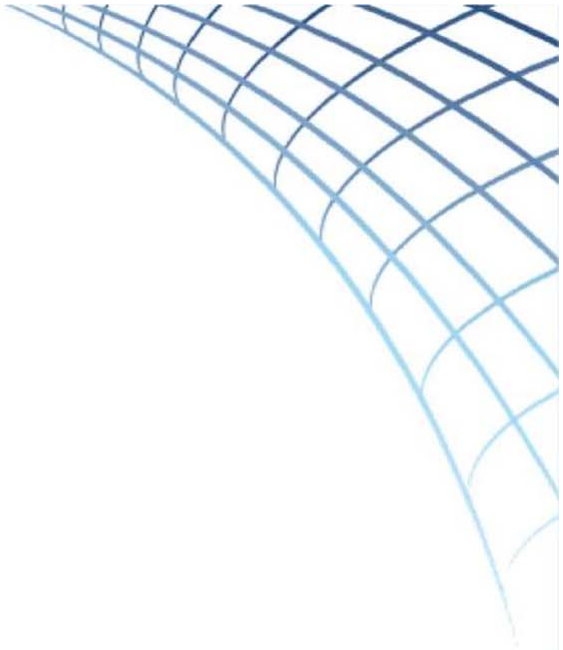

Effective Data Visualization Requires

- Large-Scale Rendering

- Visualization and Percepti

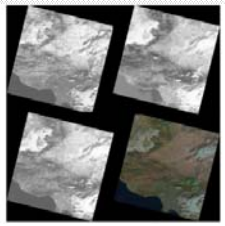

- Data Analysis

DELFT DATA SCIENCE

Visualization and Perception

- Realistic Rendering
- Perceptual Methods
- Visualization

Physically-Based Rendering

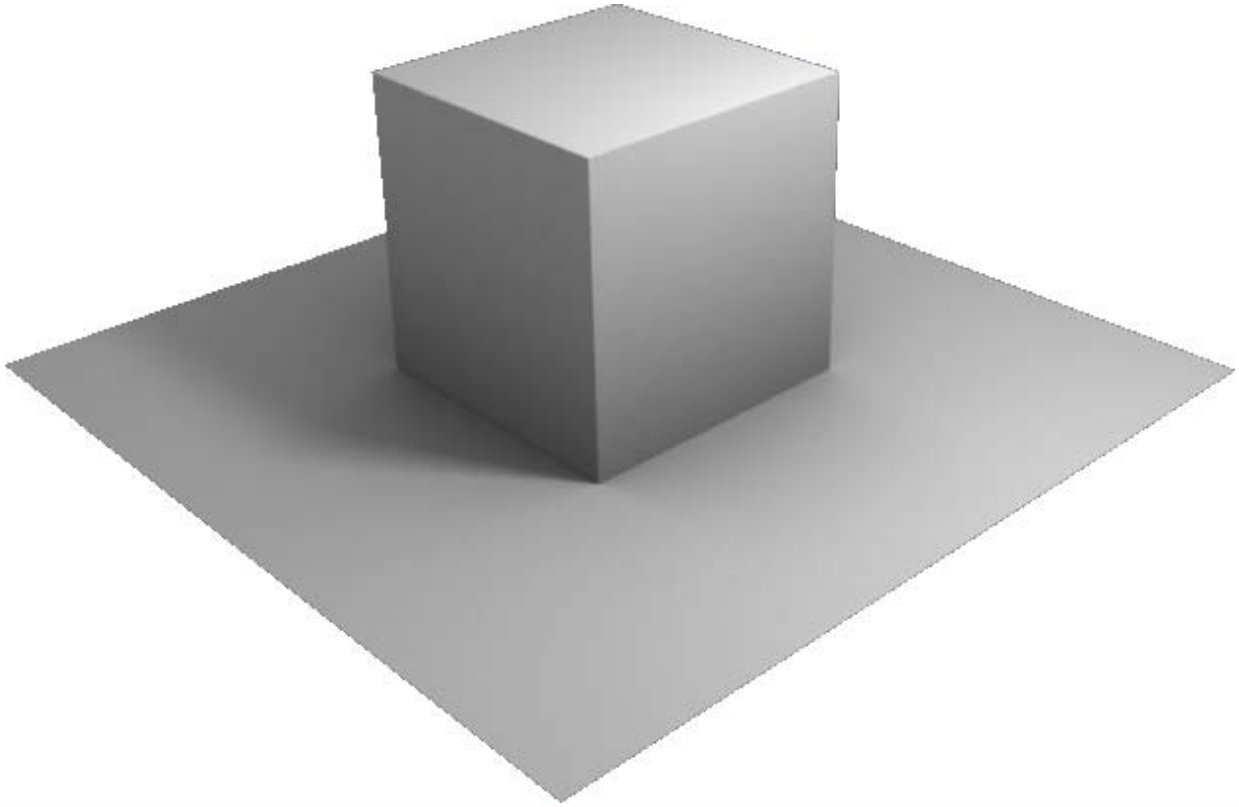

[Eisemann&Decoret - CGF 2008]
[Holländer, Ritschel, Eisemann, Boubekeur - EGSR 2011]
[Ritschel, Eisemann, Ha, Kim, Seidel - CGF 2011]
[Baboud, Eisemann, Seidel, TVCG 2012]

Light Physics are Complex

- Reflections
- Refractions
- Caustics
- Transparency
- Global Illumin.

...

DELFT DATA SCIENCE

[Lee, Eisemann, Seidel – SIGGRAPH Asia'09] [Lee, Eisemann – EGSR'13]
[Lee, Eisemann, Seidel – SIGGRAPH'10] [Kurz, Ritschel, Eisemann, Thormählen, Seidel JVRB14]
[Hullin, Lee, Eisemann, Seidel – SIGGRAPH'11] [Joo, Kwon, Lee, Eisemann, Lee – EGSR16]

Improving Camera Model

- Depth of Field
- Lens Blur
- Lens Flares

...

DELFT DATA SCIENCE

Global Illumination

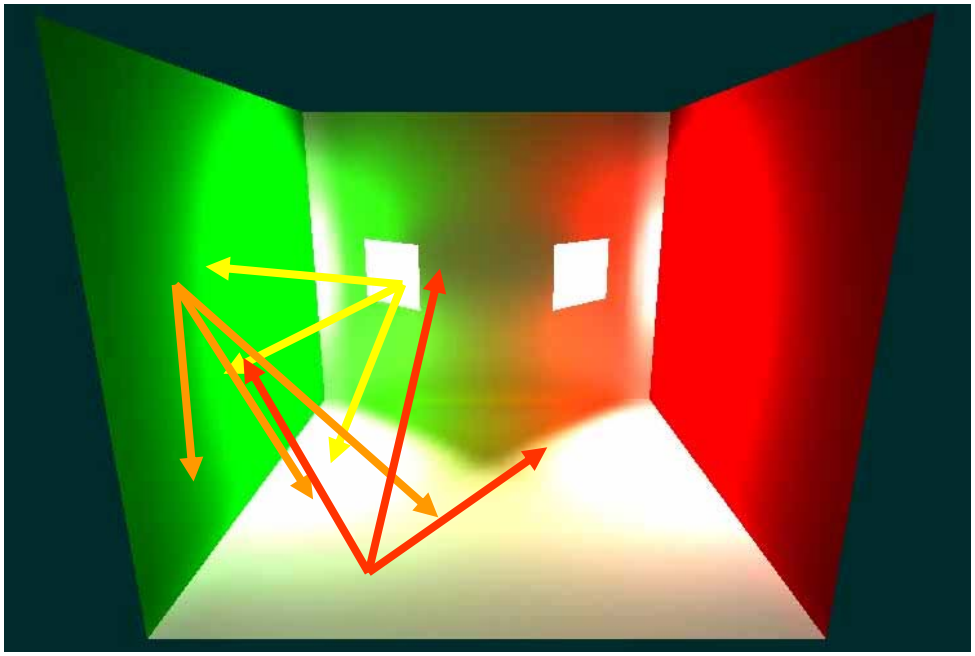

In 2002…
More than 4 hours
8 triangles

Rendering Equation

$$L(x, \theta_0, \varphi_0) = L_e(x, \theta_0, \varphi_0) + \int_{\Omega} \rho_{bd}(x, \theta_0, \varphi_0, \theta, \varphi) L_i(x, \theta, \varphi) \cos\theta \, d\omega$$

observed light = emitted light + reflected light

Rendering Equation

- Position

$$L(x, \theta_0, \varphi_0) = L_e(x, \theta_0, \varphi_0) + \int_{\Omega} \rho_{bd}(x, \theta_0, \varphi_0, \theta, \varphi) L_i(x, \theta, \varphi) \cos\theta \, d\omega$$

Spatial domain

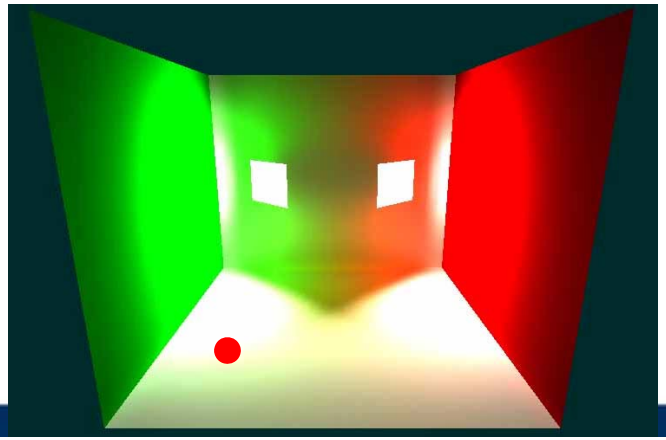

Rendering Equation

- Orientation

$$L(x, \theta_0, \varphi_0) = L_e(x, \theta_0, \varphi_0) + \int_\Omega \rho_{bd}(x, \theta_0, \varphi_0, \theta, \varphi) L_i(x, \theta, \varphi) \cos\theta \, d\omega$$

Angular domain

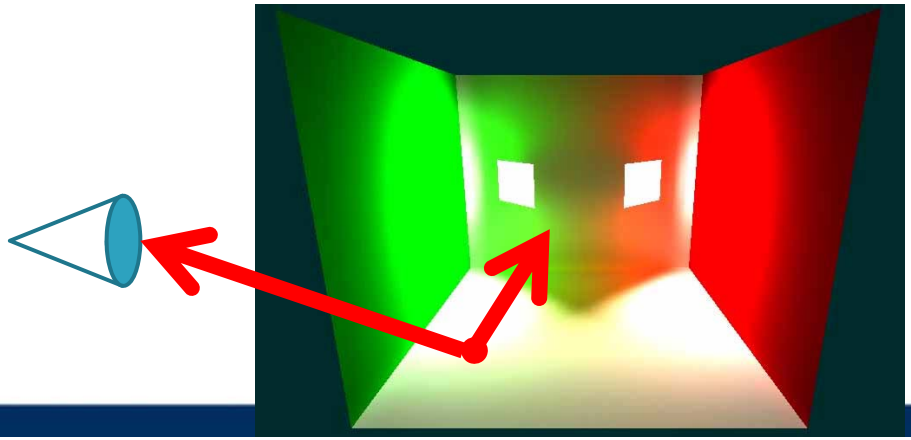

Rendering Equation

- Radiance

$$L(x,\theta_0,\varphi_0) = L_e(x,\theta_0,\varphi_0) + \int_\Omega \rho_{bd}(x,\theta_0,\varphi_0,\theta,\varphi) L_i(x,\theta,\varphi) \cos\theta d\omega$$

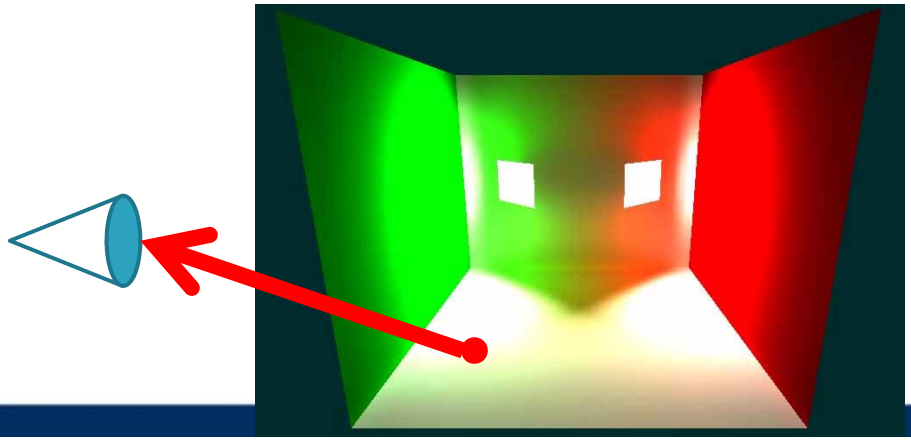

Rendering Equation

$$L(x,\theta_0,\varphi_0) = L_e(x,\theta_0,\varphi_0) + \int_\Omega \rho_{bd}(x,\theta_0,\varphi_0,\theta,\varphi) L_i(x,\theta,\varphi) \cos\theta d\omega$$

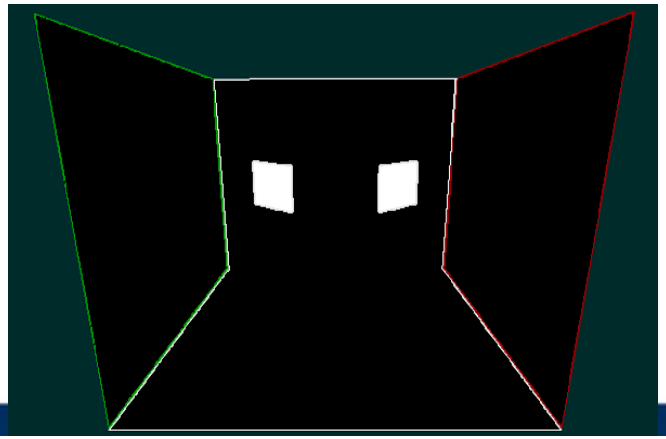

Rendering Equation

$$L(x,\theta_0,\varphi_0) = L_e(x,\theta_0,\varphi_0) + \int_\Omega \rho_{bd}(x,\theta_0,\varphi_0,\theta,\varphi) L_i(x,\theta,\varphi) \cos\theta d\omega$$

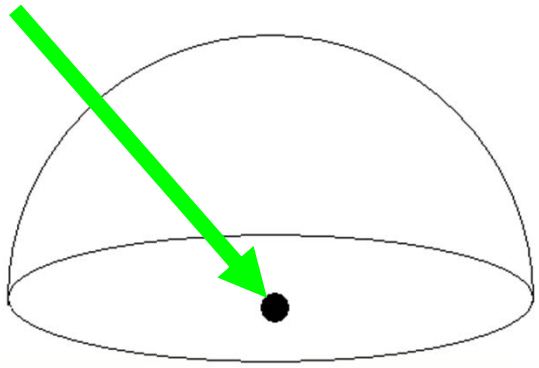

Rendering Equation

$$L(x,\theta_0,\varphi_0) = L_e(x,\theta_0,\varphi_0) + \int_\Omega \rho_{bd}(x,\theta_0,\varphi_0,\theta,\varphi) L_i(x,\theta,\varphi) \cos\theta d\omega$$

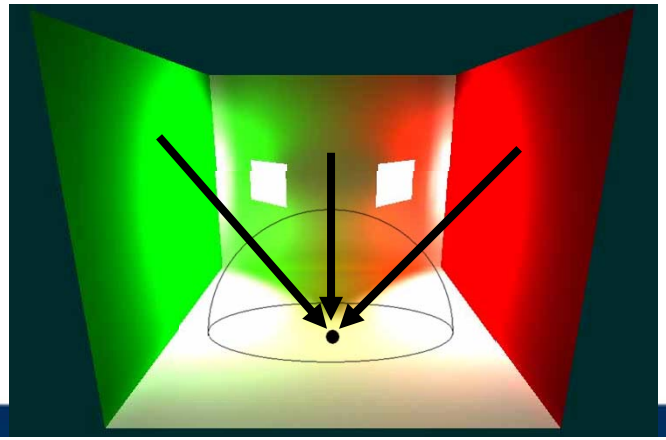

Rendering Equation

$$L(x,\theta_0,\varphi_0) = L_e(x,\theta_0,\varphi_0) + \int_\Omega \rho_{bd}(x,\theta_0,\varphi_0,\theta,\varphi) L_i(x,\theta,\varphi) \cos\theta d\omega$$

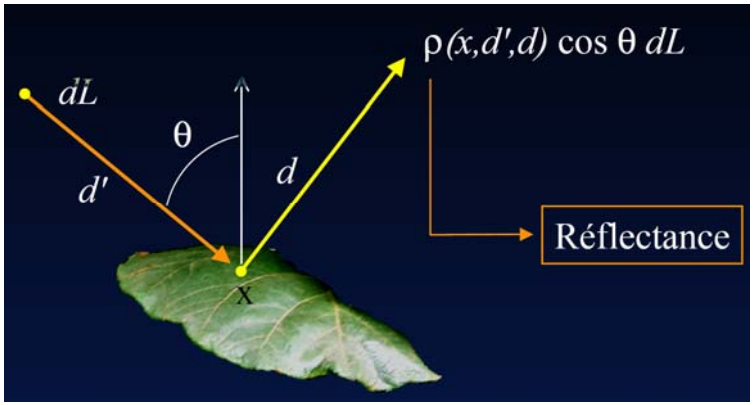

DELFT DATA SCIENCE

Rendering Equation

$$L(x,\theta_0,\varphi_0) = L_e(x,\theta_0,\varphi_0) + \int_\Omega \rho_{bd}(x,\theta_0,\varphi_0,\theta,\varphi) L_i(x,\theta,\varphi) \cos\theta d\omega$$

Rendering Equation

$$L(x,\theta_0,\varphi_0) = L_e(x,\theta_0,\varphi_0) + \int_\Omega \rho_{bd}(x,\theta_0,\varphi_0,\theta,\varphi) L_i(x,\theta,\varphi) \cos\theta d\omega$$

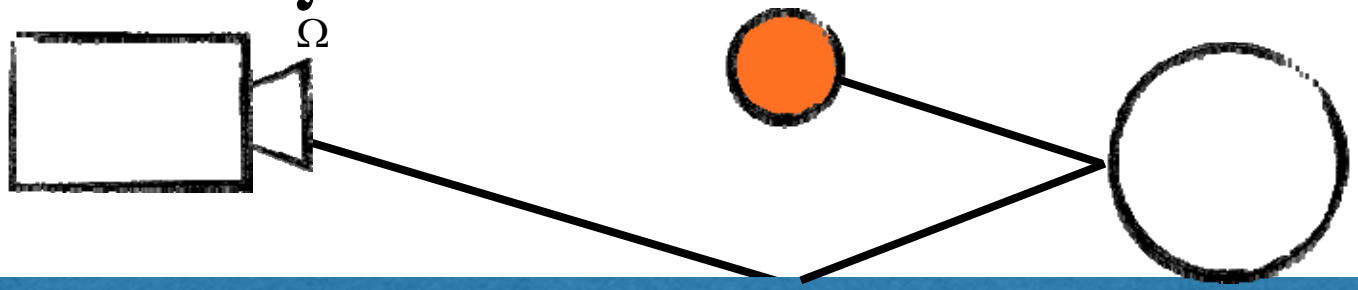

Importance Sampling

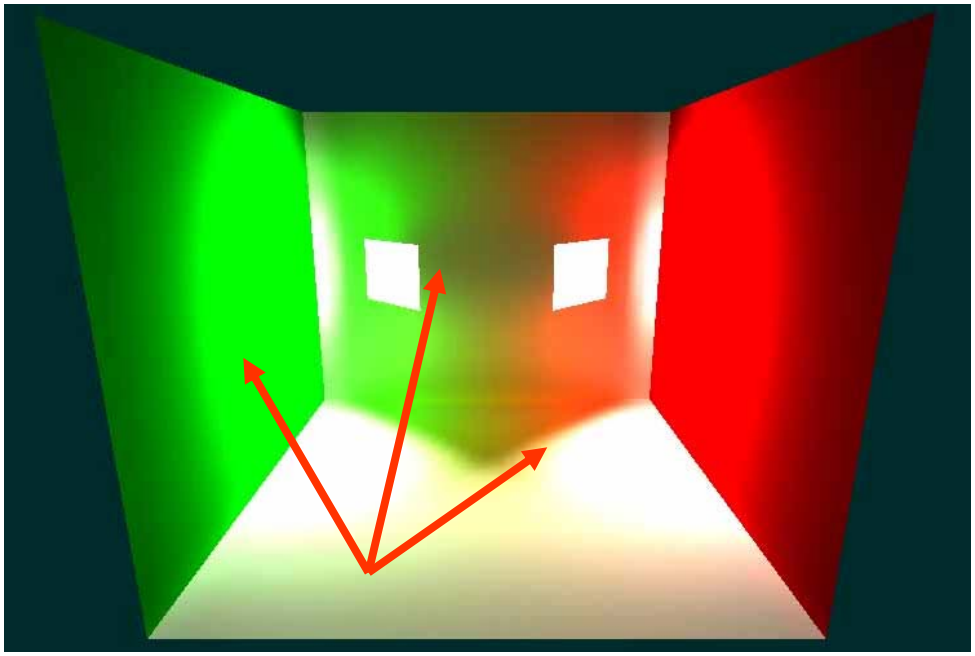

"Shoot in the right direction"

Importance Sampling

- Uniform sampling of a function?

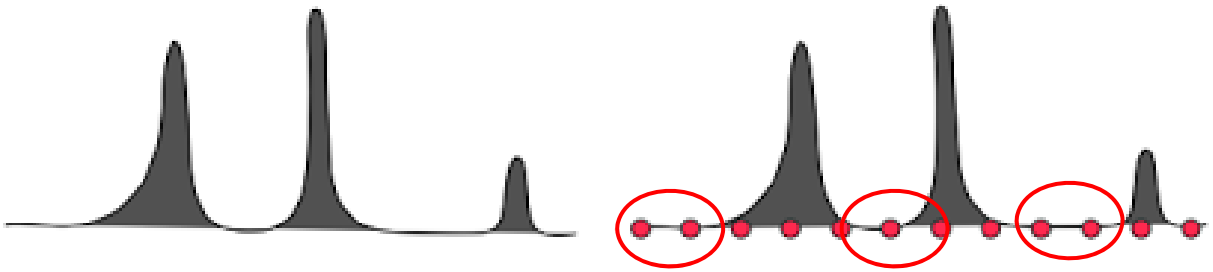

DELFT DATA SCIENCE

Importance Sampling

- Uniform sampling of a function?

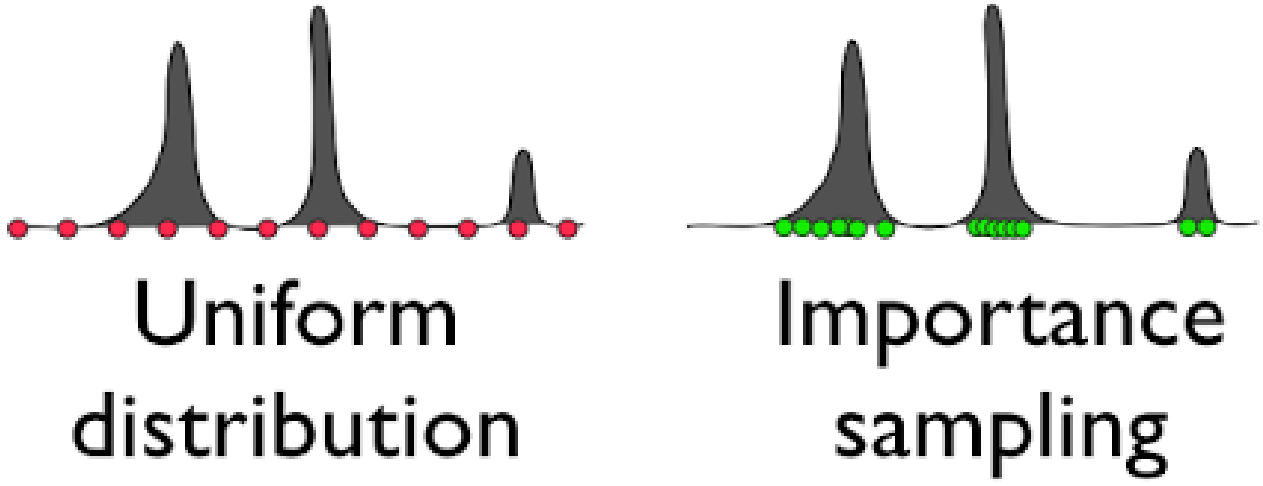

Importance Sampling

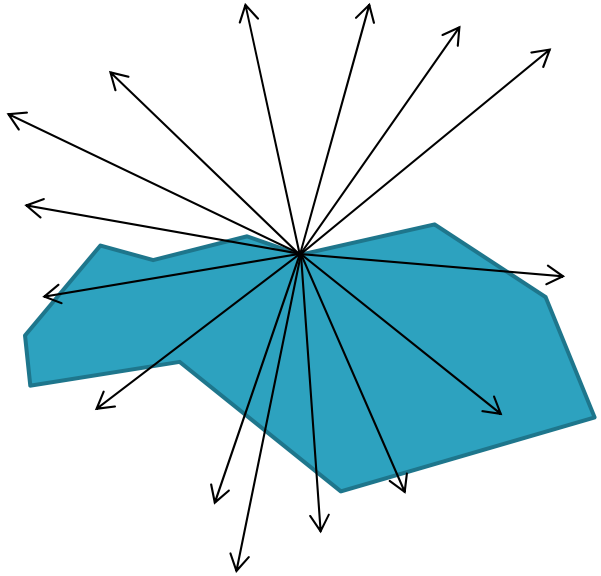

Importance Sampling

- Steer sampling via light sources?

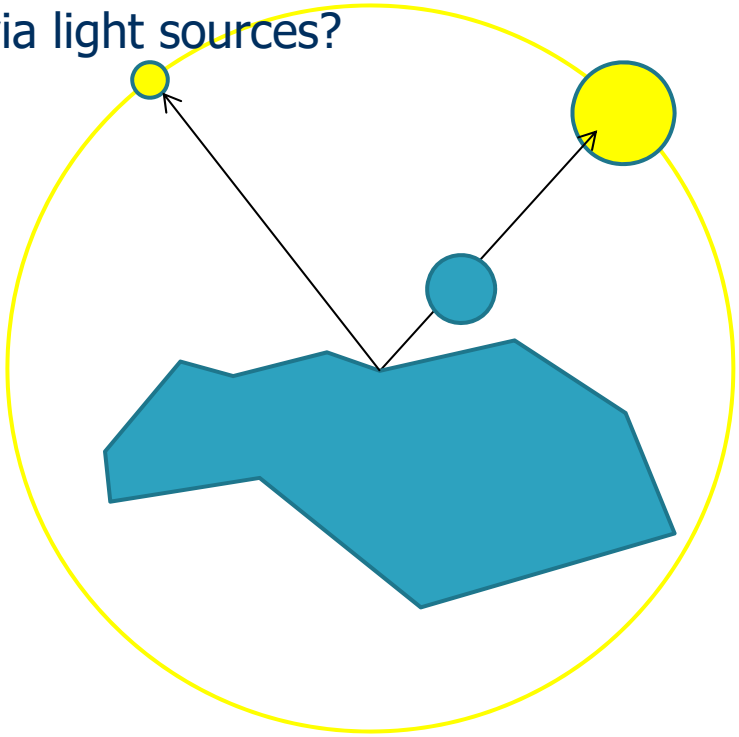

[Kroes, Eisemannx2 - GI2015] , Sweeney Award Winner
[v. Radziewski, Kroes, Eisemannx2 - TVCG 2016]

Importance Sampling

- Same amount of rays

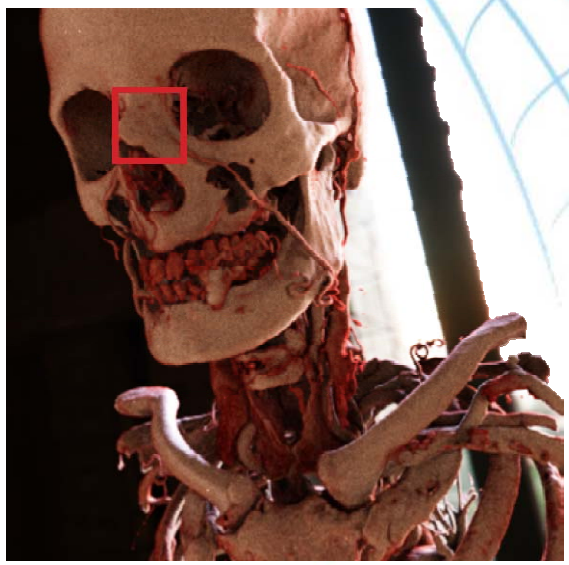

Two-step= Light and Visibility together are used for sampling

DELFT DATA SCIENCE

[Kroes, Eisemannx2 - GI2015] , Sweeney Award Winner
[v. Radziewski, Kroes, Eisemannx2 - TVCG 2016]

How to determine visibility?

- Shooting rays from every point is too expensive

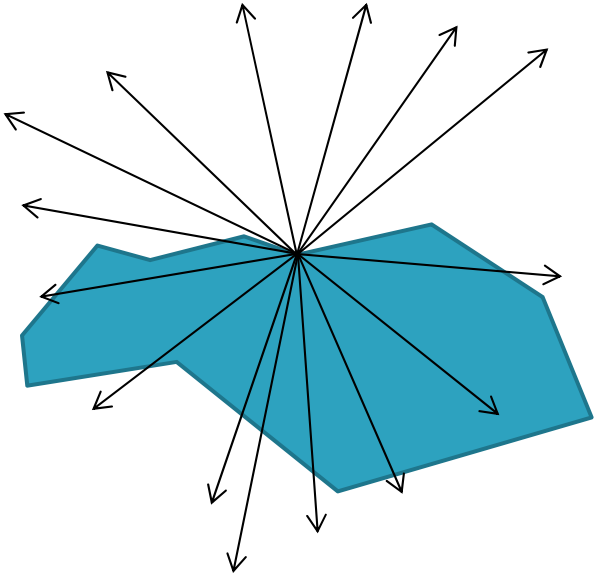

Sweeping Solution

- Fix a random direction:

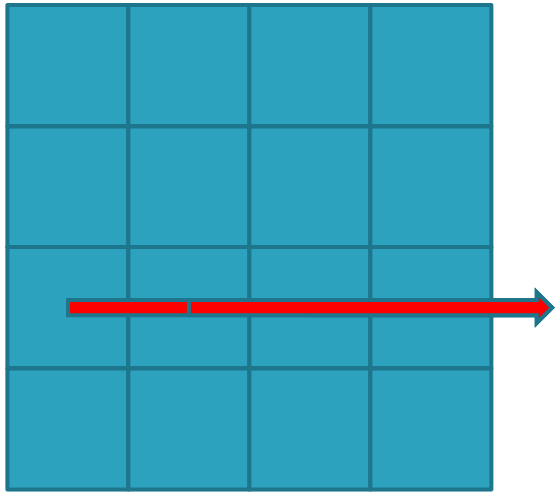

[Kroes, Eisemannx2 - GI2015]
[v. Radziewski, Kroes, Eisemannx2 - TVCG 2016]

Sweeping Solution

- Accumulate visibility in a single sweep for all voxels

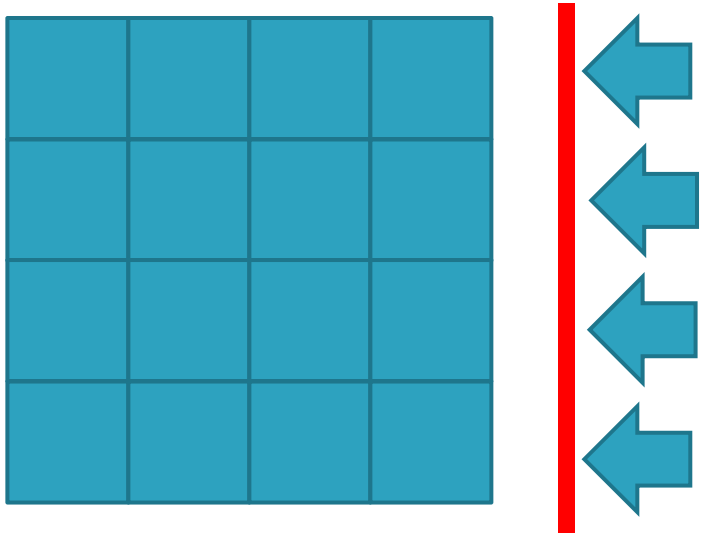

[Kroes, Eisemannx2 - GI2015]
[v. Radziewski, Kroes, Eisemannx2 - TVCG 2016]

Store Visibility

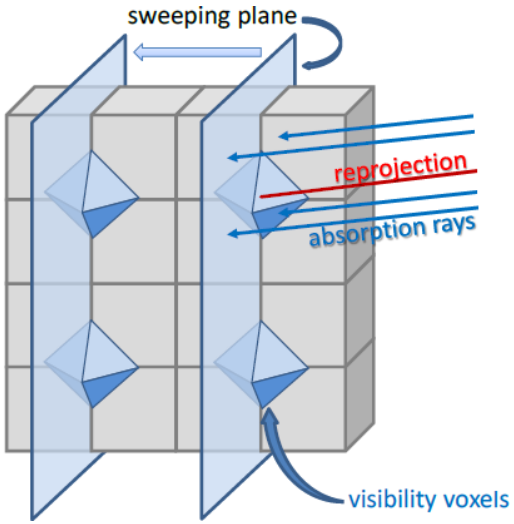

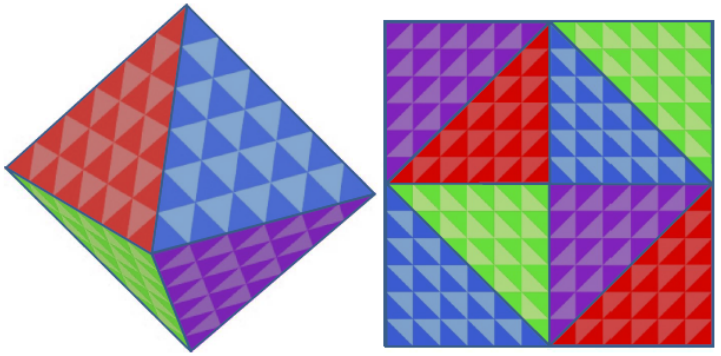

- Fast parallelized scheme for image-based visibility information

Convergence with our importance sampling

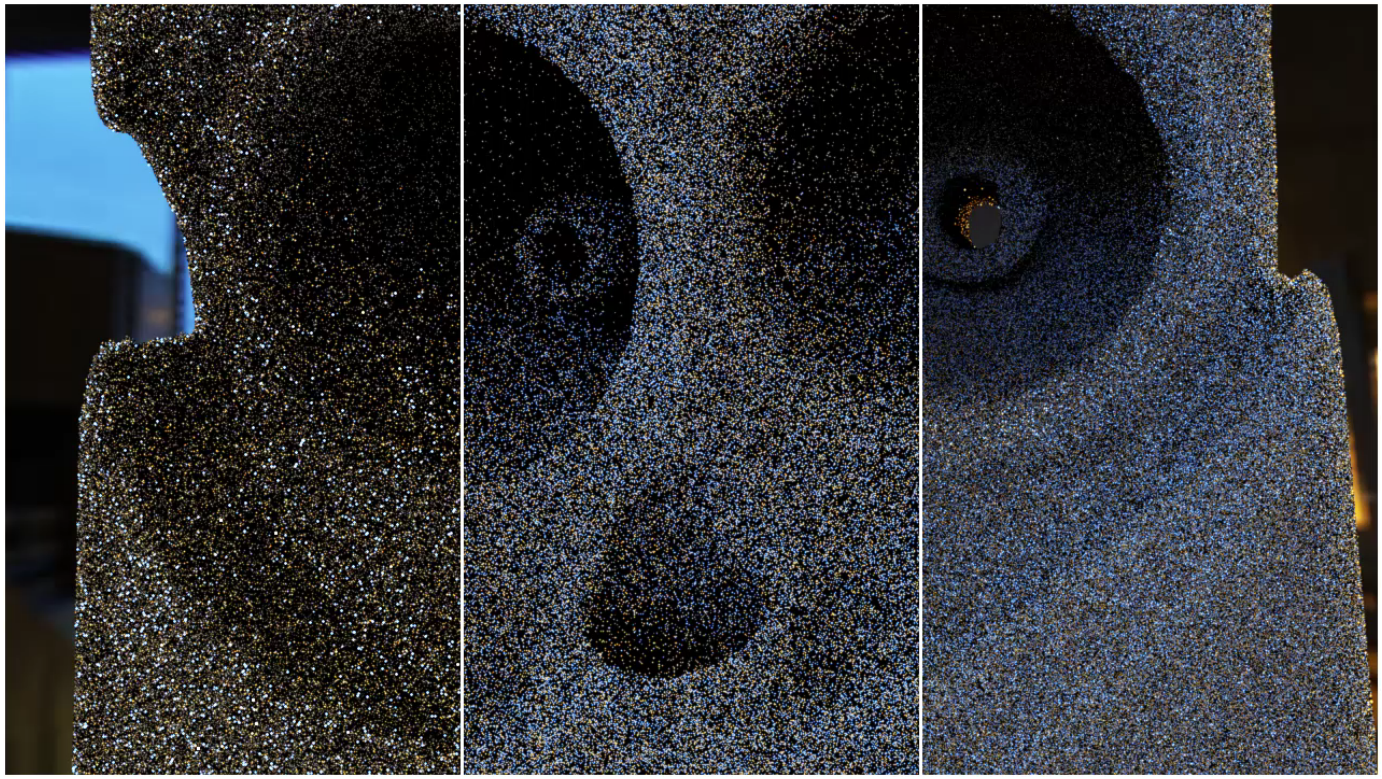

[Bauszat, Eisemannx2, Magnor – Eurographics 2015]

Light Reconstruction

10 mins:

2 mins:

"What signal is likely, given my observations?"

[Bauszat, Eisemannx2, Magnor – Eurographics 2015]

Light Reconstruction

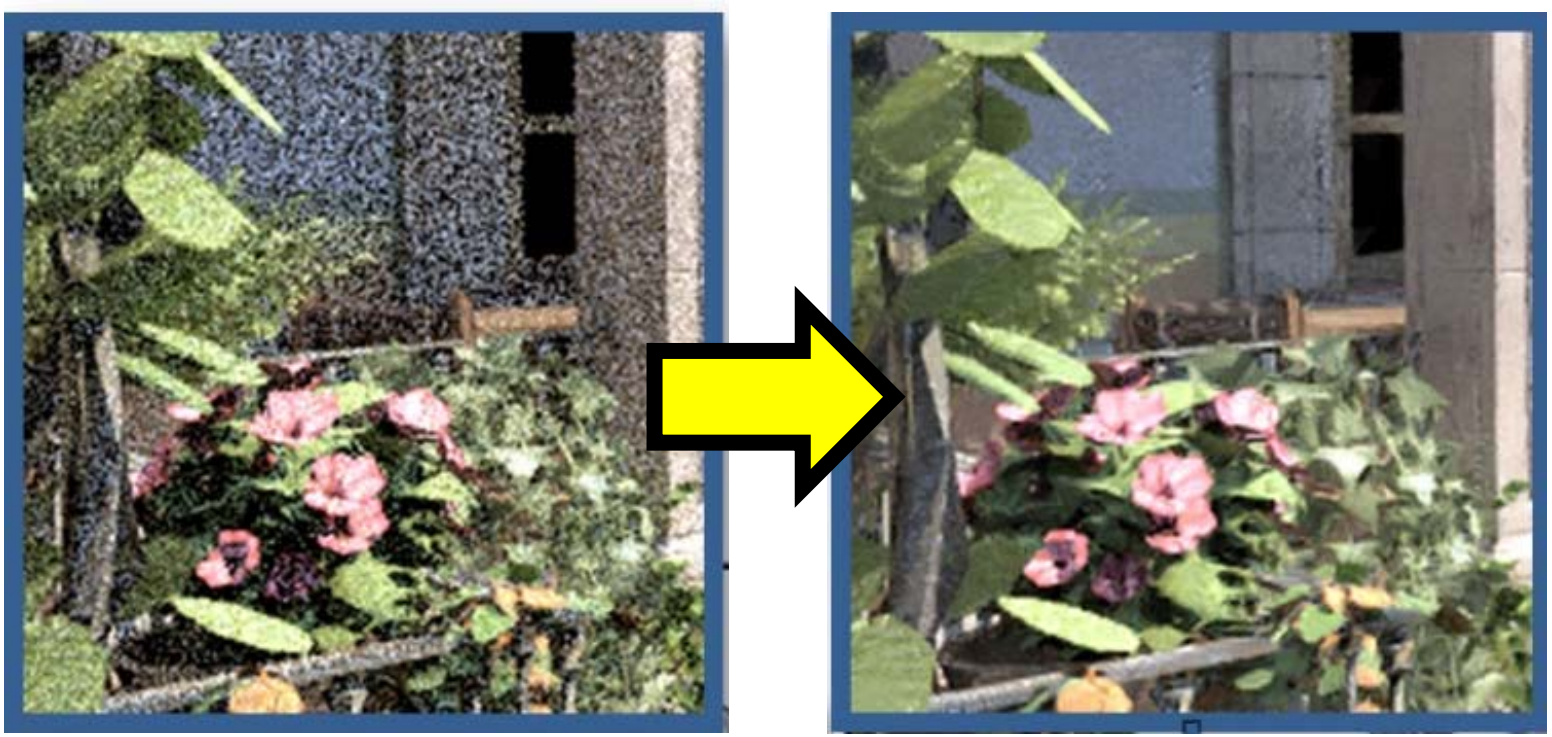

How can we go real-time?

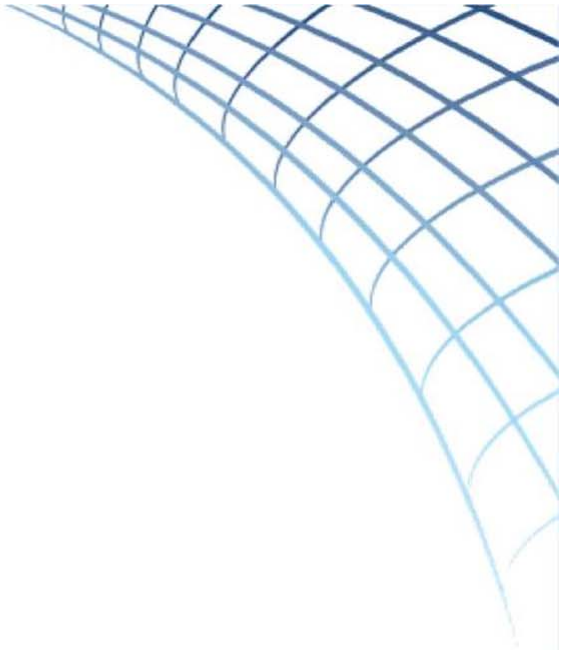

[Crassin, Neyret, Sainz, Green, Eisemann, PG 2012]

Voxel-Cone Tracing

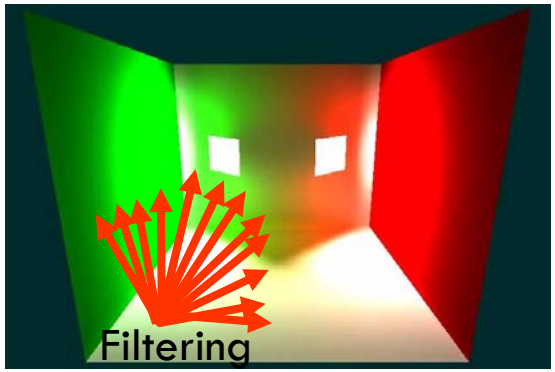

- Illumination Filtering

Filtering as approximation of many rays

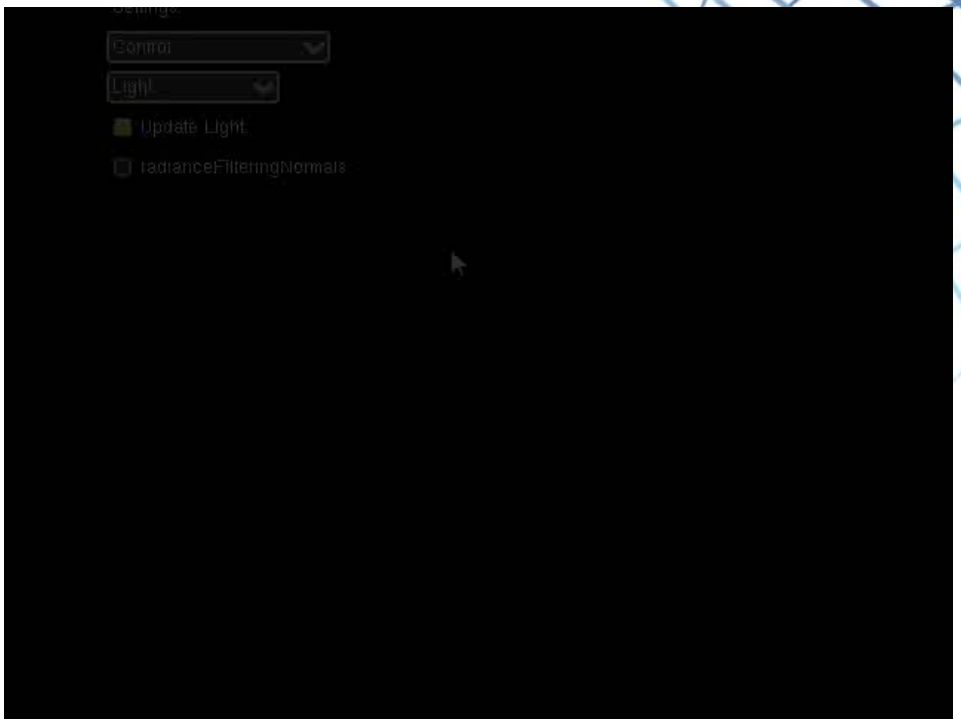

Basis for GI in Unreal Engine 4 & NVIDIA VXGI

Based on:
[Crassin, Neyret, Sainz, Green, Eisemann, PG 2012]

Global Illumination

Unreal Engine 4 / Cry Engine 2015

NVIDIA VXGI 2015 and VXAO 2016

Fig 6. Rise of the Tomb Raider rendered with VXAO and HBAO+

DELFT DATA SCIENCE

NVIDIA Debunks the Moon Landing…

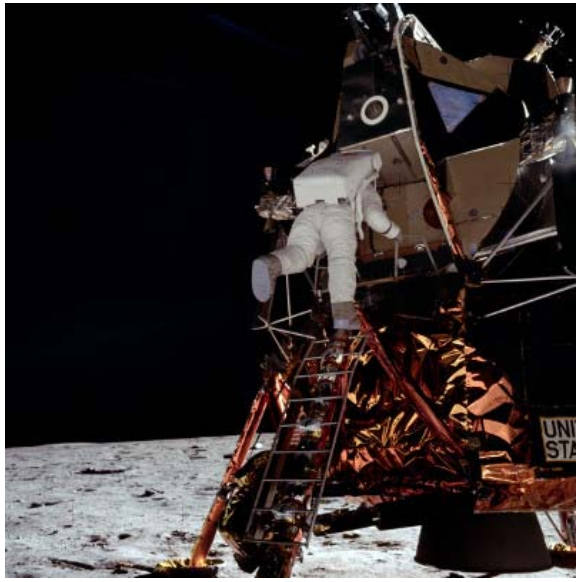

Copyright NVIDIA

NVIDIA Debunks the Moon Landing…

Copyright NVIDIA

Exposure Render

- Integrates similar methods for volume rendering

How many images per second do we need?

Questions?

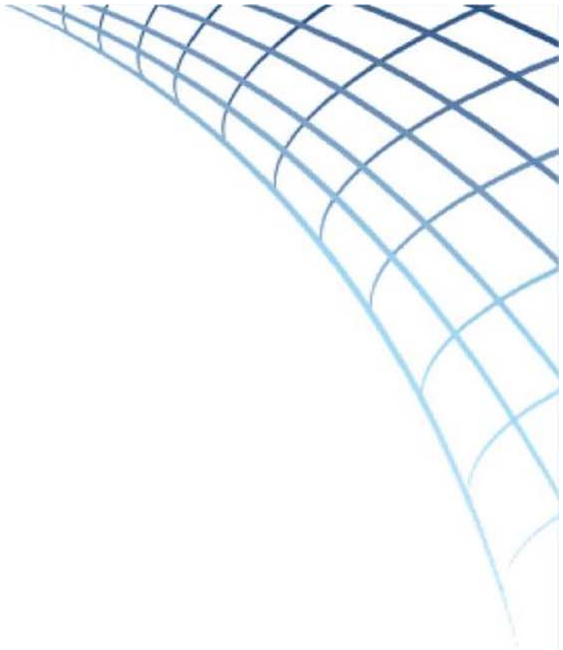

How many images per second do we need?

- More images:
  - Increased detail perception / Blur reduction
  - Increased task performance

40 Hz

120 Hz

DELFT DATA SCIENCE

[Stengel, Bauszat, Eisemann x2, Magnor – TVCG2015]

[Didyk, Eisemann, Ritschel, Myszkowski, Seidel – Eurographics 2010]

P

REALITY

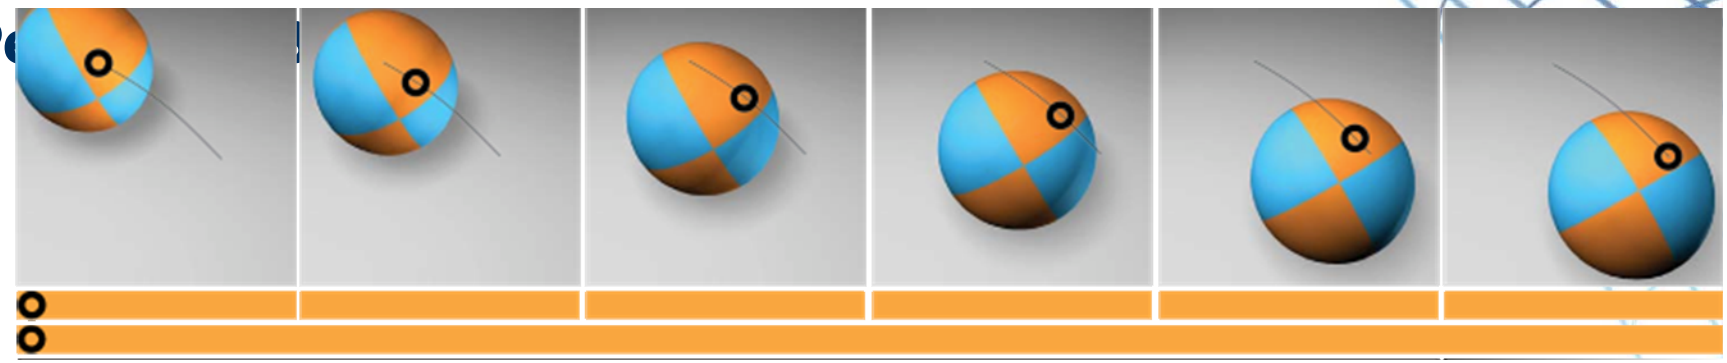

“Simple” solution: Produce more frames…

DELFT DATA SCIENCE

Koninklijk Instituut Van Ingenieurs

[Didyk, Eisemann, Ritschel, Myszkowski, Seidel – Eurographics 2010]

Perceptual Upsampling

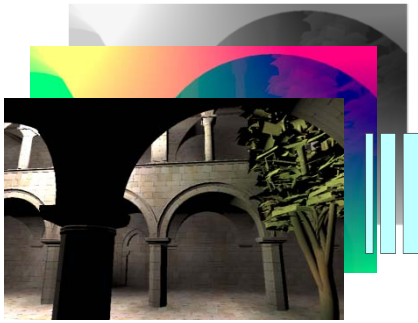

original frames
+ motion flow & depth
(40Hz)

Efficient
Perceptually-motivated
GPU extrapolation

Reduced blur
(120 Hz)

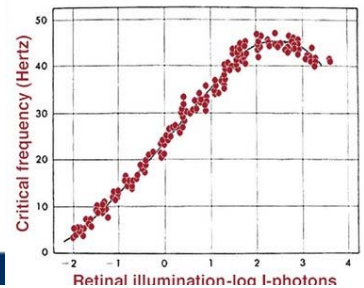

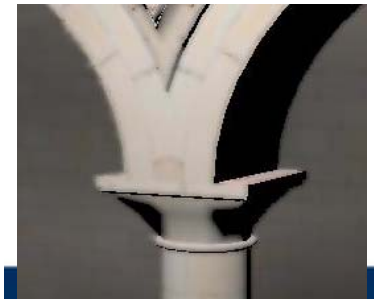

DELFT DATA SCIENCE

Image Decomposition

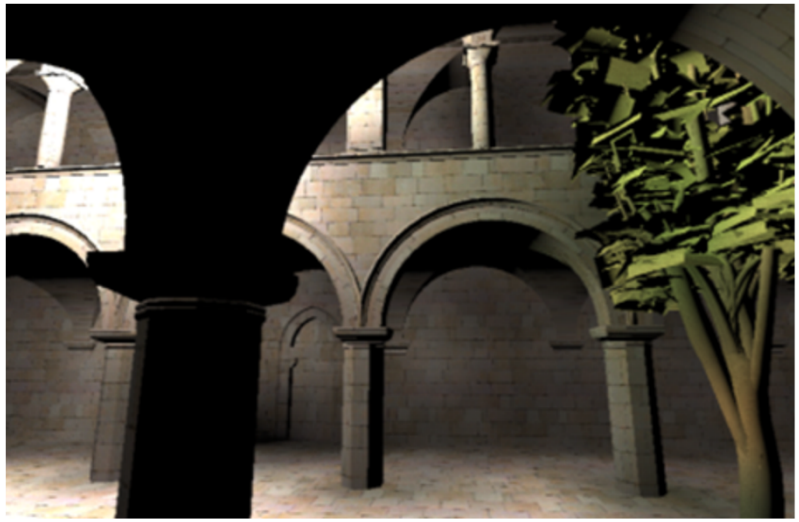

Combine high-frequency information

Frame 0

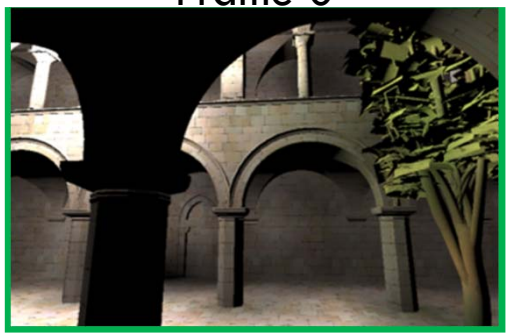

Frame 1

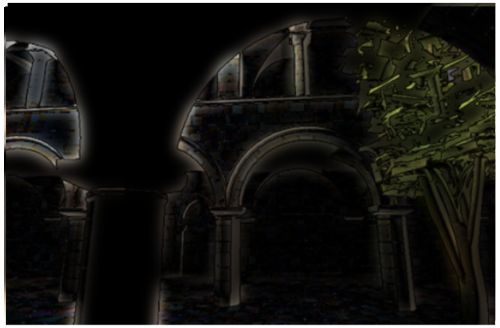

Frame 2

- At 120 Hz indistinguishable from original

DELFT DATA SCIENCE

Combine high-frequency information

Frame 0

Frame 1

Frame 2

- Our idea: Turn the principle around

DELFT DATA SCIENCE

Produce frames via simple warping

- Transform grid via motion flow (respect depth)

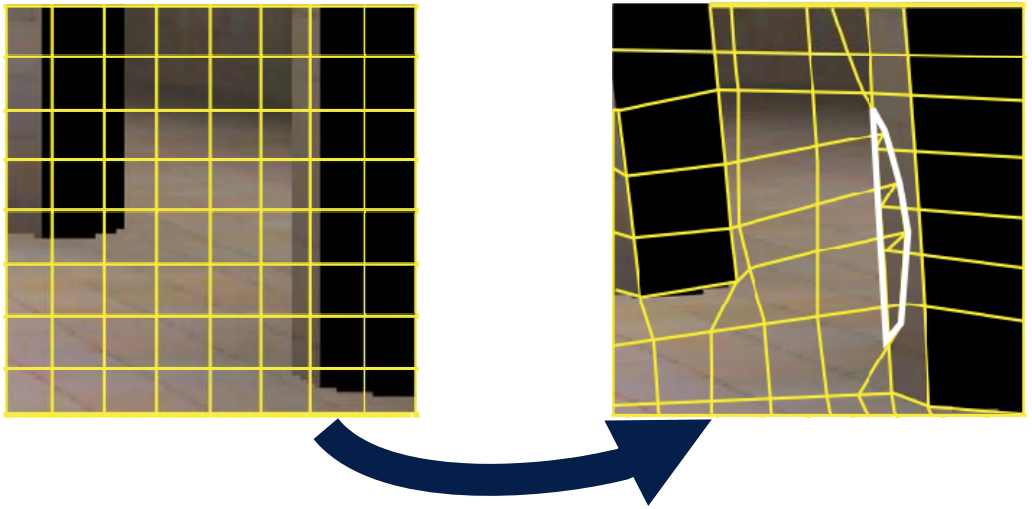

~1 ms in full HD

DELFT DATA SCIENCE

Koninklijk Instituut Van Ingenieurs

[Didyk, Eisemann, Ritschel, Myszkowski, Seidel – Eurographics 2010]

Perceptual Upsampling

- Quality Performance

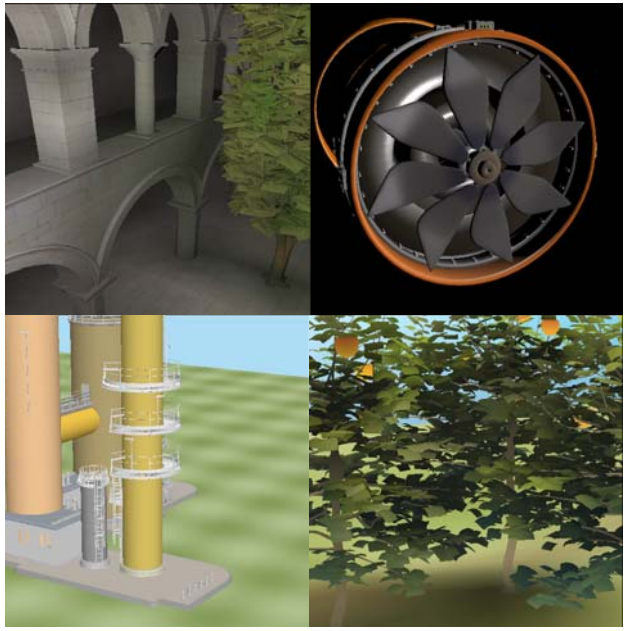

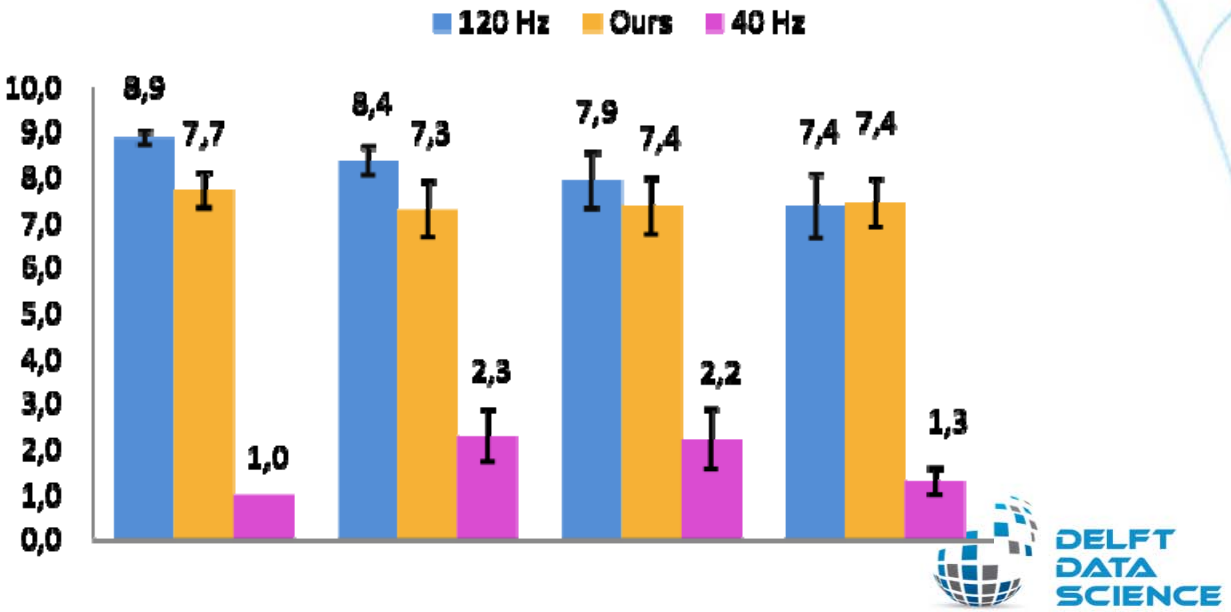

DELFT DATA SCIENCE

[Didyk, Eisemann, Ritschel, Myszkowski, Seidel – Eurographics 2010]

Perceptual Upsampling

- Quality Performance

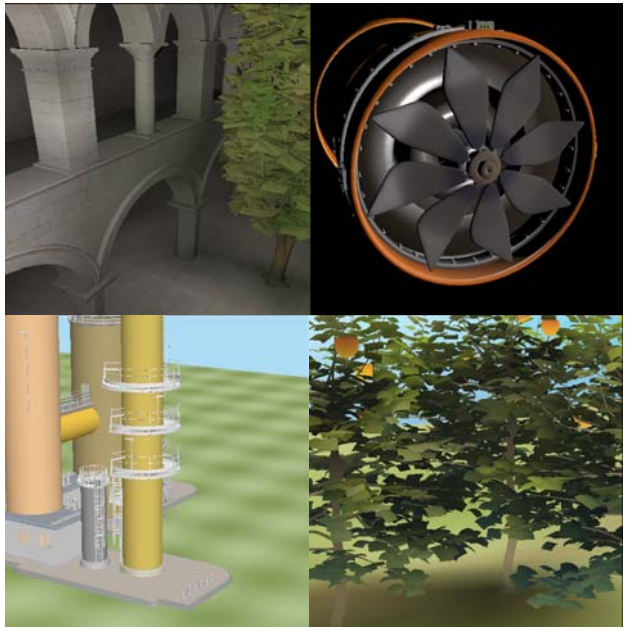

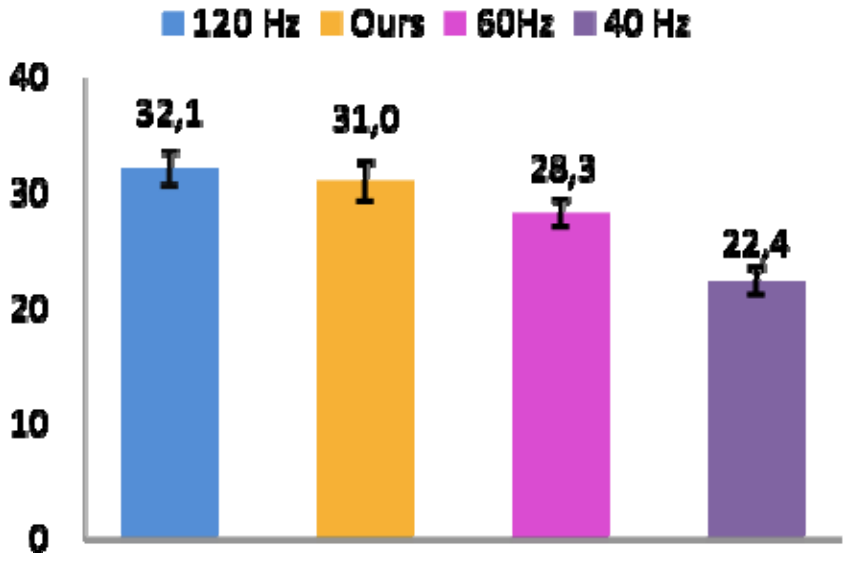

DELFT DATA SCIENCE

[Stengel, Bauszat, Eisemannx2, Magnor TVCG2015]

Ghosting in HFR Videos

60
fps

DELFT
DATA
SCIENCE

[Stengel, Bauszat, Eisemannx2, Magnor TVCG2015]

Blur mismatch between Camera and Eye

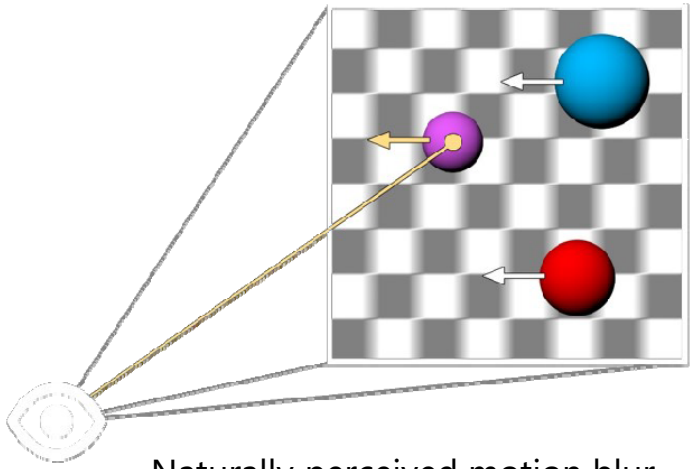

Naturally perceived motion blur

- eye tracking → high-detail
- static background is blurred
- which leads to ghosting for

'sharp' backgrounds over duration of a frame

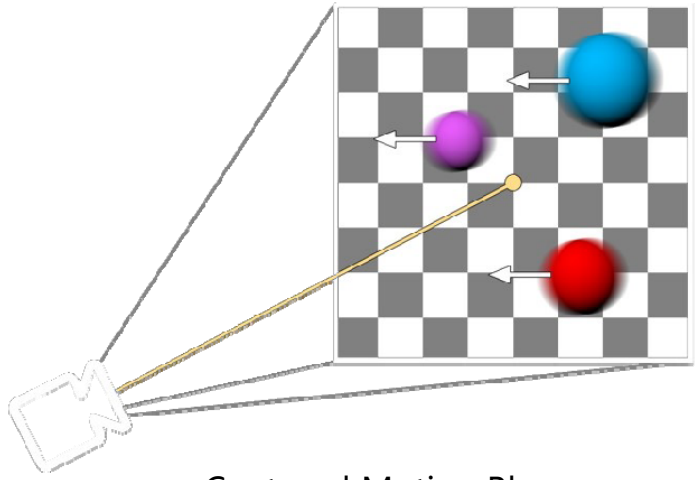

Captured Motion Blur

- static camera
- moving objects

DELFT DATA SCIENCE

[Stengel, Bauszat, Eisemannx2, Magnor TVCG2015]

Blur mismatch between Camera and Eye

[Stengel, Bauszat, Eisemannx2, Magnor TVCG2015]

Blur mismatch between Camera and Eye

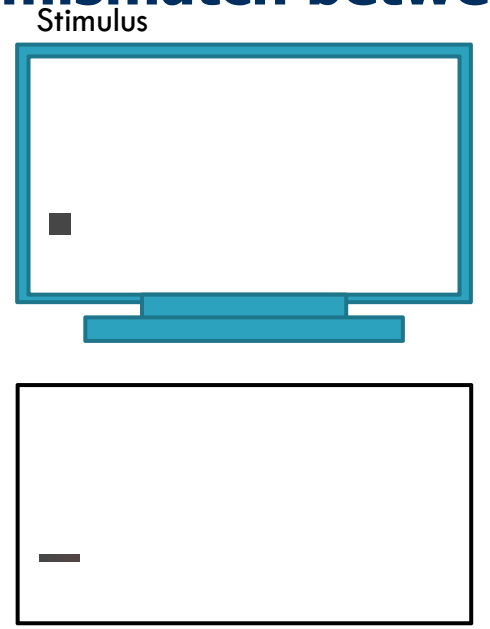

Motion blur recorded
by camera

[Stengel, Bauszat, Eisemannx2, Magnor TVCG2015]

Blur mismatch between Camera and Eye

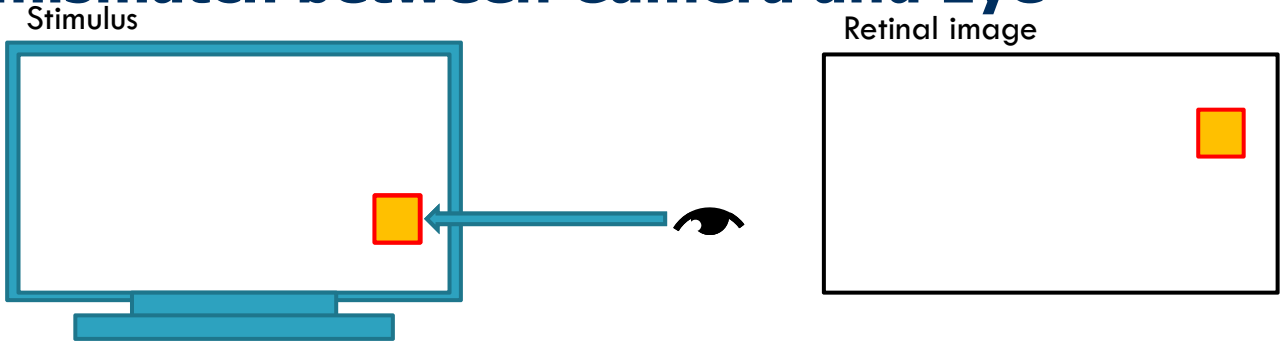

[Stengel, Bauszat, Eisemannx2, Magnor TVCG2015]

Blur mismatch between Camera and Eye

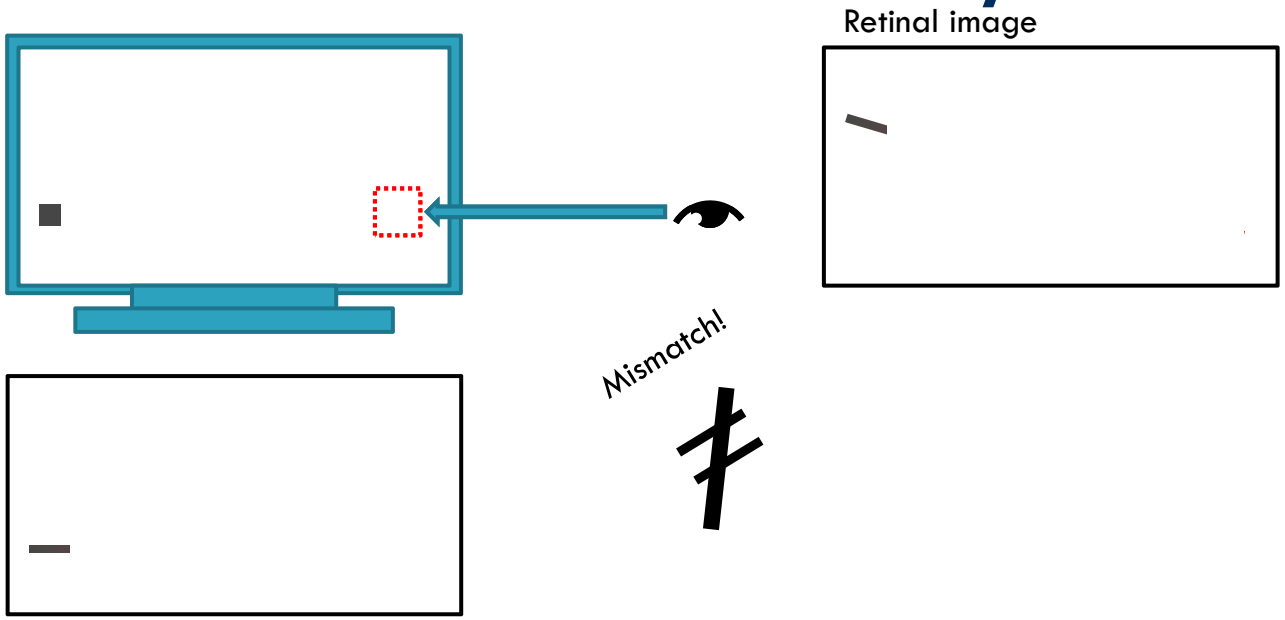

[Stengel, Bauszat, Eisemannx2, Magnor TVCG2015]

Blur mismatch between Camera and Eye

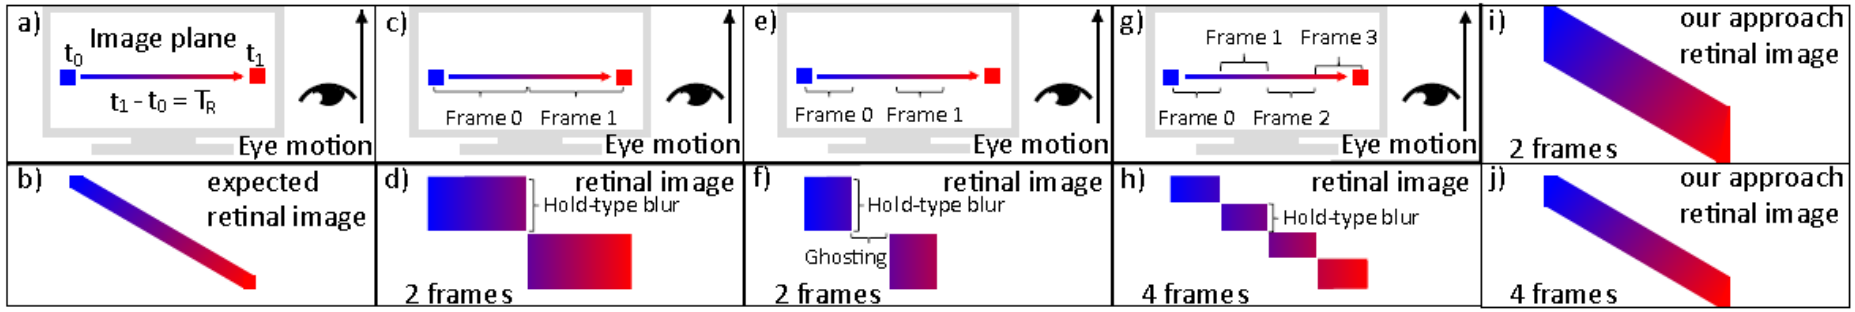

Please follow the face of the statue in the following videos.


Short exposure / 60 fps


Long exposure / 60 fps


Perceptual filter / 60 fps

DATA SCIENCE

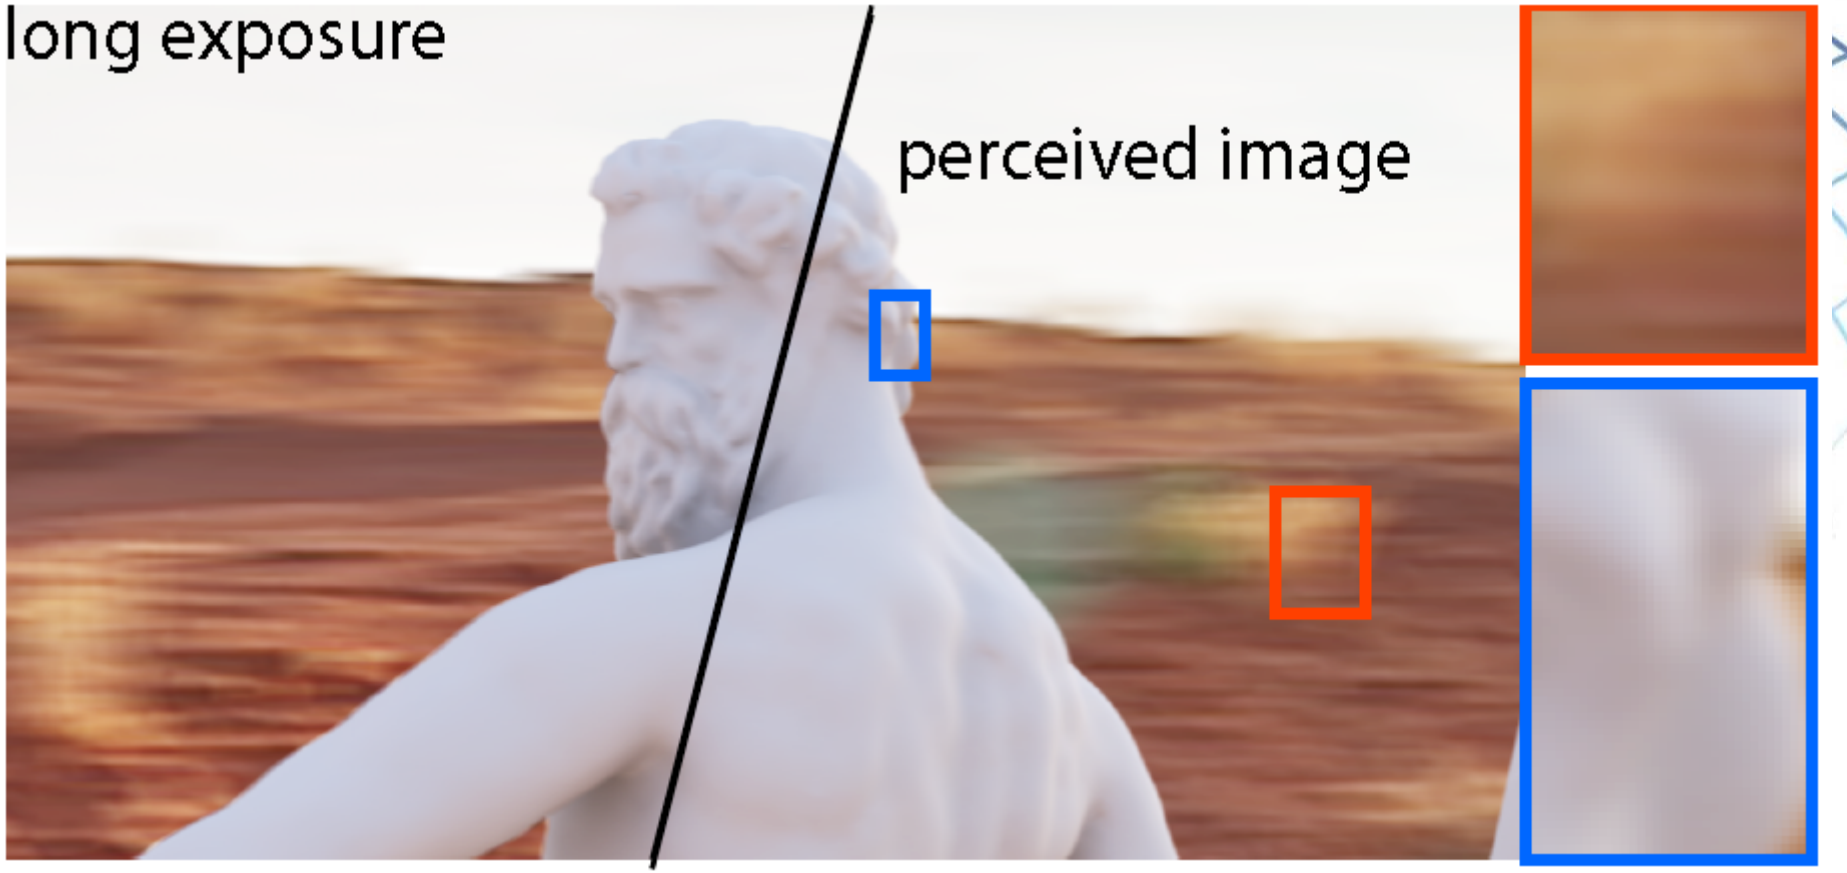

DATA SCIENCE

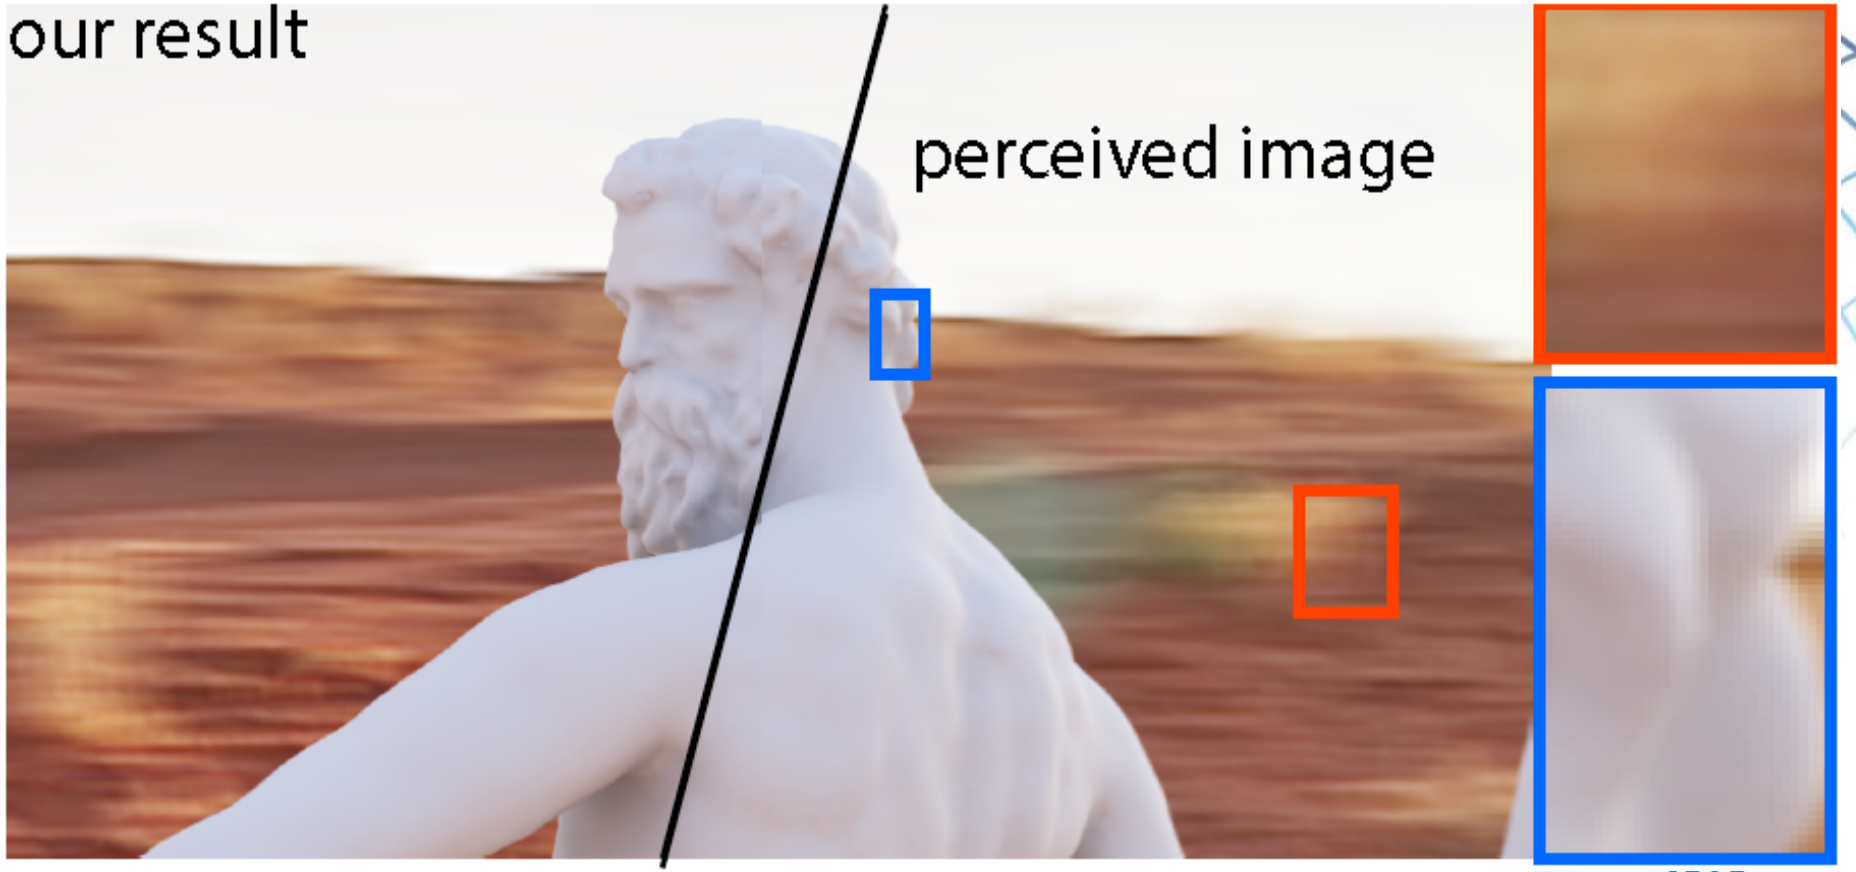

DATA SCIENCE

Can we use our perceptual filter to influence eye motion?

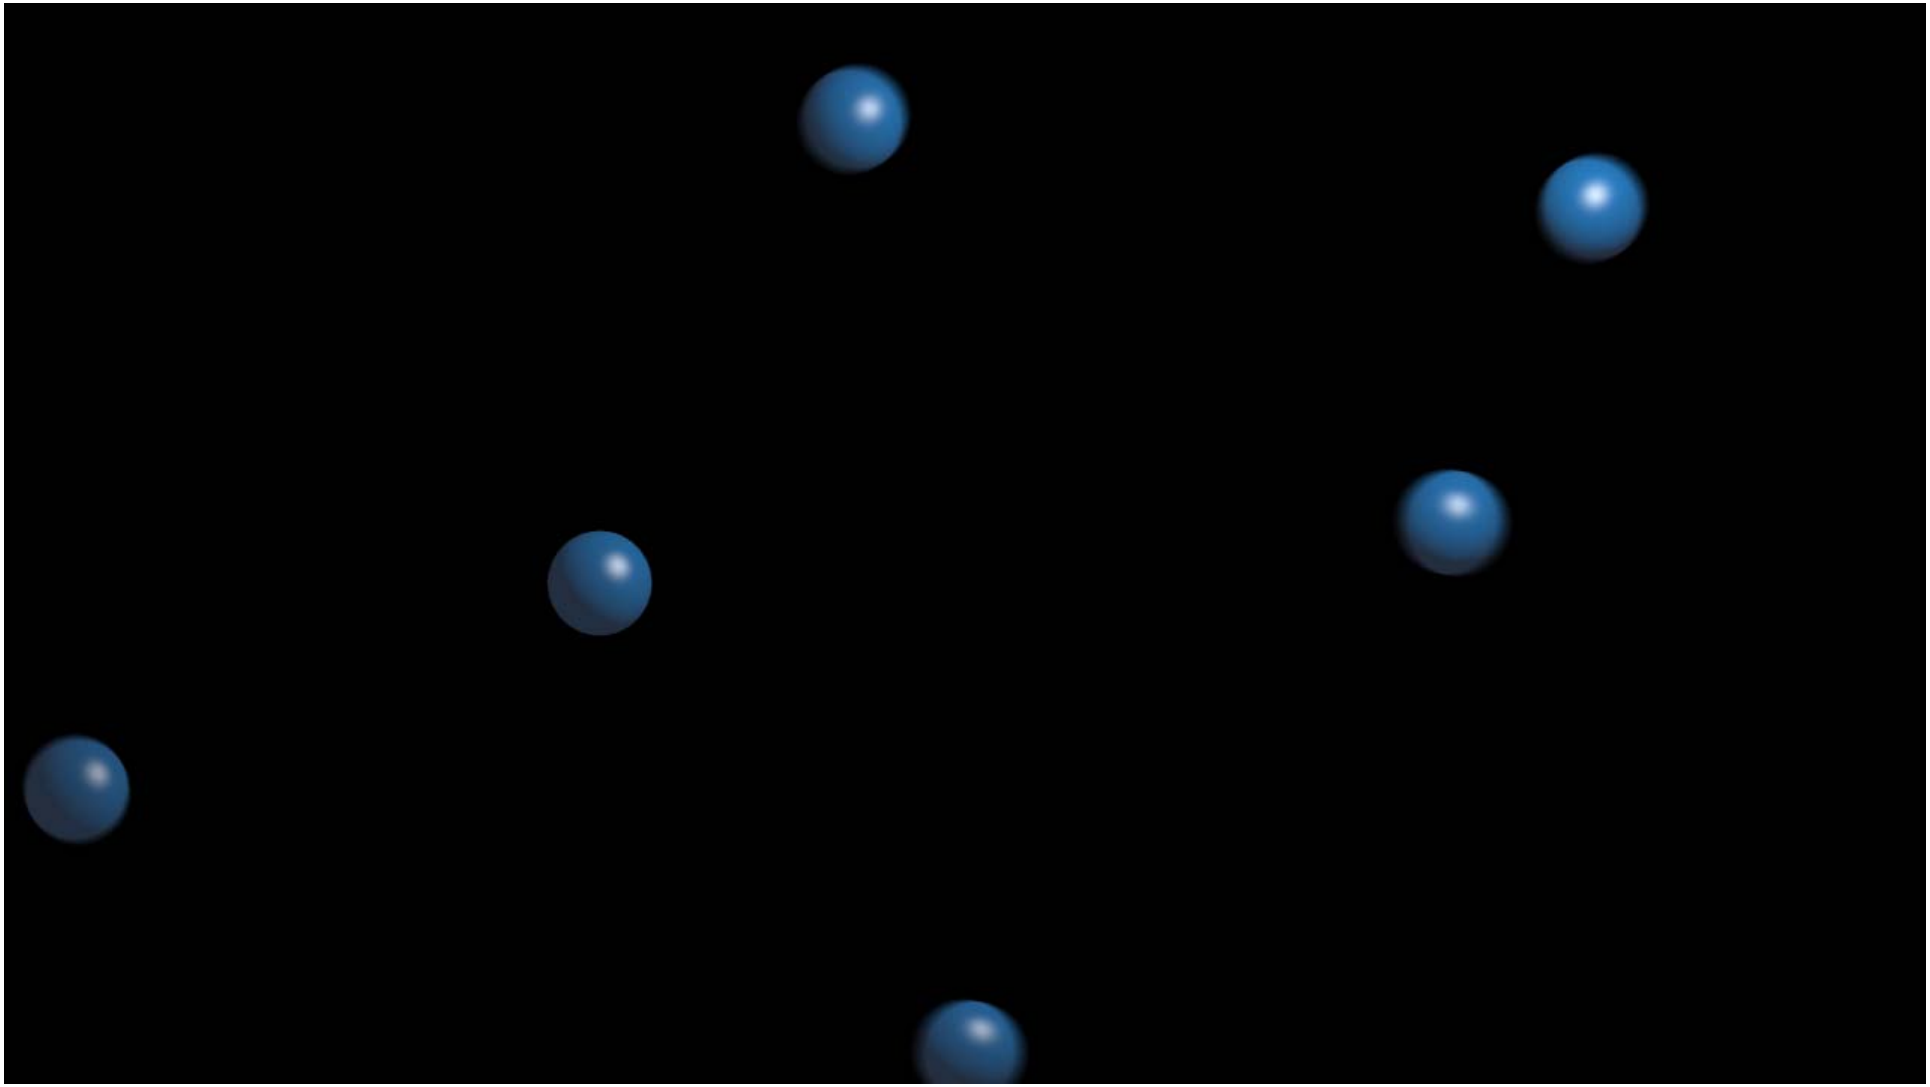

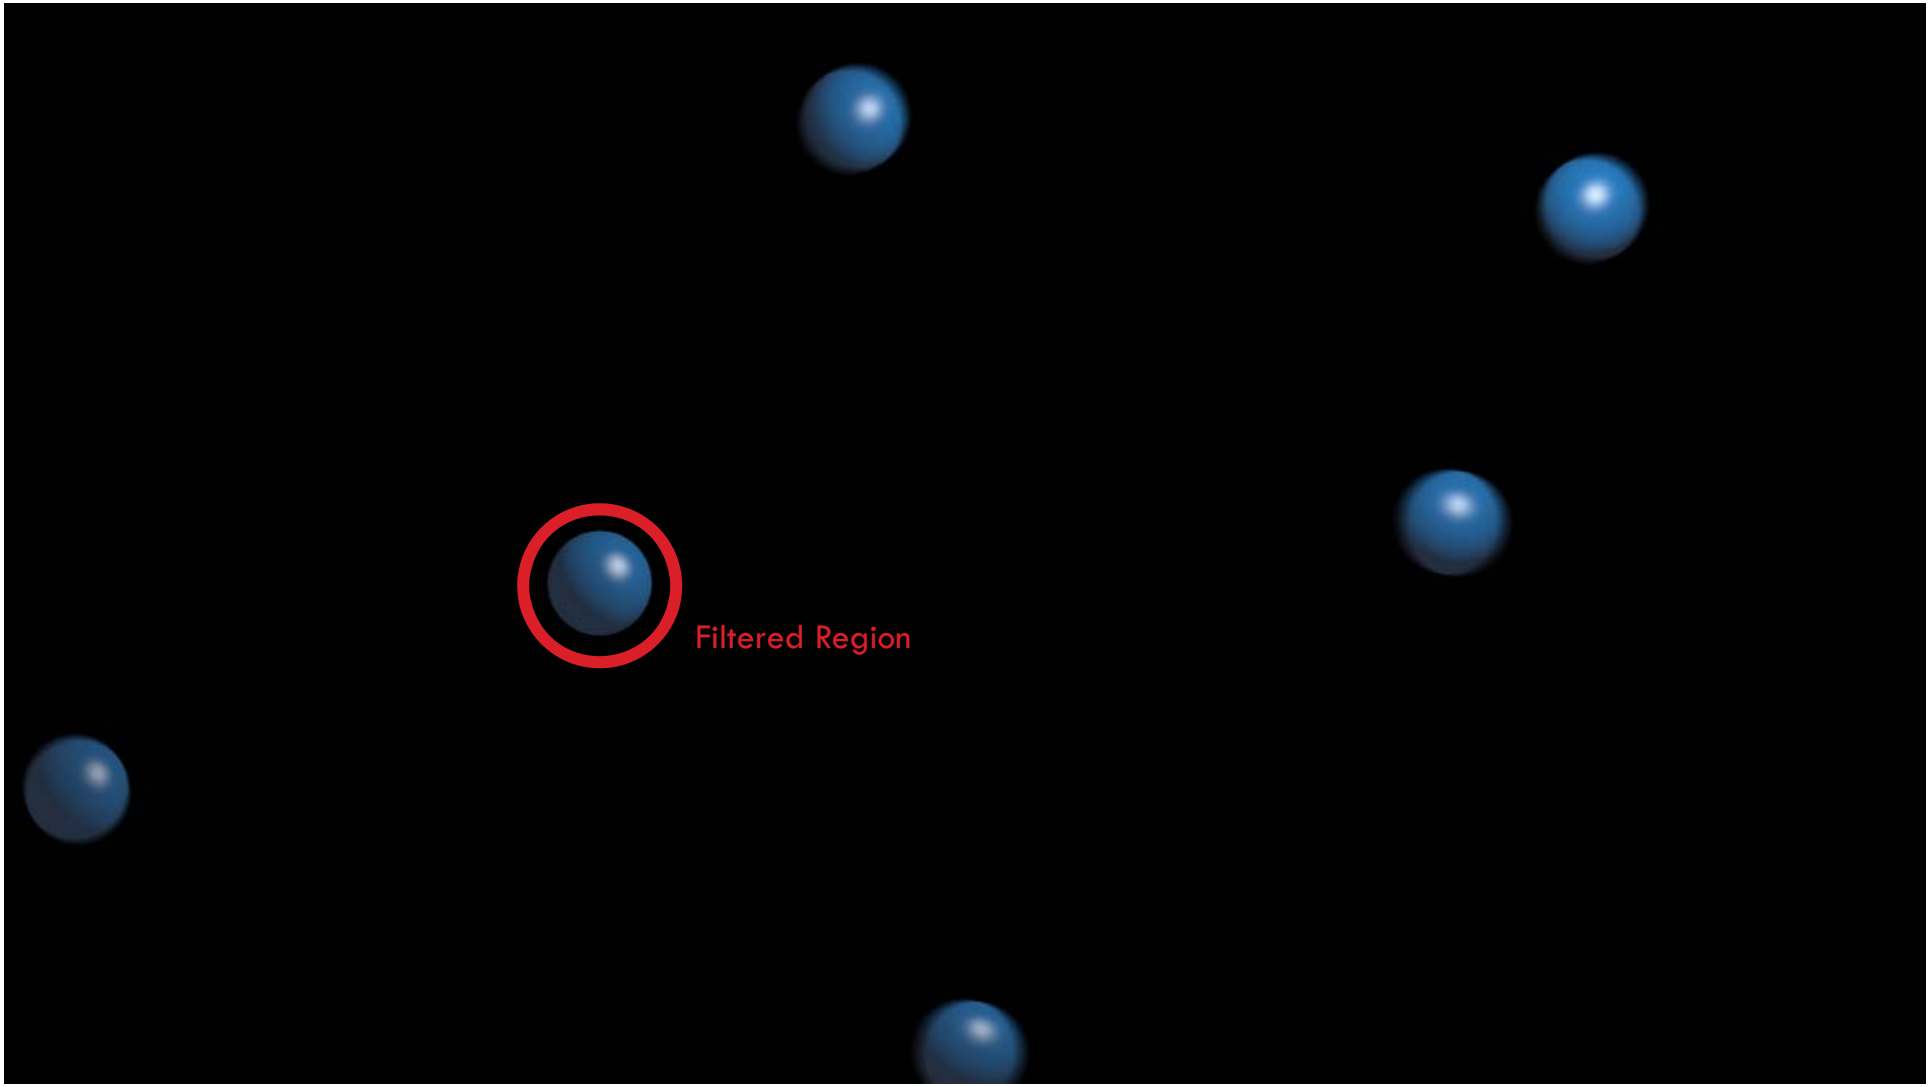
Filtered Region

Subtle Gaze Guidance

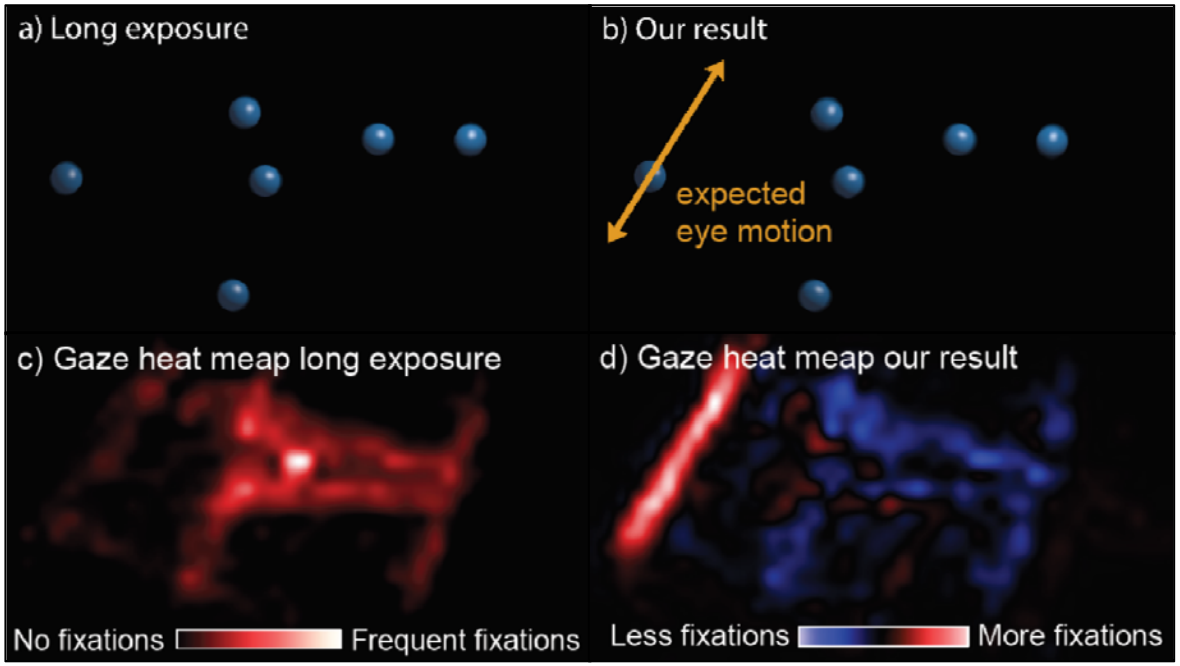

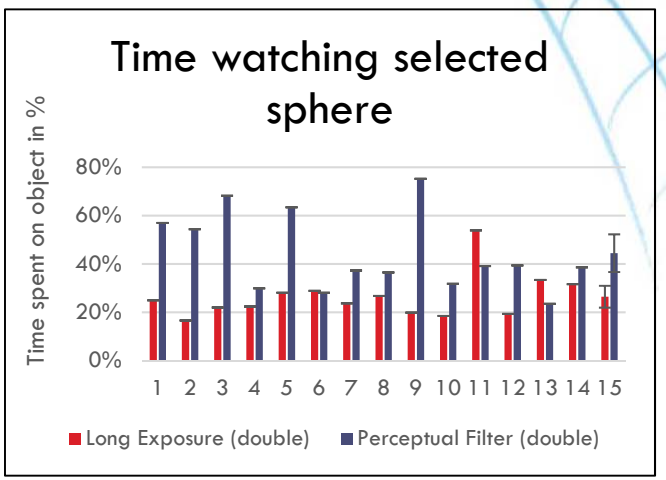

DELFT DATA SCIENCE


Love

- We can influence where you look!
- Can we influence what you see?

Overcoming Physical Limitations

- Apparent Resolution Enhancement
  [Didyk, Eisemann, Ritschel, Myszkowski, Seidel – SIGGRAPH 2010]
  [Templin, Didyk, Ritschel, Eisemann, Myszkowski, Seidel - SCCG 2011]
  Resolution
- A Computational Model of Afterimages
  [Ritschel & Eisemann - Eurographics'12]
  Brightness

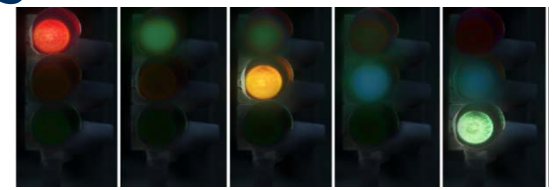

- A Perceptual Model for Disparity
  [Didyk, Ritschel, Eisemann, Myszkowski, Seidel - SIGGRAPH 2011/SIGAsia2012]
  Stereo illusion

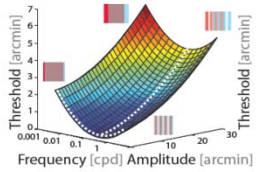

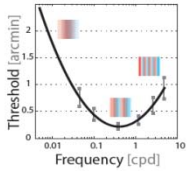

Questions?

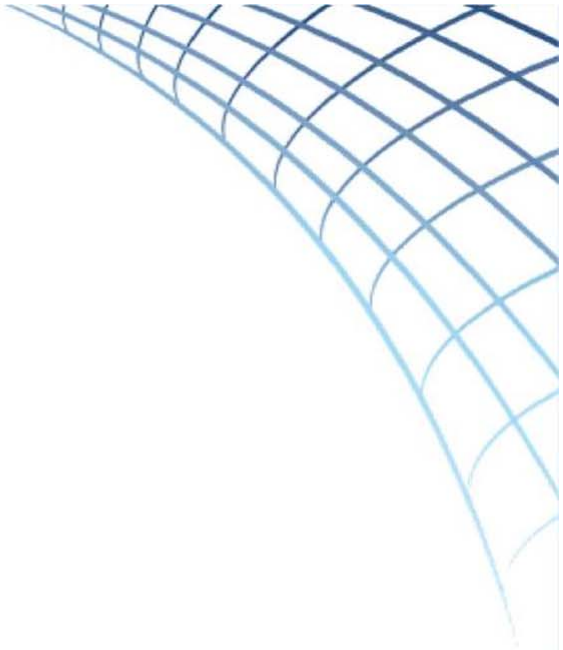

[Didyk, Eisemann, Ritschel, Myszkowski, Seidel – SIGGRAPH 2010]

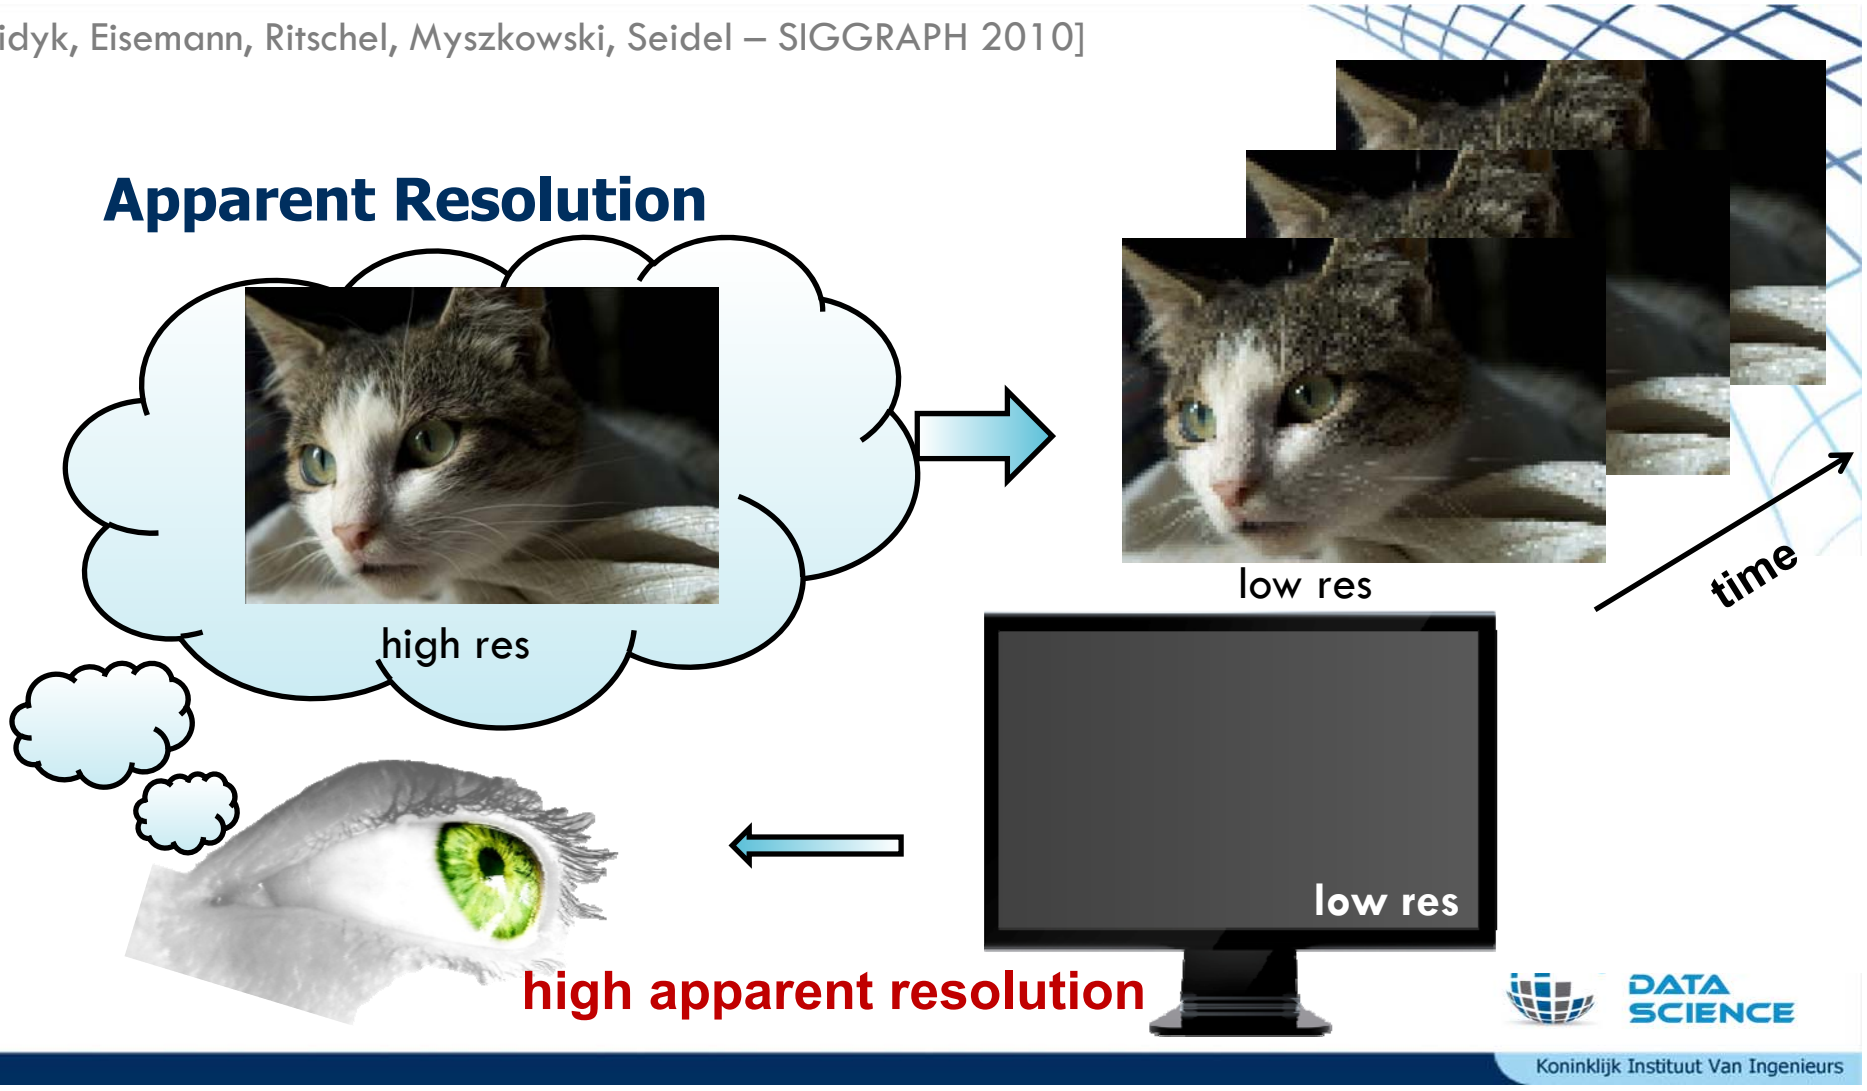

Temporal Domain

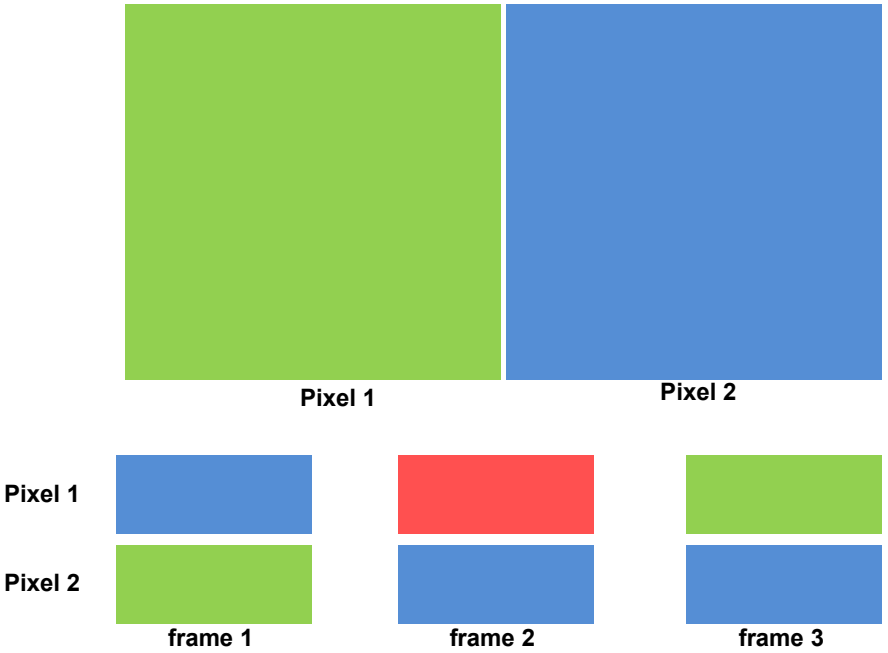

Temporal Domain – static case

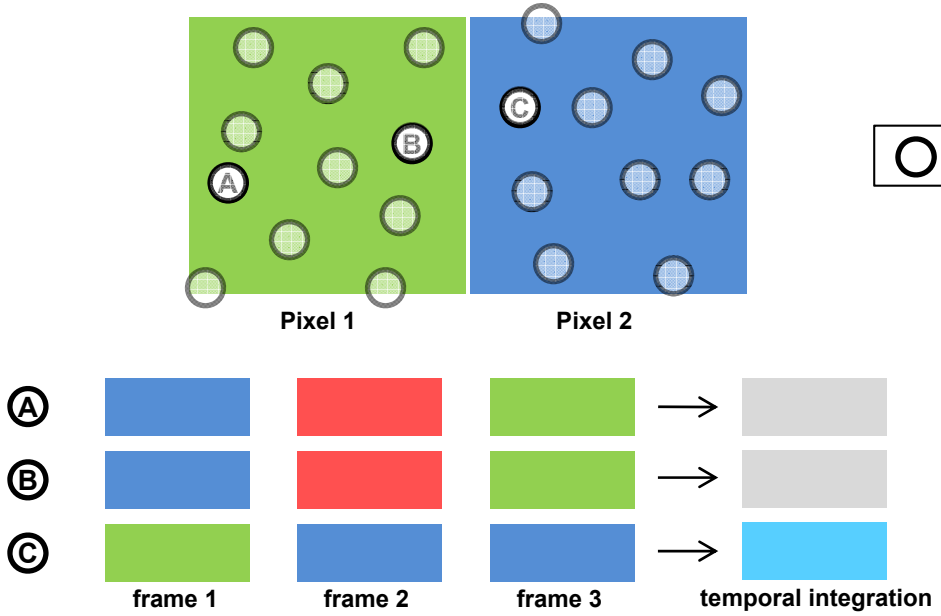

Temporal Domain – dynamic case

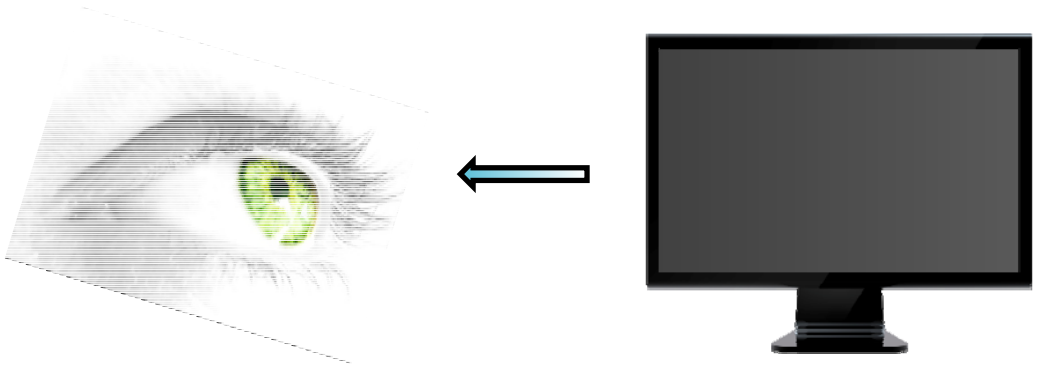

Temporal Domain – dynamic case

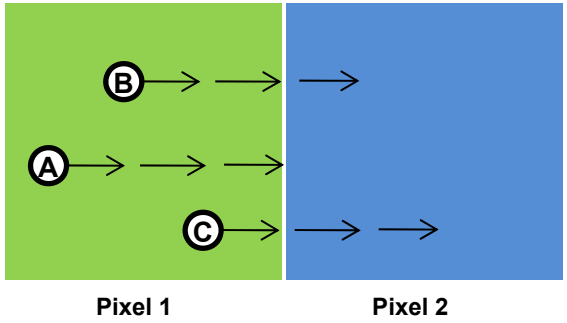

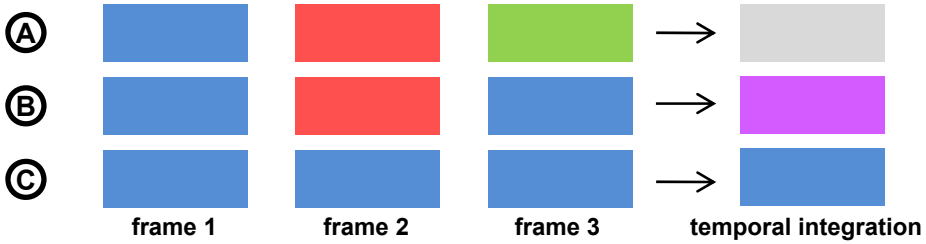

DELFT DATA SCIENCE

Temporal Integration Model

Receptor signal:

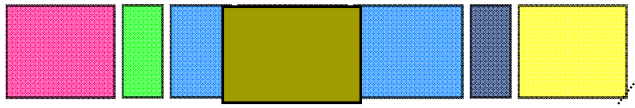

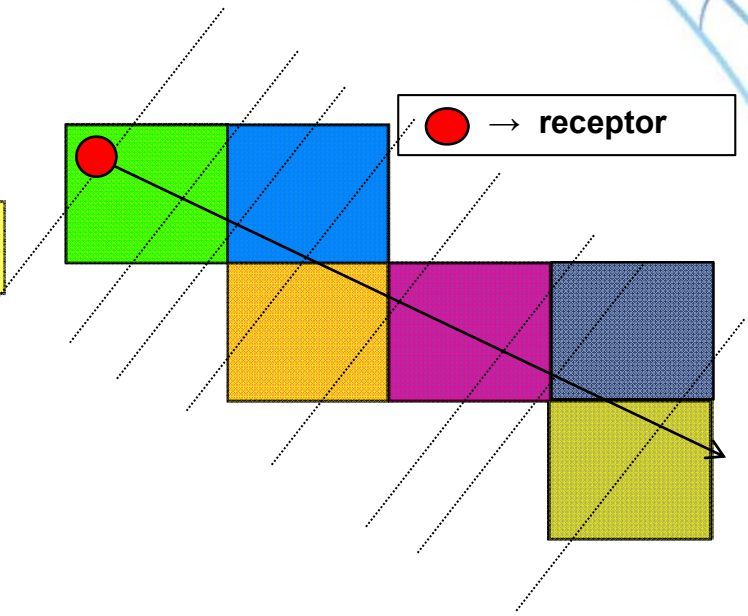

DELFT DATA SCIENCE

Temporal Integration Model

Receptor signal:

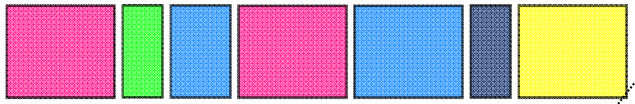

$$\sum_{i=0}^{N} w_i \, I(p(i), i)$$

segment length

time segment

Observed color for segment i

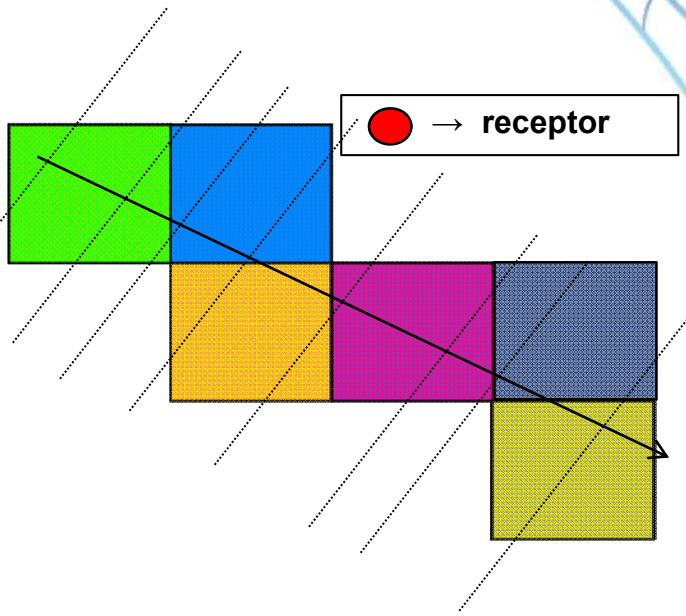

DELFT DATA SCIENCE

Prediction in Equations

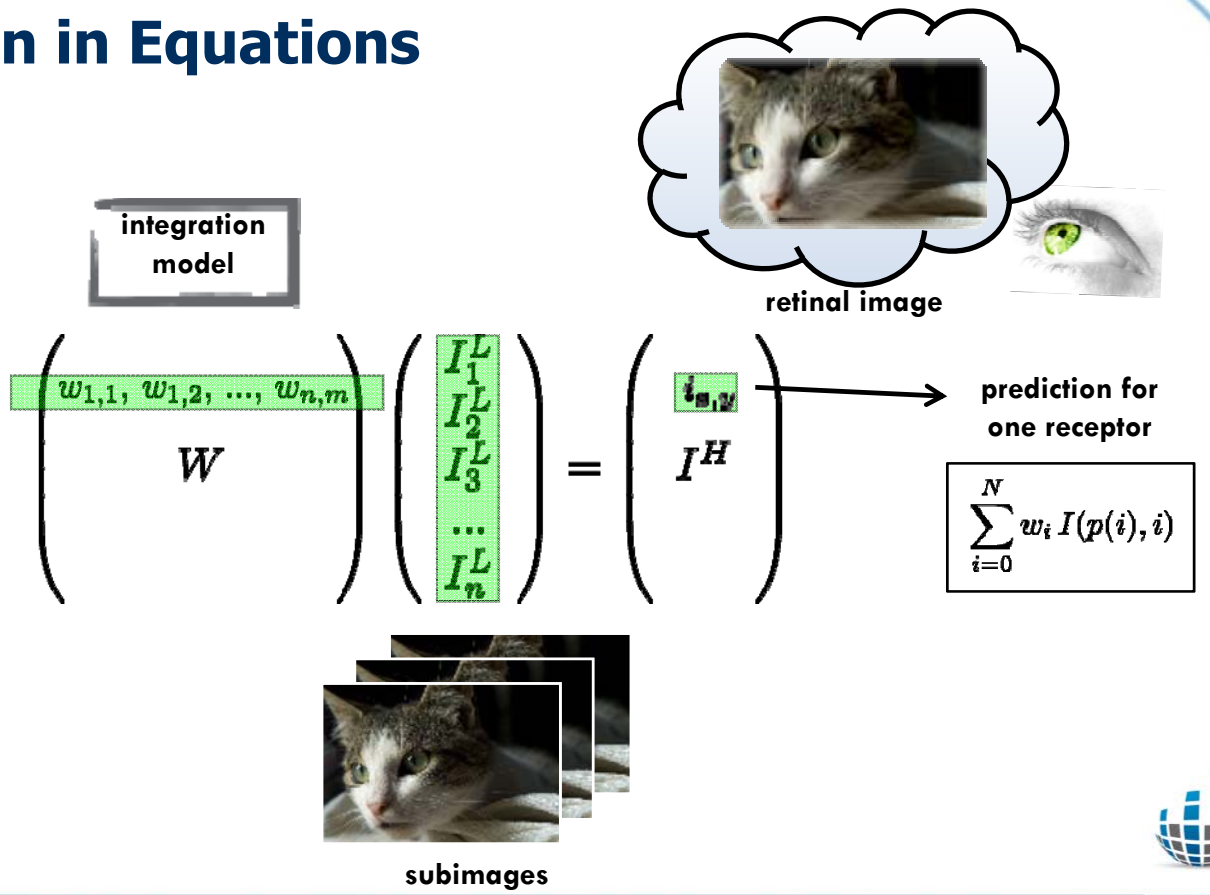

Optimization Problem

$$\begin{pmatrix} W \end{pmatrix} \begin{pmatrix} I_1^L \\ I_2^L \\ I_3^L \\ \dots \\ I_n^L \end{pmatrix} = \begin{pmatrix} I^H \end{pmatrix}$$

integration model

subimages

high resolution image

DELFT DATA SCIENCE

Prediction in Equations

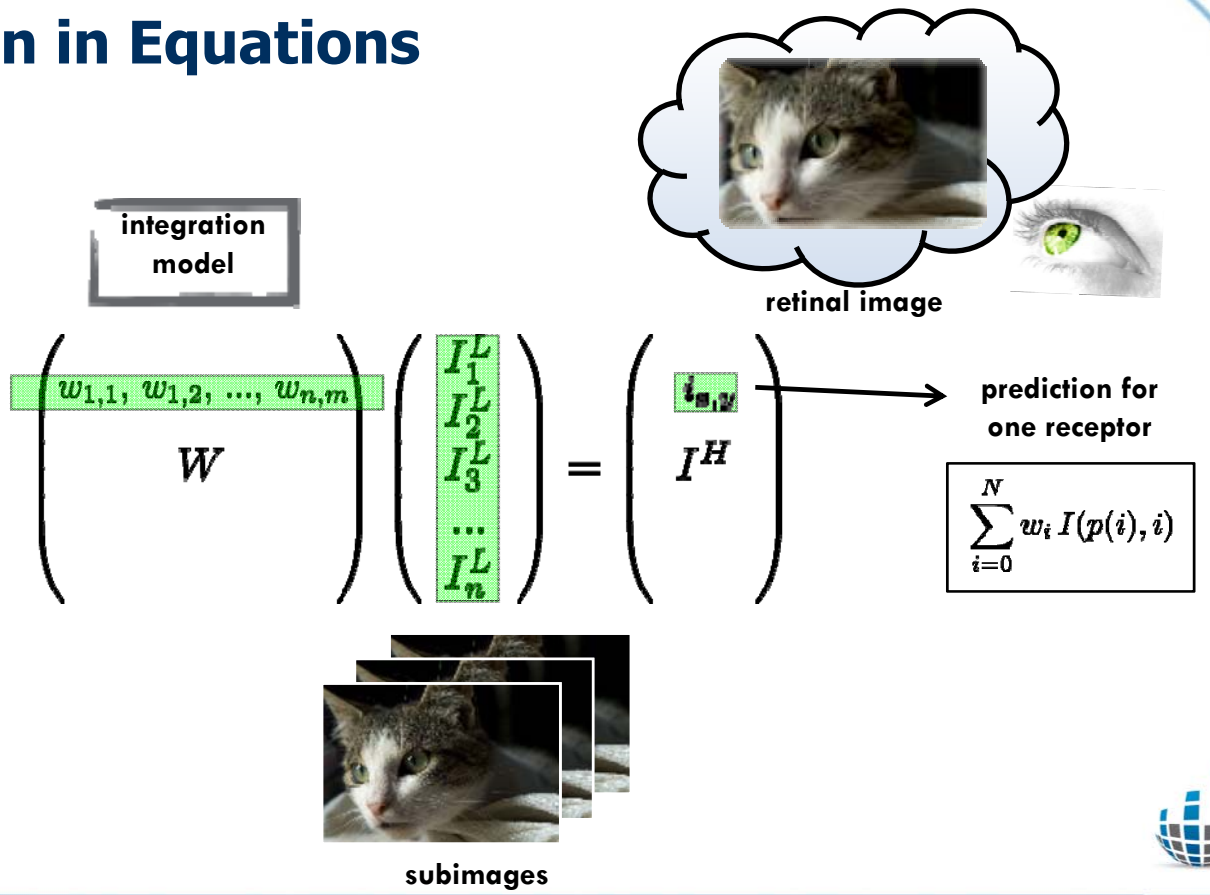

Optimization Problem

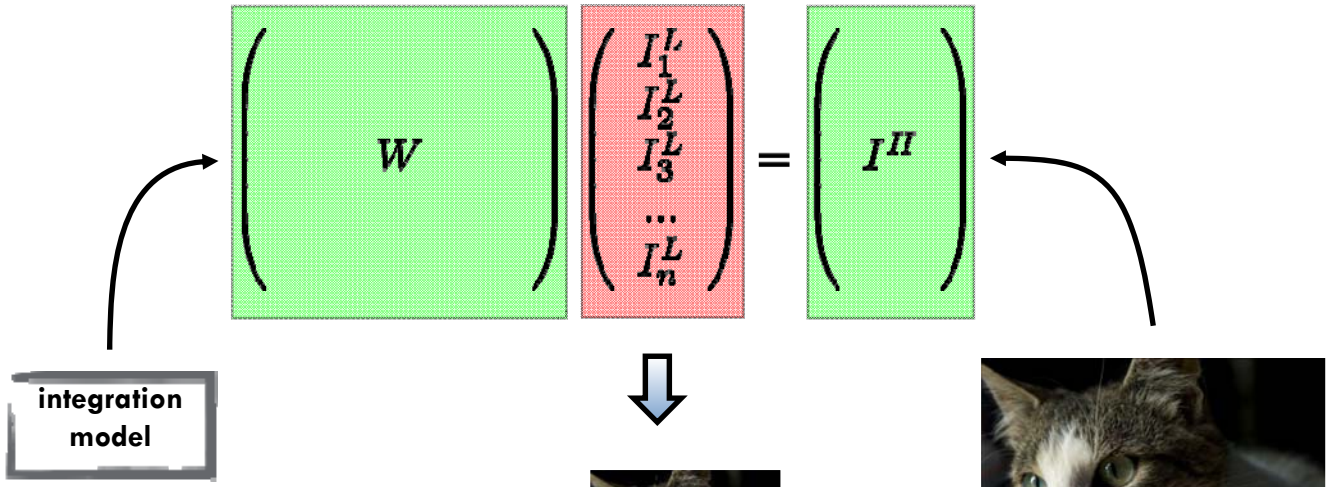

DELFT DATA SCIENCE

Fusion Frequency

- Depends on
  - temporal contrast
  - spatial extent

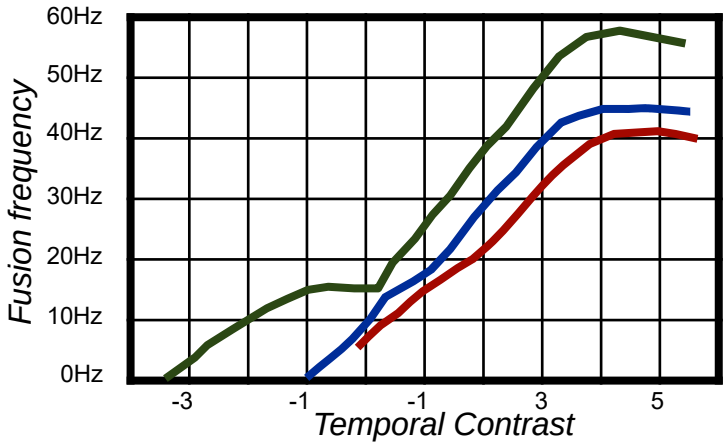

Critical Flicker Frequency - Hecht and Smith's data from
Brown J. L. *Flicker and Intermittent Simulation*

DELFT DATA SCIENCE

Optimization Result

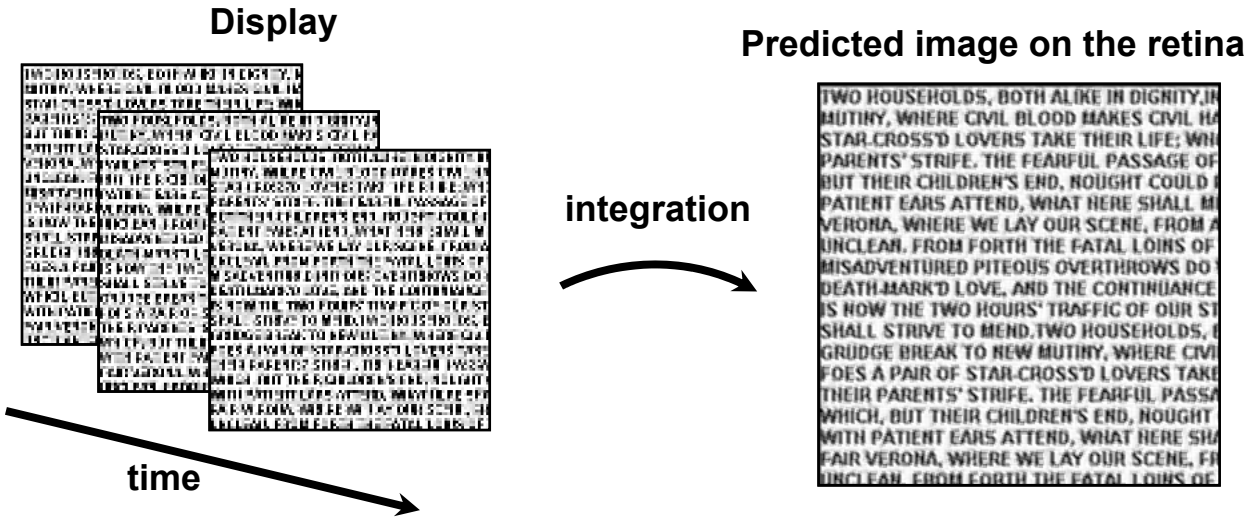

DELFT DATA SCIENCE

[Didyk, Eisemann, Ritschel, Myszkowski, Seidel – SIGGRAPH 2010]

Our vs. Previous Downsampling Techniques

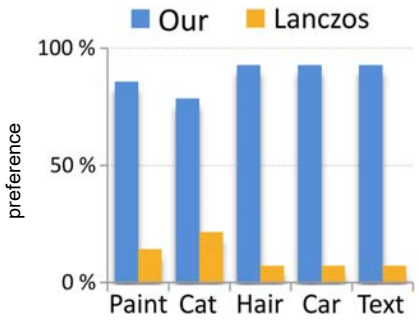

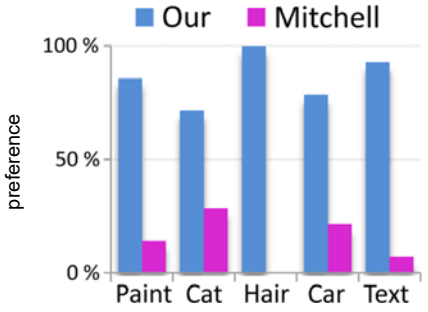

DELFT DATA SCIENCE

[Didyk, Eisemann, Ritschel, Myszkowski, Seidel – SIGGRAPH 2010]

Example: Alphabet

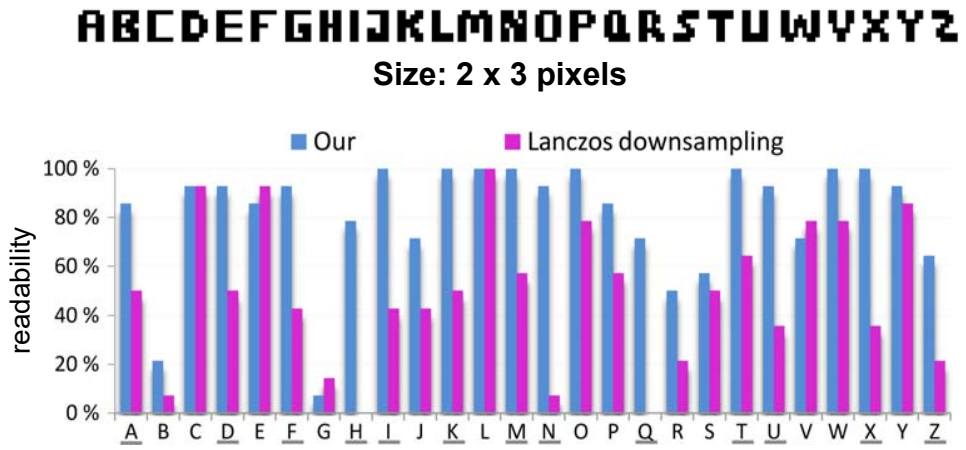

- Applications:
  - scrolling text or maps on low resolution devices
  - stock tickers, news headlines

Overcoming Physical Limitations

- Apparent Resolution Enhancement
  [Didyk, Eisemann, Ritschel, Myszkowski, Seidel – SIGGRAPH 2010]
  [Templin, Didyk, Ritschel, Eisemann, Myszkowski, Seidel - SCCG 2011]
  Resolution
- A Computational Model of Afterimages
  [Ritschel & Eisemann - Eurographics'12]
  Brightness

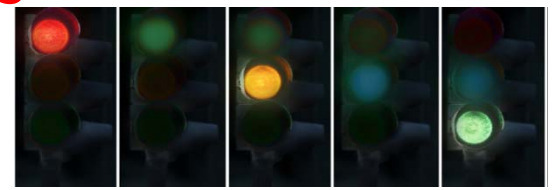

- A Perceptual Model for Disparity
  [Didyk, Ritschel, Eisemann, Myszkowski, Seidel - SIGGRAPH 2011/SIGAsia2012]
  Stereo illusion

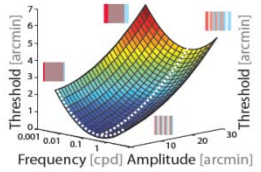

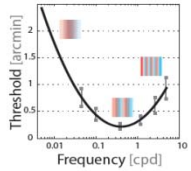

Example

- Focus on the dark spot…

[Ritschel and Eisemann - Eurographics2012]

Chromatic Adaptation

- Local bleaching

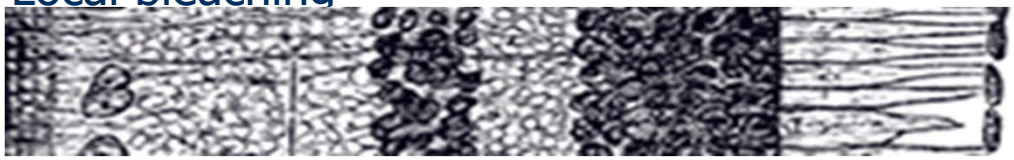

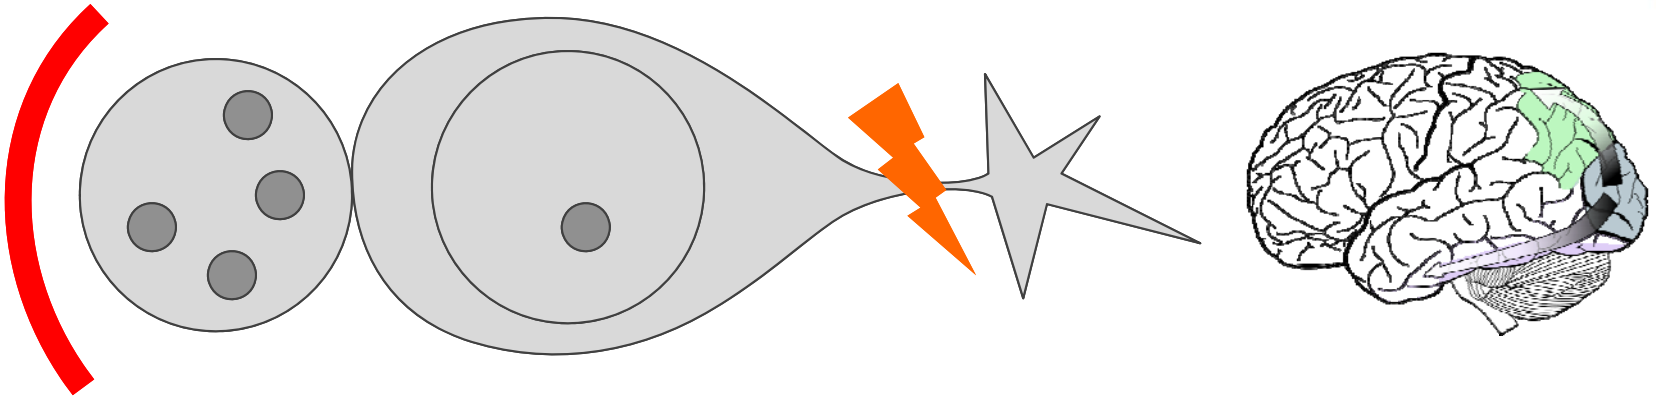

Adaptation state relates to the bleaching in the retinal cells

[Ritschel and Eisemann - Eurographics2012]

Overcoming Physical Limitations

- Apparent Resolution Enhancement
  [Didyk, Eisemann, Ritschel, Myszkowski, Seidel – SIGGRAPH 2010]
  [Templin, Didyk, Ritschel, Eisemann, Myszkowski, Seidel - SCCG 2011]
  Resolution
- A Computational Model of Afterimages
  [Ritschel & Eisemann - Eurographics'12]
  Brightness

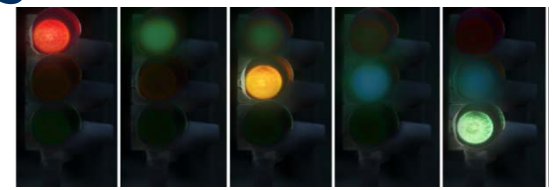

- A Perceptual Model for Disparity
  [Didyk, Ritschel, Eisemann, Myszkowski, Seidel - SIGGRAPH 2011/SIGAsia2012]
  Stereo illusion

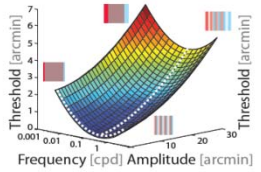

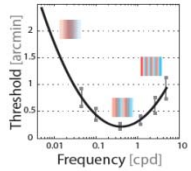

Rebirth of Stereo

St

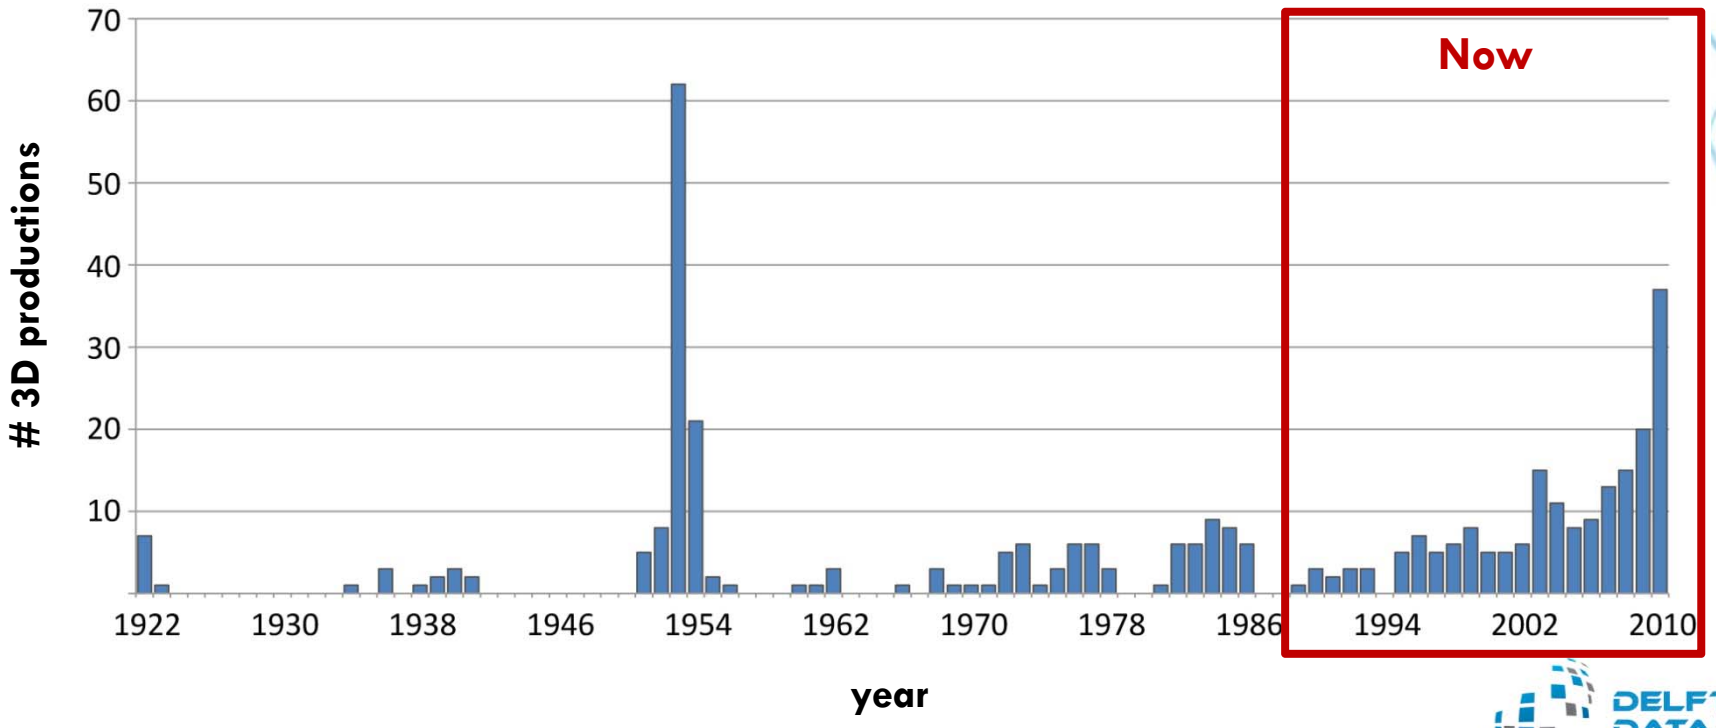

DELFT DATA SCIENCE

Stereo Vision

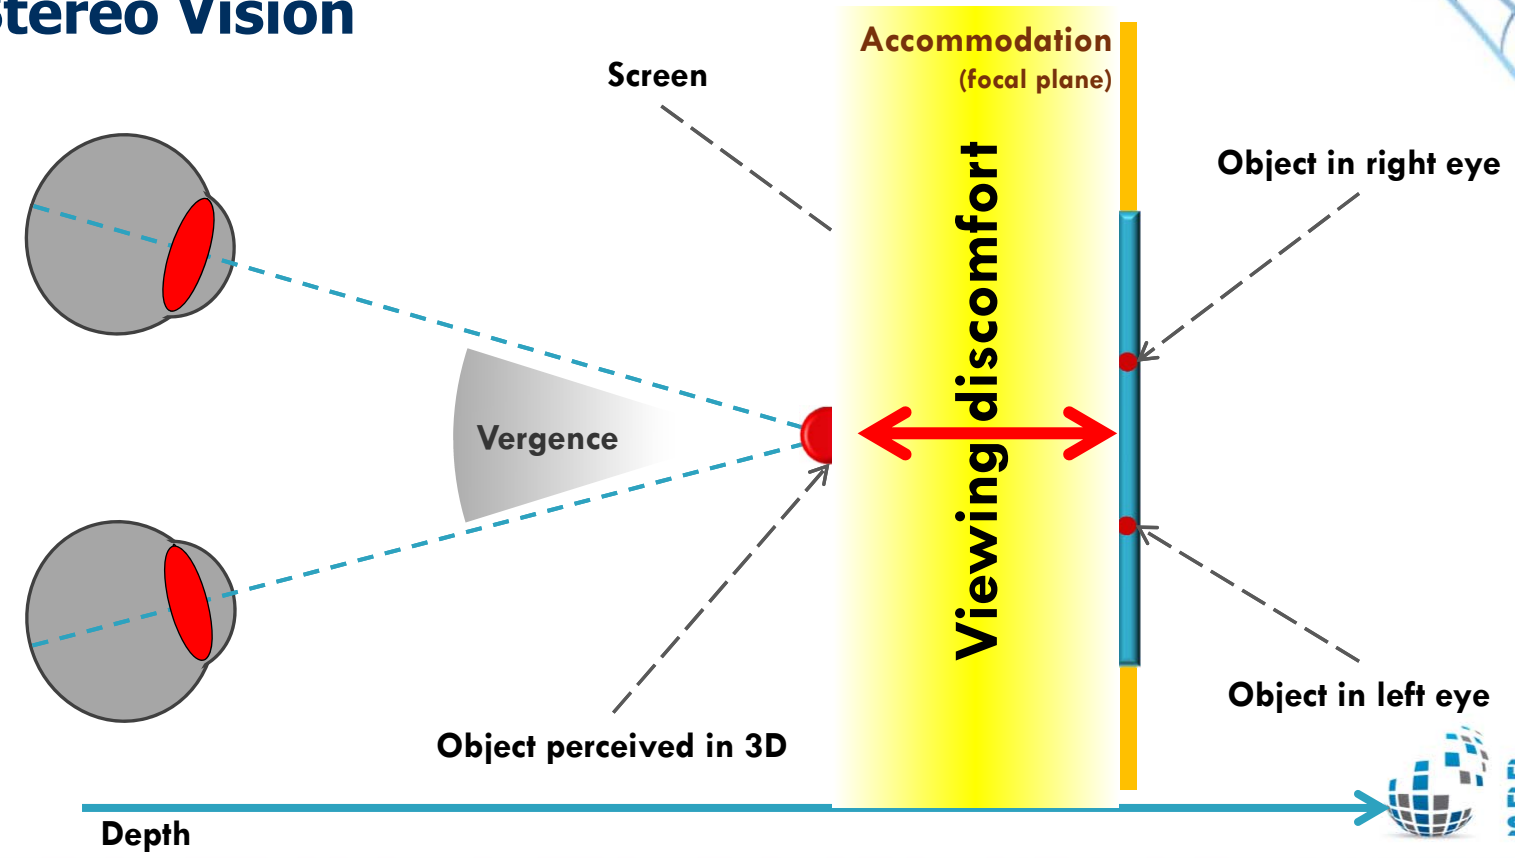

DELFT DATA SCIENCE

Stereo Vision

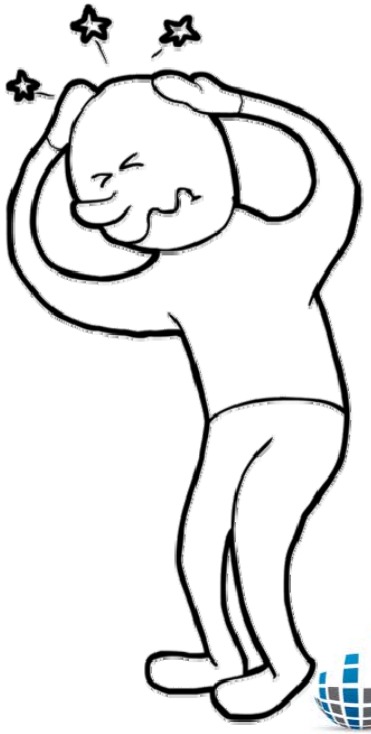

DELFT DATA SCIENCE

Stereo Vision

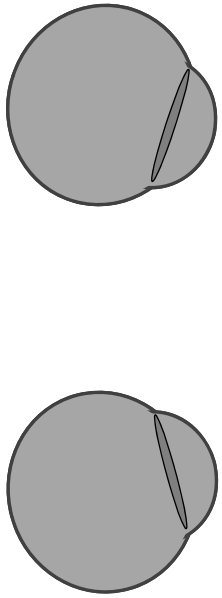

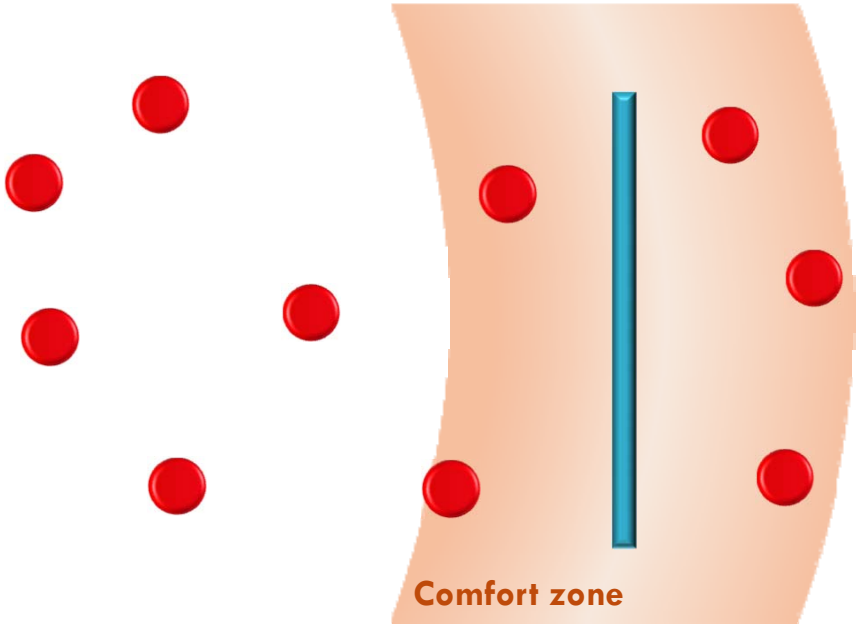

Scene manipulation

~~Viewing discomfort~~ → Viewing comfort

DELFT DATA SCIENCE

[Didyk, Ritschel, Eisemann, Myszkowski, Seidel – SIGGRAPH 2011]
[Didyk, Ritschel, Eisemann, Myszkowski, Seidel – SIGGRAPH Asia 2012]

Perception Model

Parameter space:

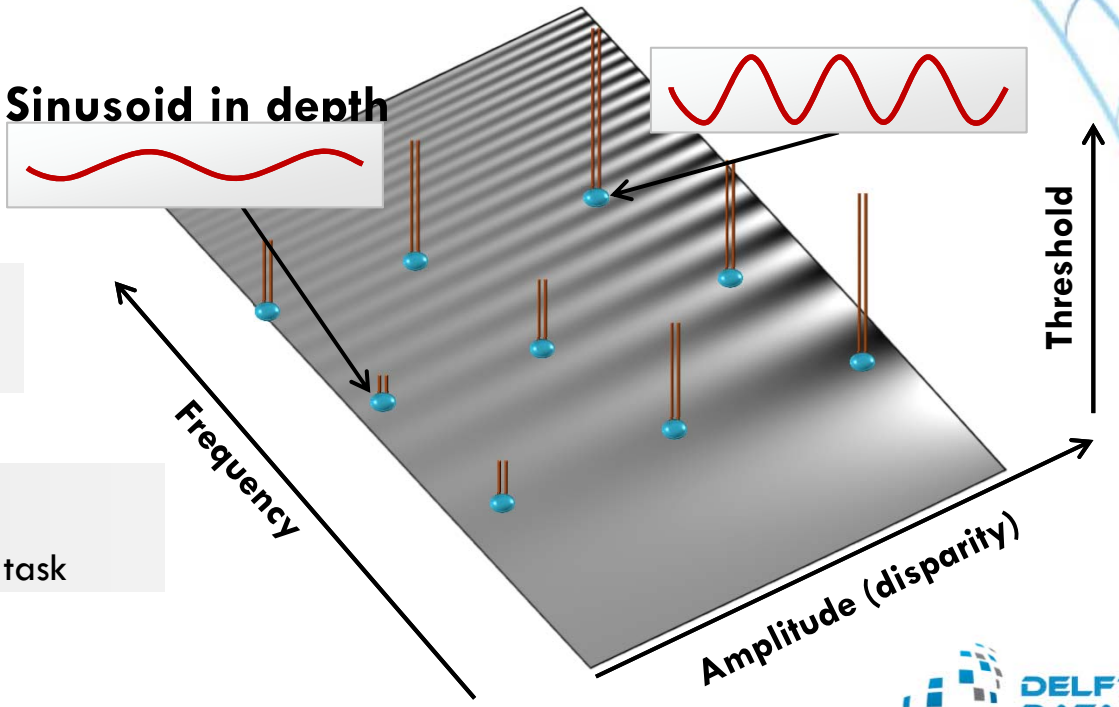

Sample the space

Measure thresholds

- Experiment with adjustment task

DELFT DATA SCIENCE

Measurements

Thresholds measurement:

- Two sinusoidal corrugations
- Which one has more depth? (left/right)
- Amplitude adjustment (PEST with 2AFC)

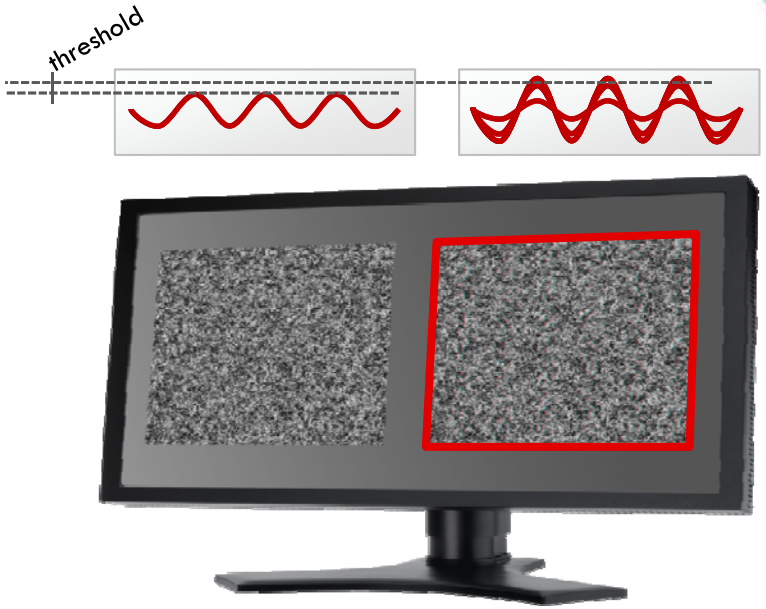

DELFT DATA SCIENCE

[Didyk, Ritschel, Eisemann, Myszkowski, Seidel – SIGGRAPH 2011]
[Didyk, Ritschel, Eisemann, Myszkowski, Seidel – SIGGRAPH Asia 2012]

Stereo Vision – Metric

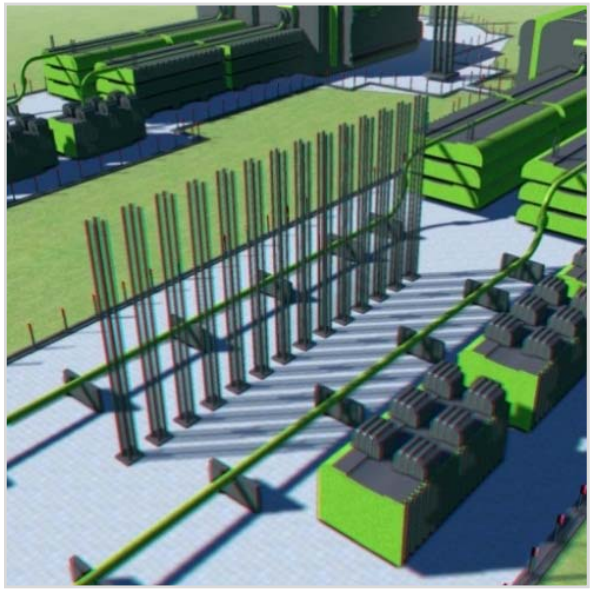

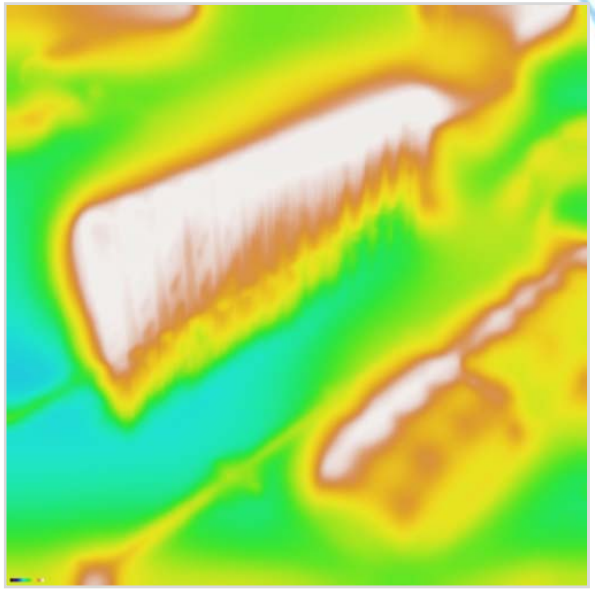

DELFT DATA SCIENCE

[Didyk, Ritschel, Eisemann, Myszkowski, Seidel – SIGGRAPH 2011]
[Didyk, Ritschel, Eisemann, Myszkowski, Seidel – SIGGRAPH Asia 2012]

Stereo Vision – Metric

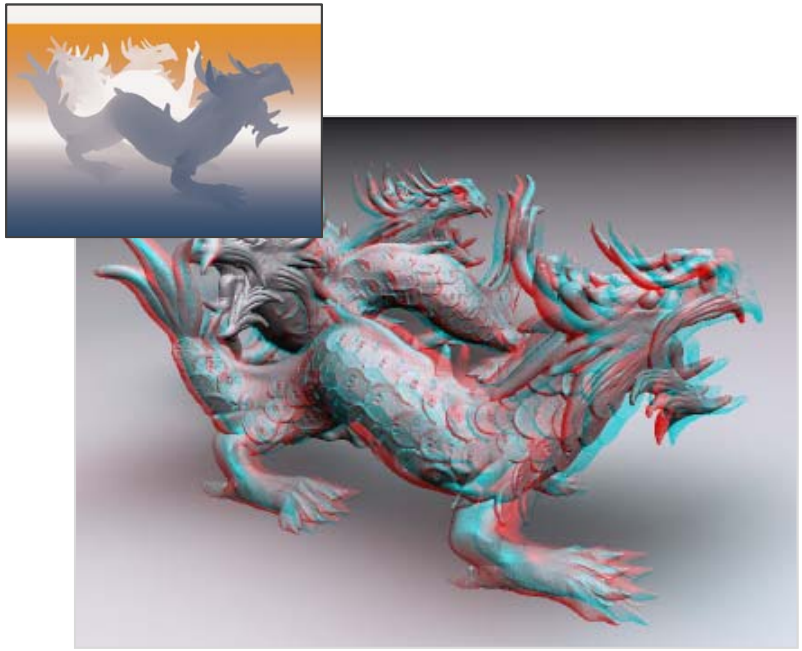

Standard stereo

Backward-compatible stereo

DELFT DATA SCIENCE

Perceptual Rendering

Taking perception into account

- can reduce computational cost
- can increase quality/comfort
- is advantageous even on modern displays

[Pajak, Herzog, Eisemann, Myszkowski, Seidel - EG 2014]

Streaming of Rendered Content with Depth Compression

H264

Our solution + depth

[Pajak, Herzog, Eisemann, Myszkowski, Seidel - EG 2014]

Remote Computing

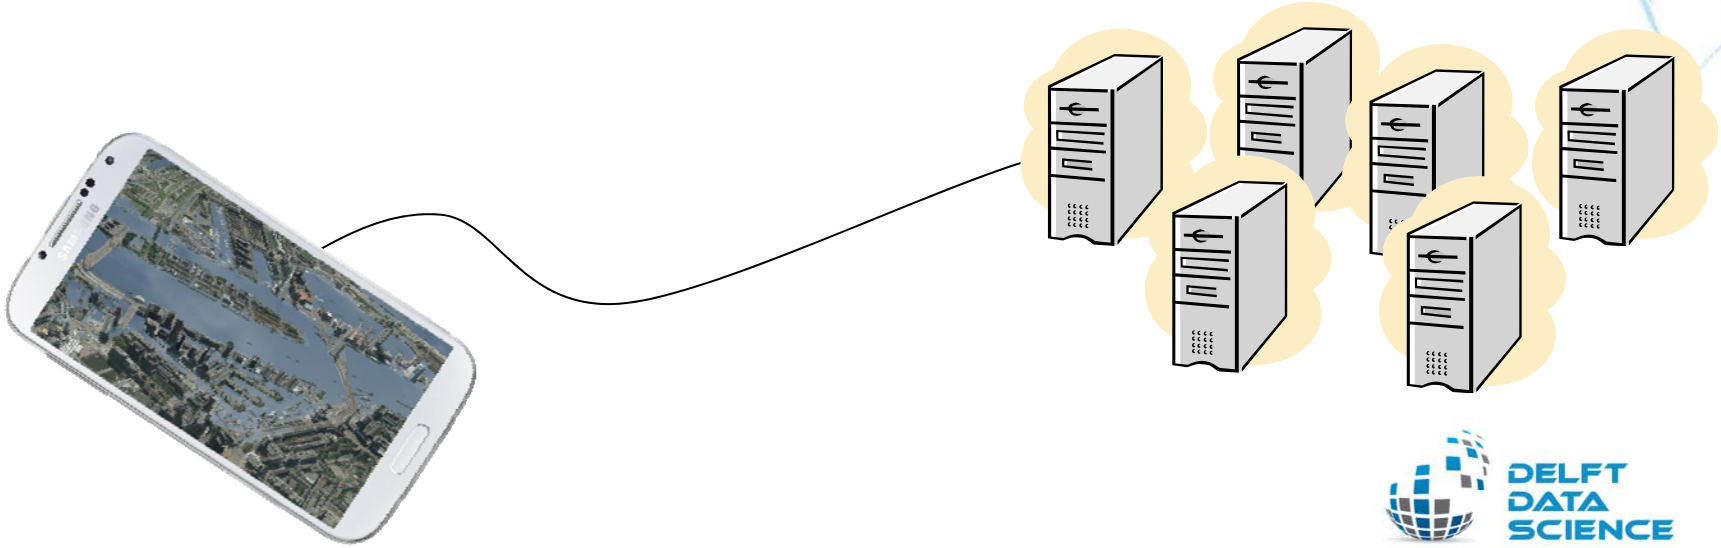

Enlighten your Research
SurfSara

Remote Computing

General Image Synthesis Solution

DELFT DATA SCIENCE

Remote Computing

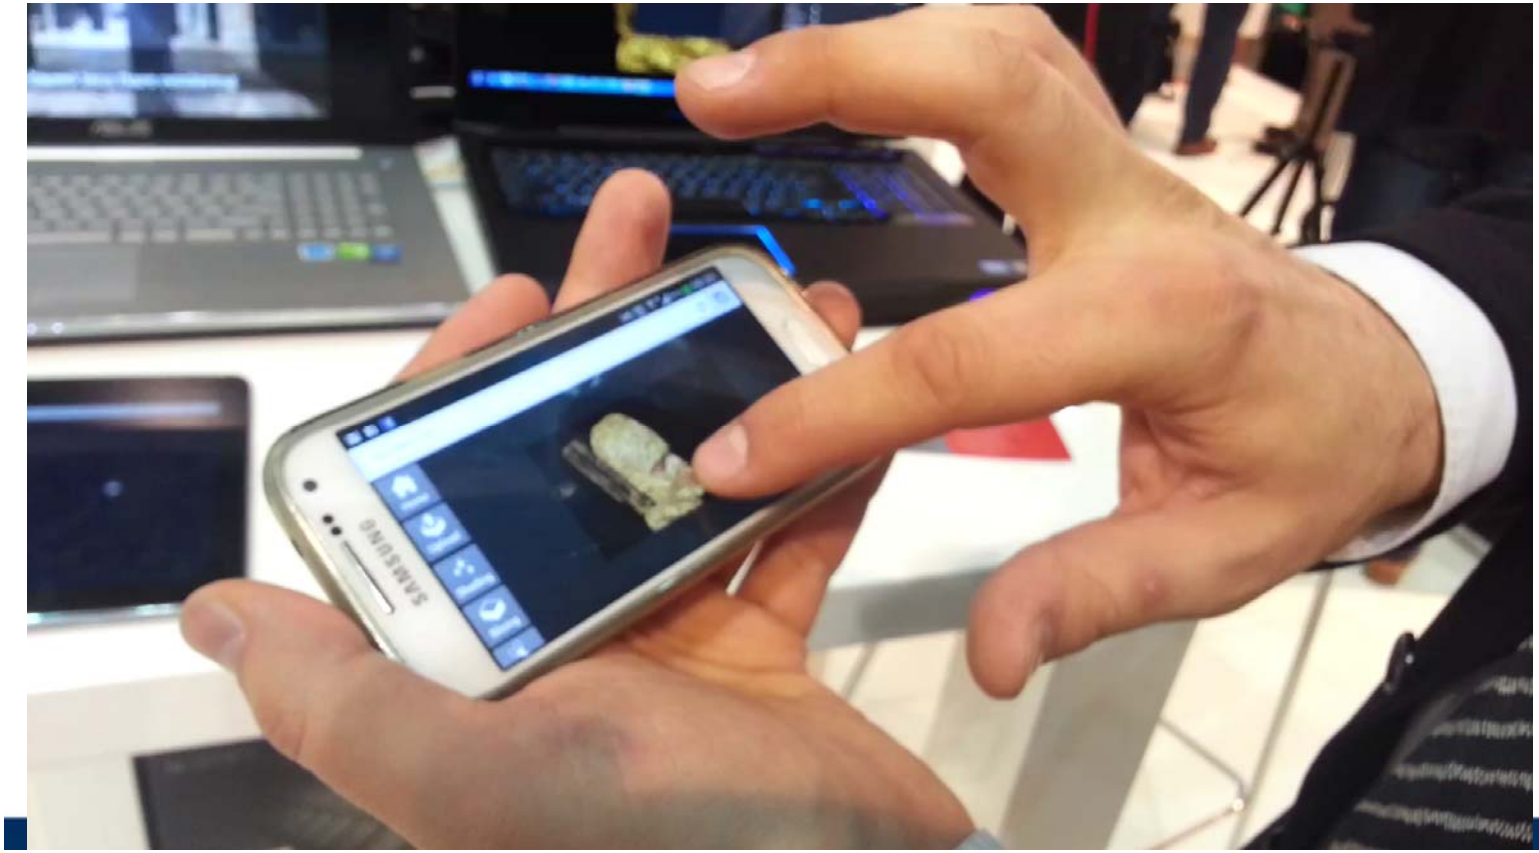

DELFT DATA SCIENCE

Koninklijk Instituut Van Ingenieurs

[Tutenel et al. GeoWorldForum 2013]

Interaction: Highlighting and Markup

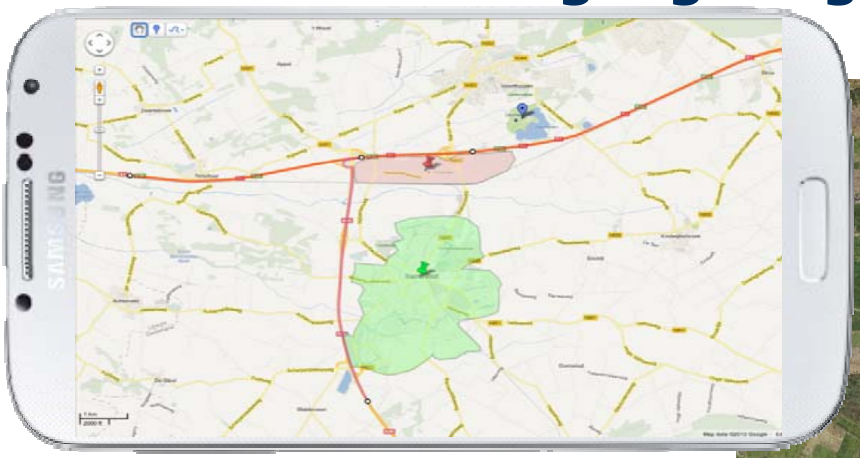

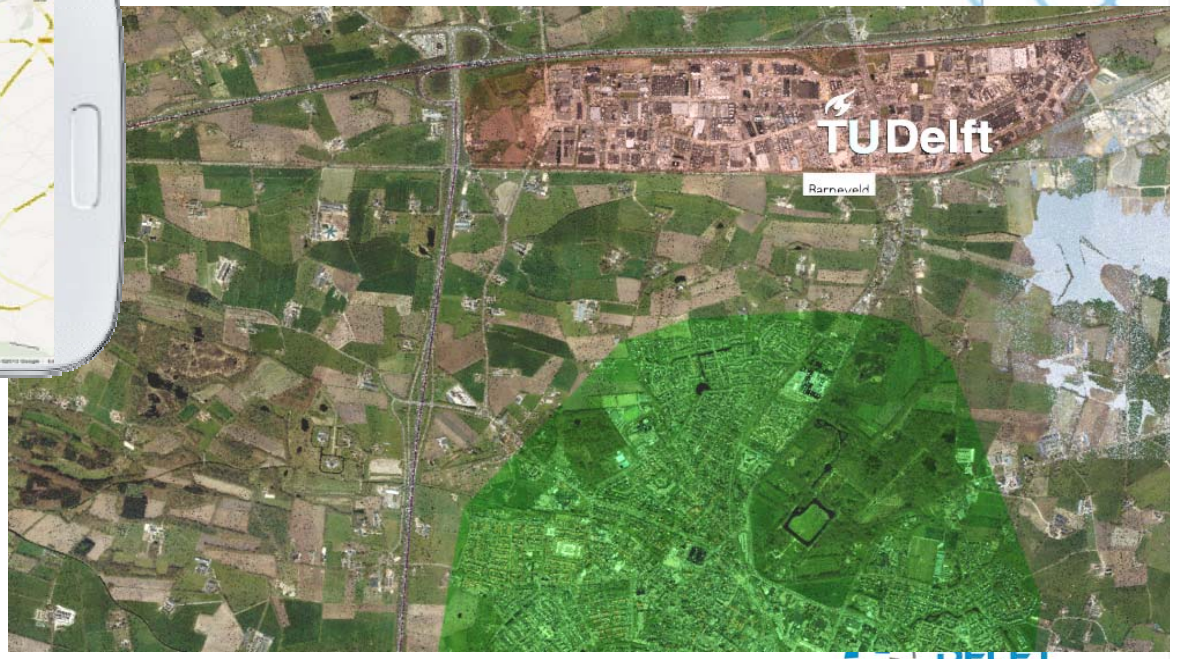

Example: Improved Navigation

- Simplified representation for recognition
- Visual Context for orientation

Grabler et al. 2008

Example: Improved Navigation – Existing solutions

Top-down

Street-level

Bird's eye

3D

Combined

DELFT DATA SCIENCE

[Kol et al. VMV 2014]

Example: Improved Navigation

- Give a Context
- Optimize for building recognition
  - Canonical View
  - Real-time constraints

DELFT DATA SCIENCE

Regular

Canonical

DATA SCIENCE

[Kol et al. VMV 2014]

Example: Improved Navigation

- Number of wrong turns in navigation tasks:

Canonical View	Street-level View	Top-down View
2/12	3/12	6/12

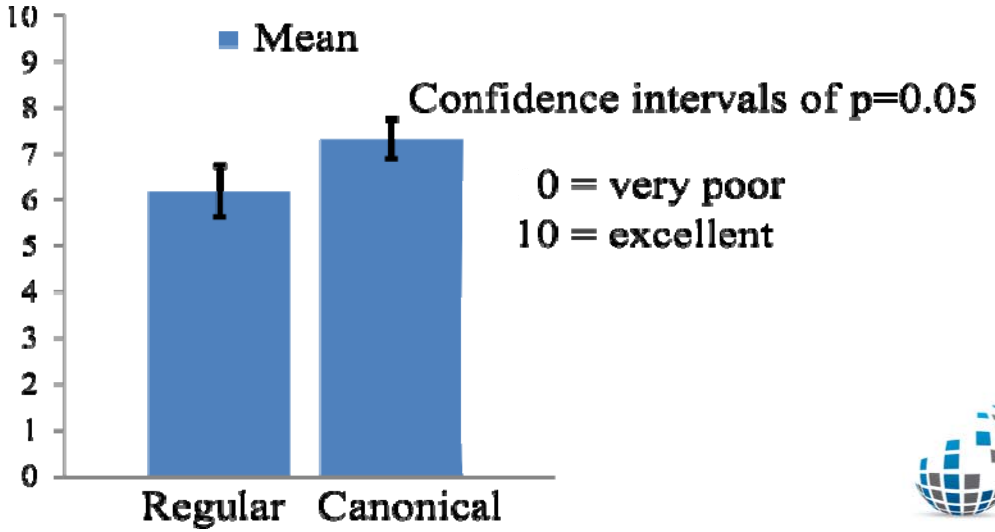

DELFT DATA SCIENCE

Novel Interaction Possibilities

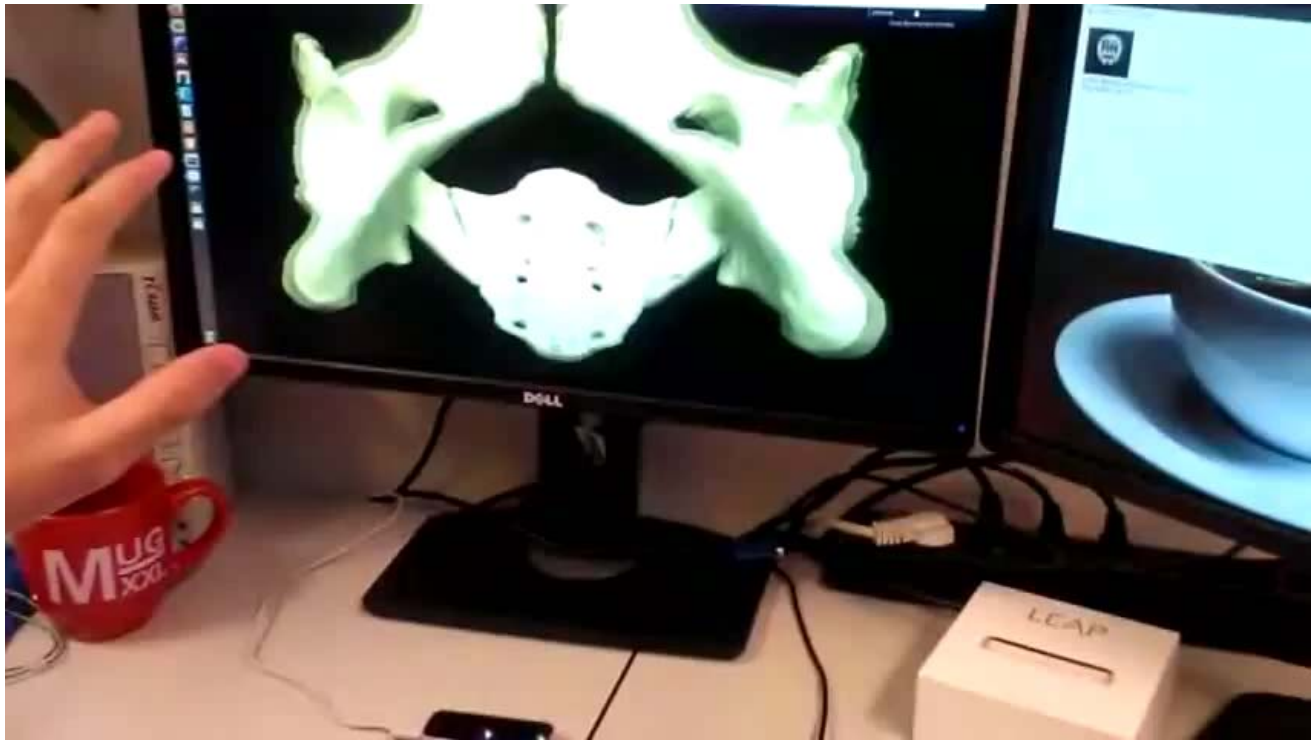

[Thiery, Guy, Boubekeur, Eisemann – TOG2016]

Encode Motion Efficiently - Animated Sphere Meshes

input mesh

sphere-mesh

interpolated sphere-mesh

DELFT DATA SCIENCE

[Thiery, Guy, Boubekeur, Eisemann – TOG2016]

Encode Motion Efficiently - Animated Sphere Meshes

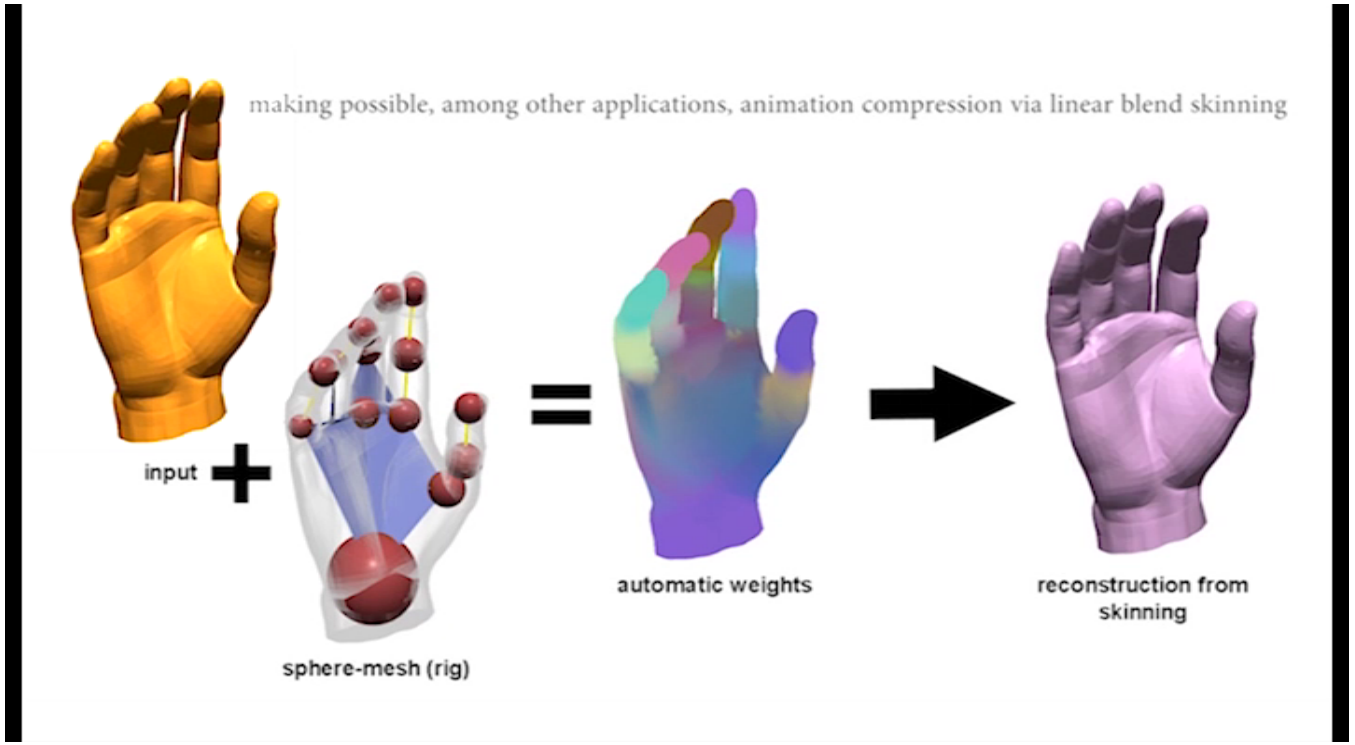

[Stengel, Grogorick, Eisemannx2, Magnor – ACM MM 2015]

Novel Display Devices

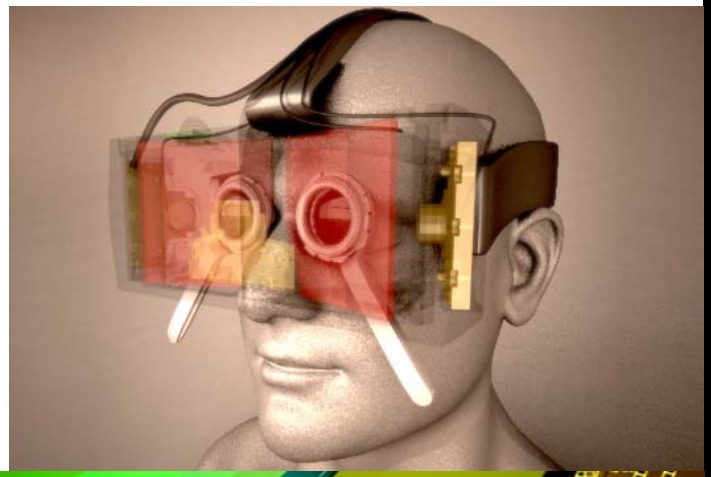

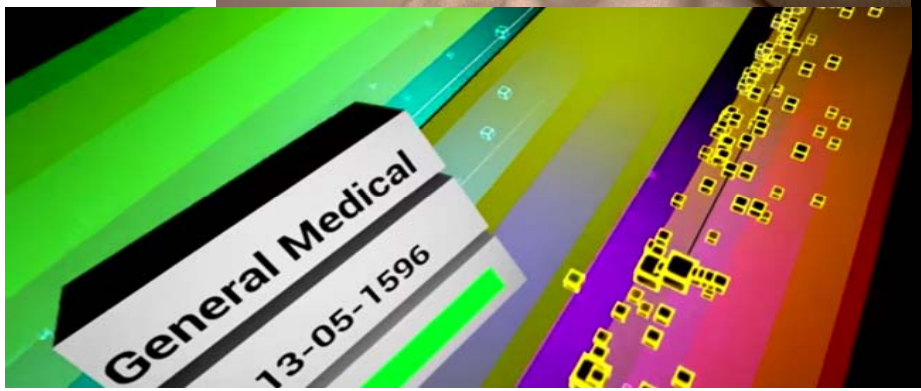

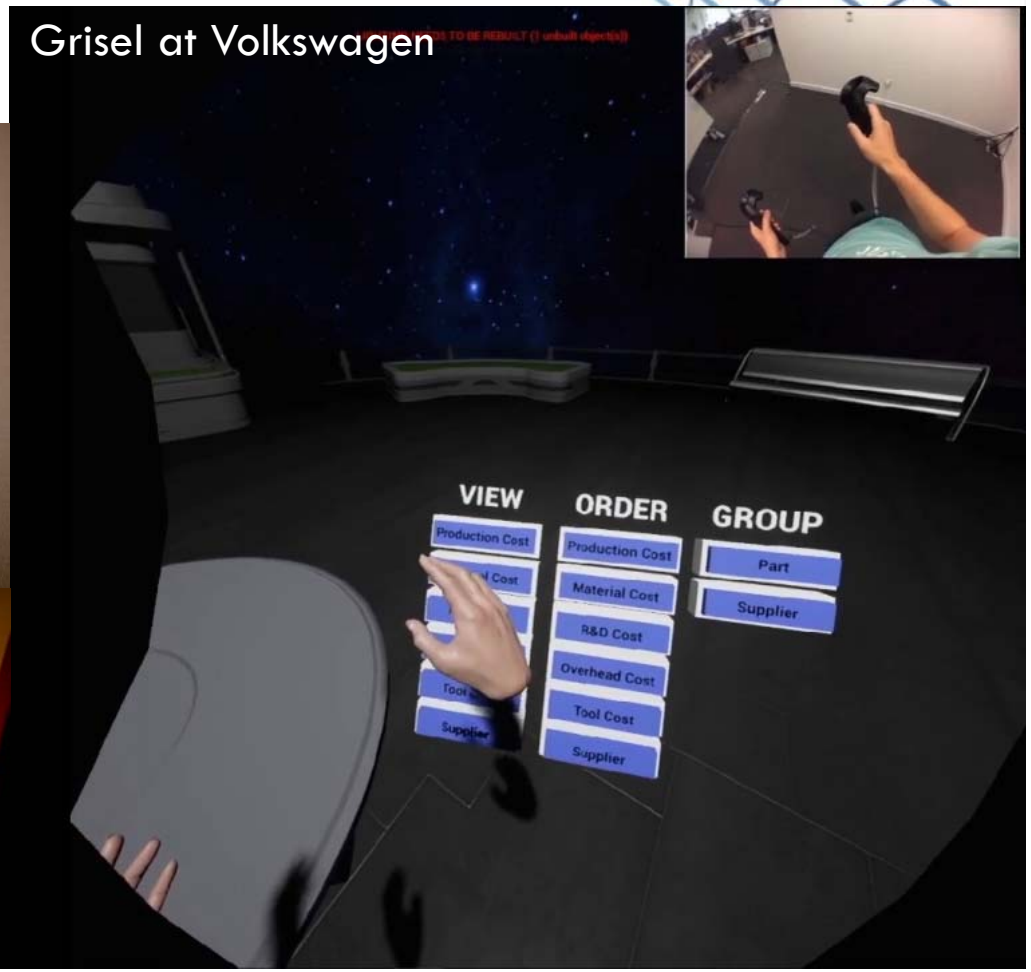

[Stengel, Grogorick, Eisemannx2, Magnor – ACM MM 2015]

Novel Display Devices

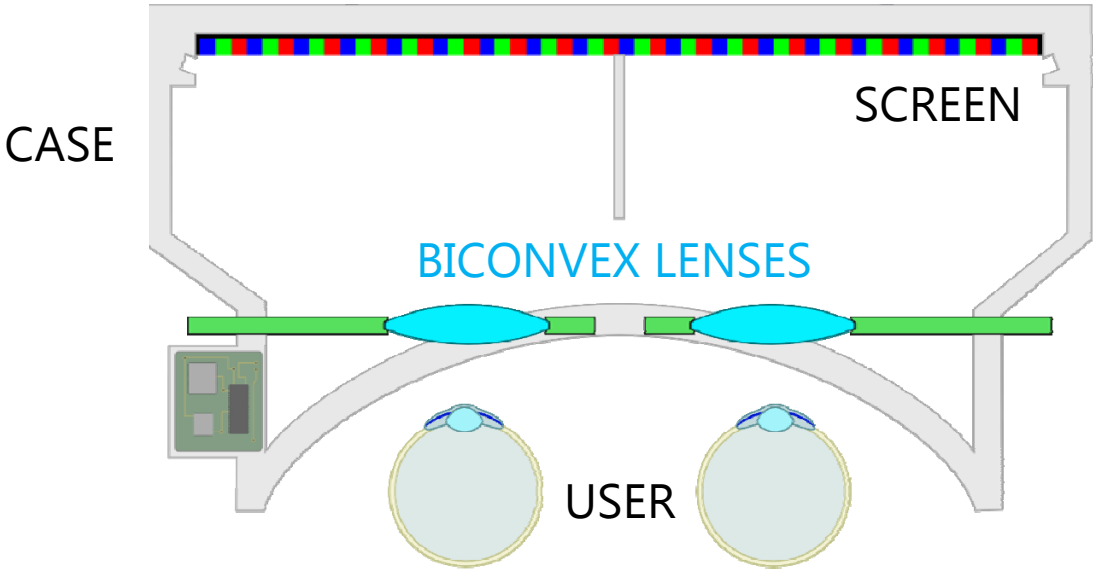

[Stengel, Grogorick, Eisemannx2, Magnor – ACM MM 2015]

Novel Display Devices

DICHROIC MIRROR

DICHROIC MIRROR

EYE TRACKING CAMERA

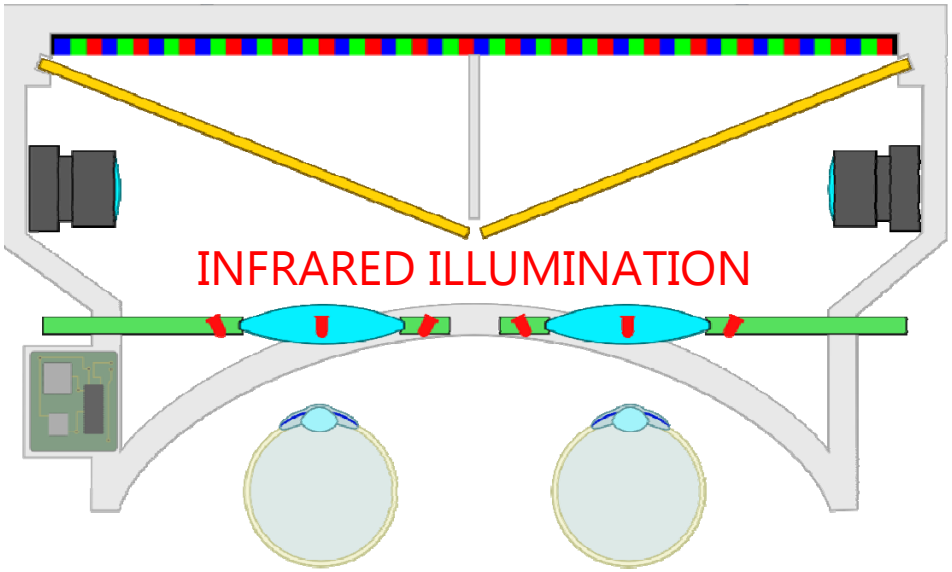

EYE TRACKING CAMERA

DELFT DATA SCIENCE

[Stengel, Grogorick, Eisemannx2, Magnor – ACM MM 2015]

CALIBRATION / USER

- Calibration via Inverse Rendering

A photon accurate model of the human eye [Deering'05]

The Physiology of the Eye [Adler et al., '08]

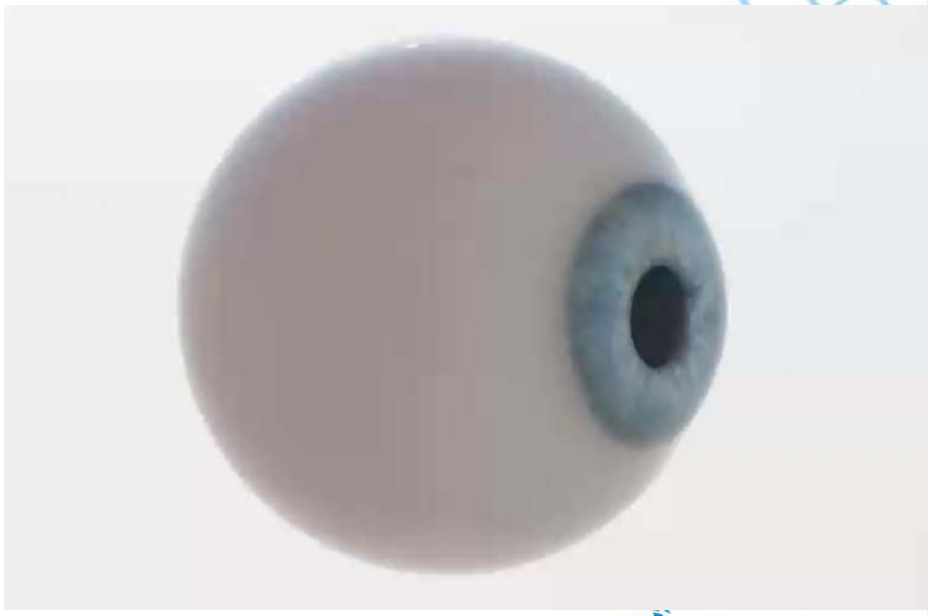

PHYSICAL EYE MODEL

[Stengel, Grogorick, Eisemannx2, Magnor – ACM MM 2015]

CALIBRATION / LENS LOCATION

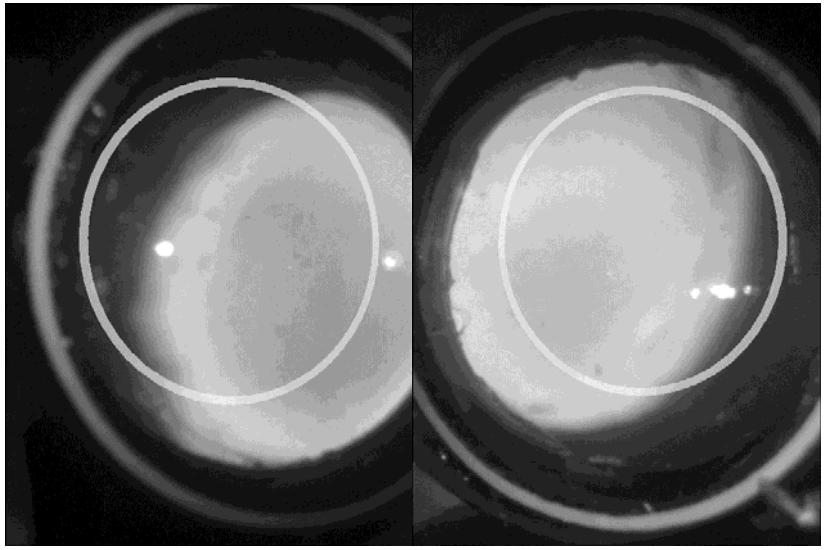

LEFT EYE RIGHT EYE

CAMERA CAPTURE

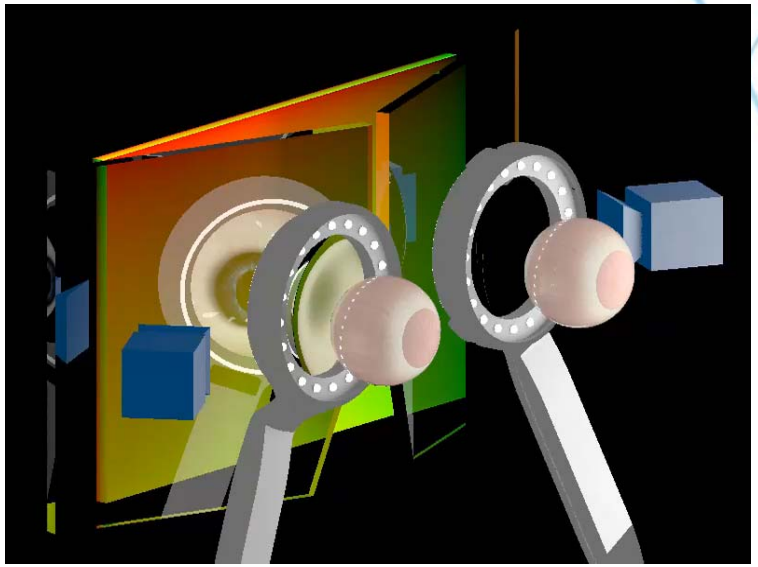

SIMULATED MODEL

DELFT DATA SCIENCE

[Stengel, Grogorick, Eisemannx2, Magnor – ACM MM 2015]

CALIBRATION / EYE BALL LOCATION

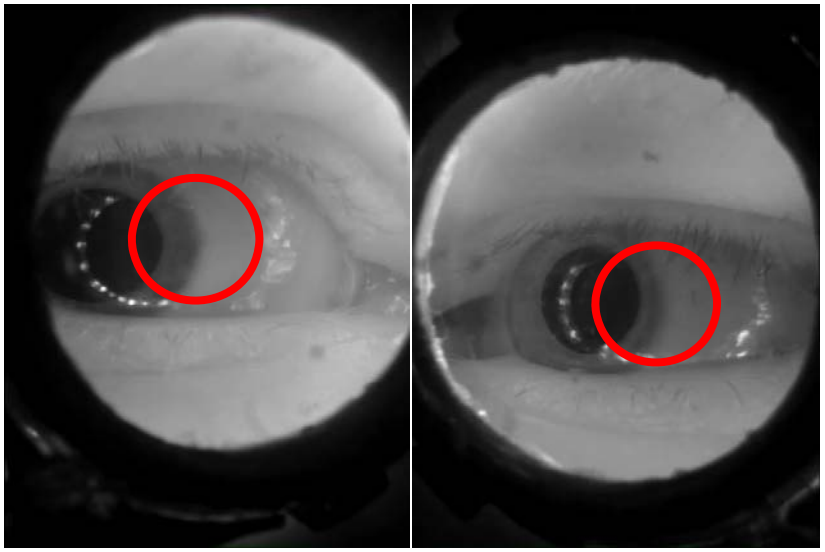

LEFT EYE RIGHT EYE

CAMERA CAPTURE

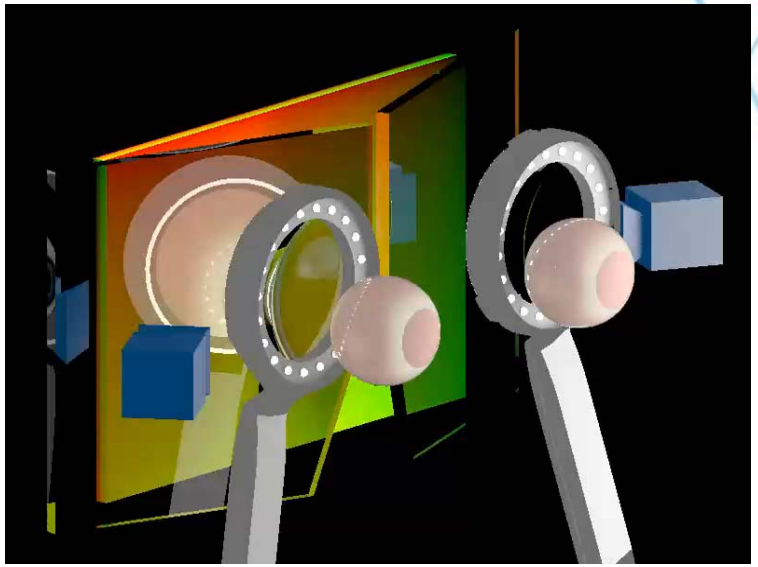

SIMULATED MODEL

[Stengel, Grogorick, Eisemannx2, Magnor – ACM MM 2015]

CALIBRATION / USER

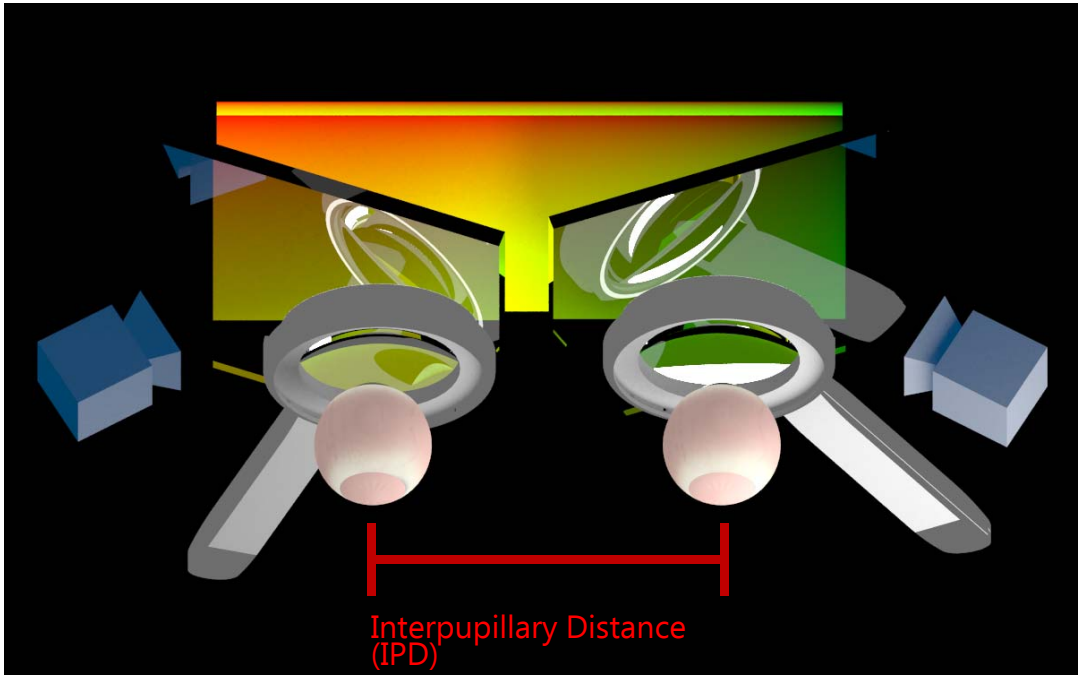

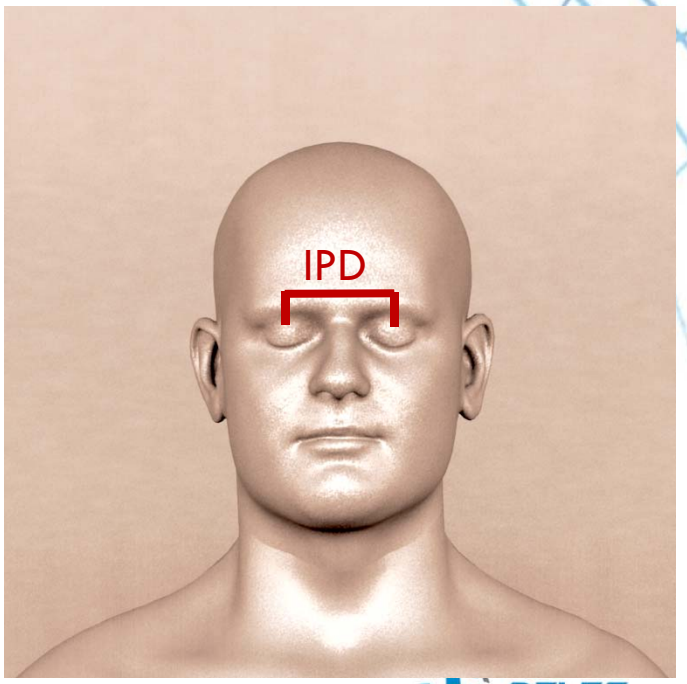

DELFT DATA SCIENCE

[Stengel, Grogorick, Eisemannx2, Magnor – ACM MM 2015]

PUPIL TRACKING

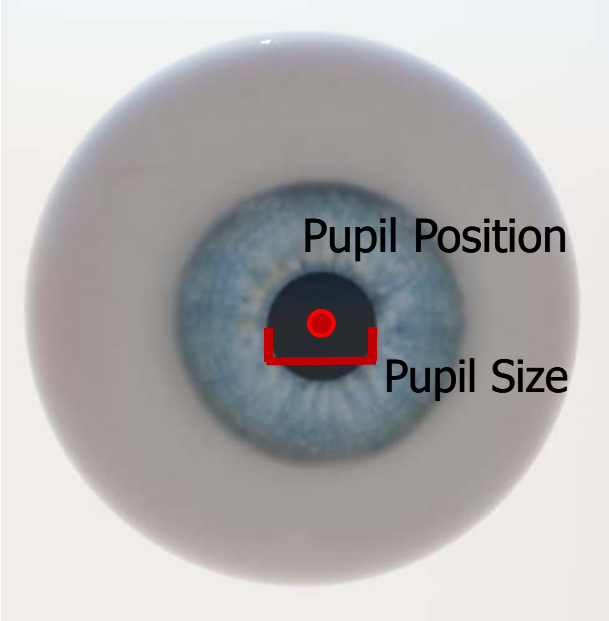

PHYSICAL EYE MODEL

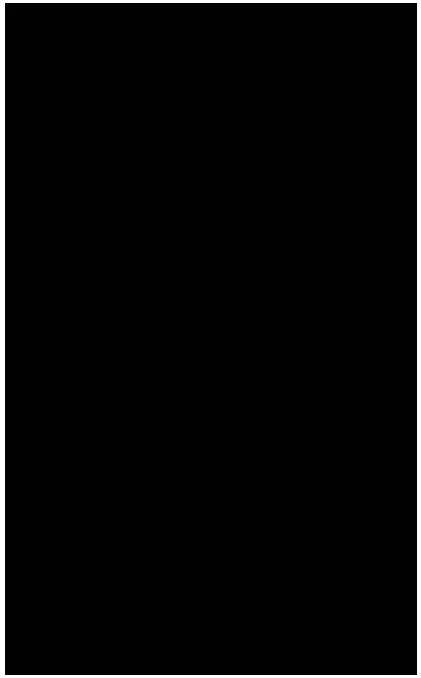

SCREEN CAPTURE

DELFT DATA SCIENCE

[Stengel, Grogorick, Eisemannx2, Magnor – ACM MM 2015]

TRACKING PERFORMANCE / DESKTOP

- two cores on i7@3.5Ghz

Processing step	Timings
Camera frame capture	11 ms
~~Occlusion metric~~	~~2 ms~~
~~Pupil Detection (occluded)~~	~~2 ms~~
Pupil Detection (visible)	0.9 ms
Gaze computation	< 0.1ms
	~~ 16.0 ms (total)~~
	~ 12.0 ms (total)

may be skipped for many users

~5 ms

~ 1 ms

Accommodation Simulation

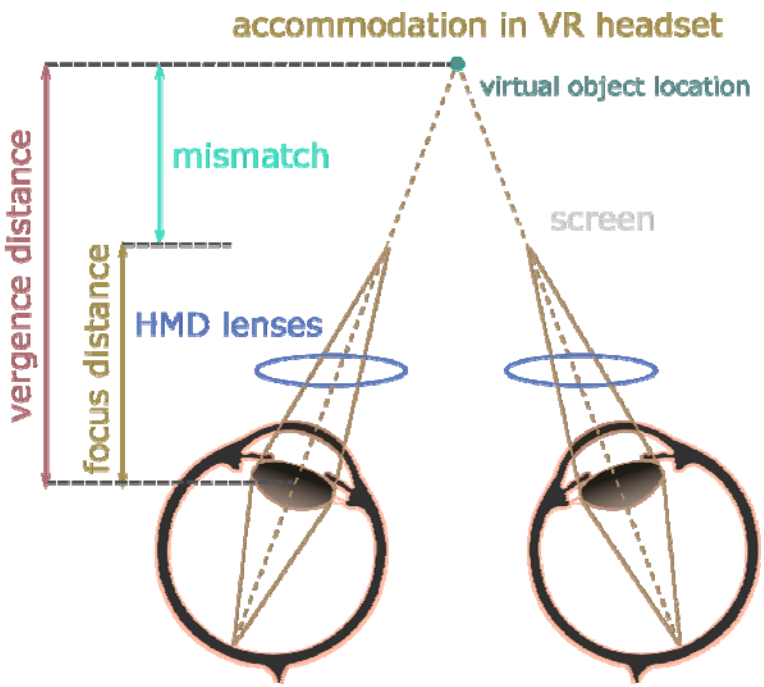

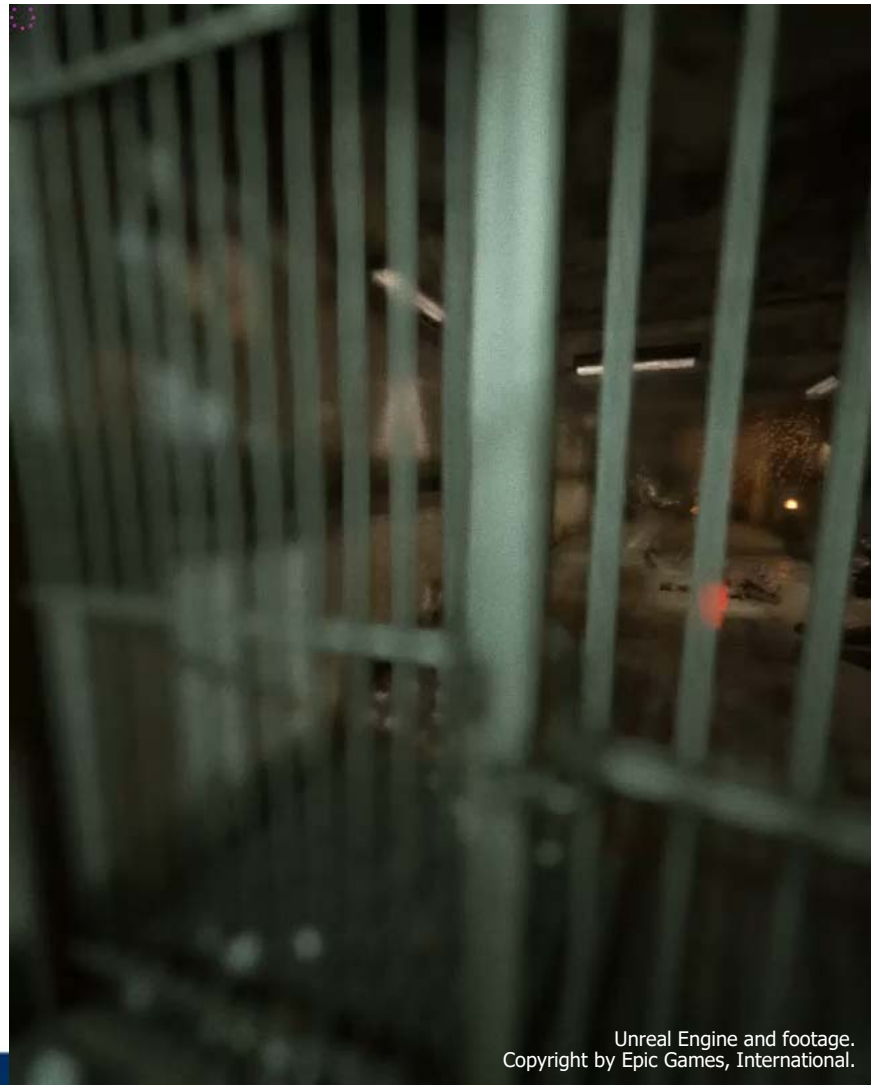

Visualization and Perception

- Realistic Rendering
- Perceptual Methods
- Visualization & Interface

Effective Data Visualization Requires

- Large-Scale Rendering

- Visualization and Percepti

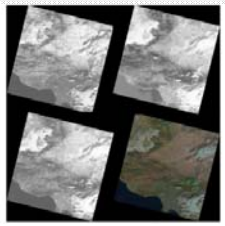

- Data Analysis

DELFT DATA SCIENCE

Questions?

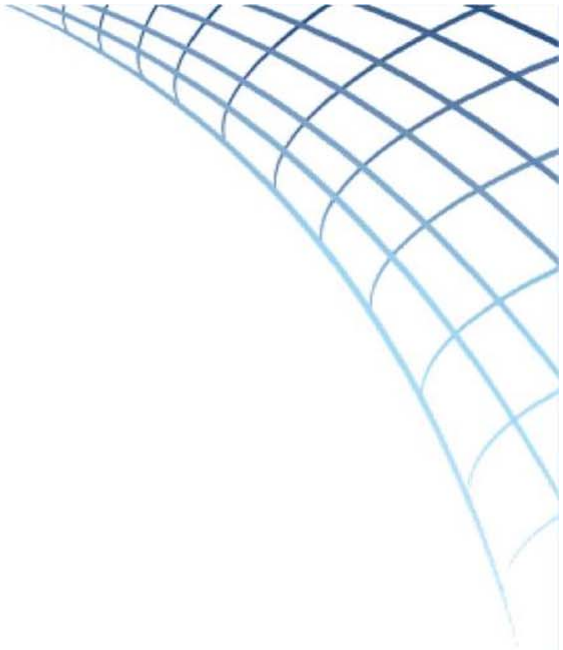

Effective Data Visualization Requires

- Large-Scale Rendering

- Visualization and Percepti

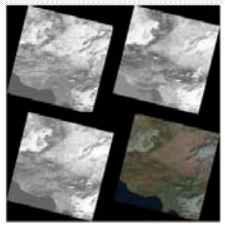

- Data Analysis

DELFT DATA SCIENCE

Data Analysis

- High-dimensional/Heterogeneous Data
- Dimensionality Reduction
- Visual Analytics

Data Analysis

DELFT DATA SCIENCE

Koninklijk Instituut Van Ingenieurs

[Sunkel, Jansen, Wand, Eisemann, Seidel EG 2012]

Supported Recognition

- here 128 M points, find all windows: 2 min

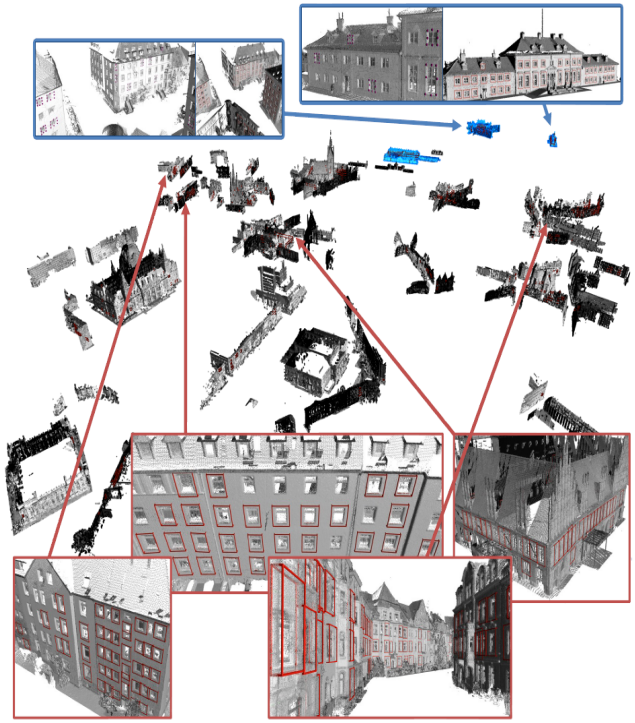

DELFT DATA SCIENCE

Koninklijk Instituut Van Ingenieurs

[Silva, Eisemann, Bidarra, Coelho - PCG 2015]

Model Construction: Blueprints + Rules +…

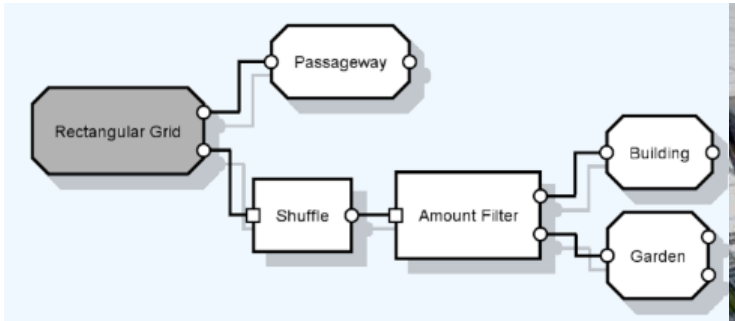

DELFT DATA SCIENCE

[Baboud, Cadik, Eisemann, Seidel – CVPR'11]

Model Construction

- Heterogeneous sources

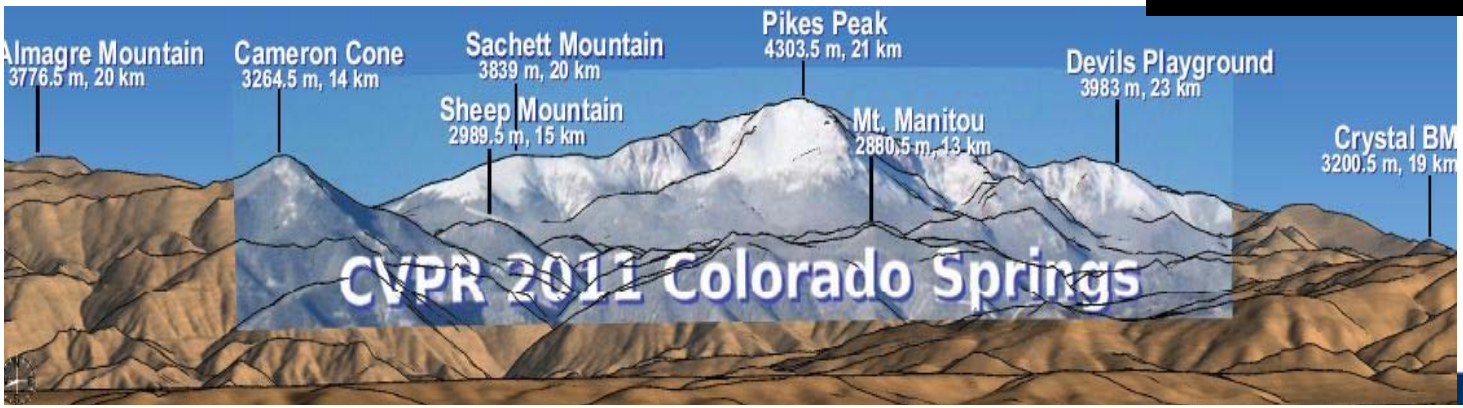

DELFT DATA SCIENCE

Koninklijk Instituut Van Ingenieurs

[Donchyts, Schellekens, Winsemius, Eisemann, van de Giesen - Remote Sens. 16]

Model Construction

- Integrating Data Sets for River Validation

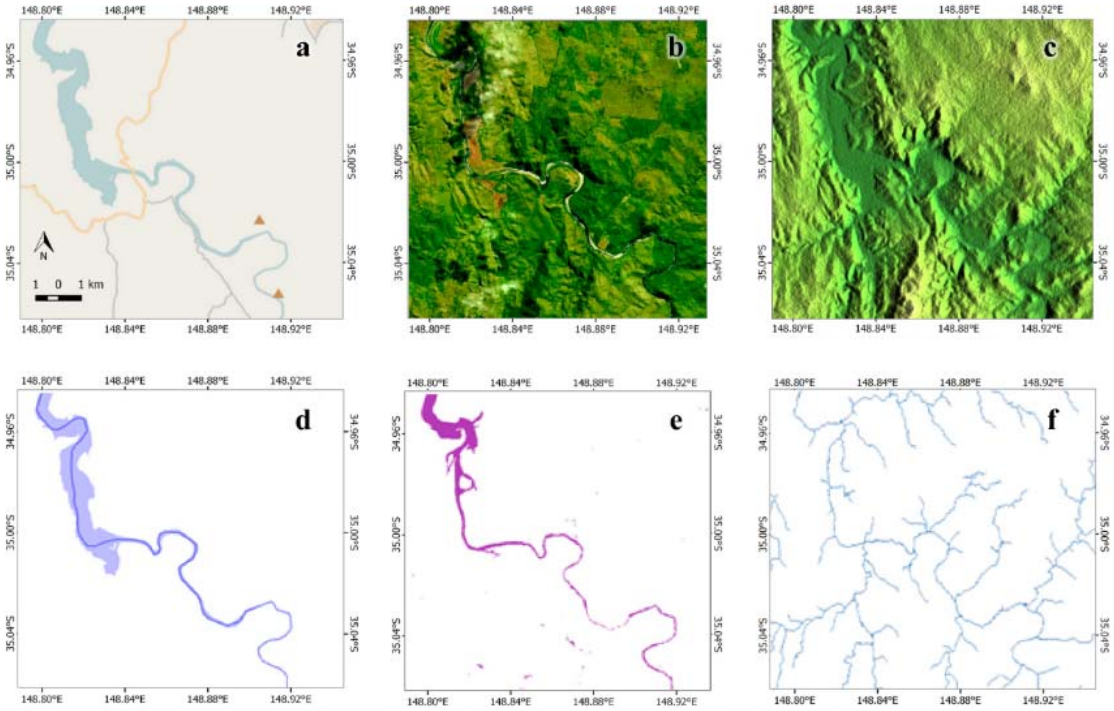

DELFT DATA SCIENCE

Data growth = increase in information?

- What if you do not yet know what "useful" is?

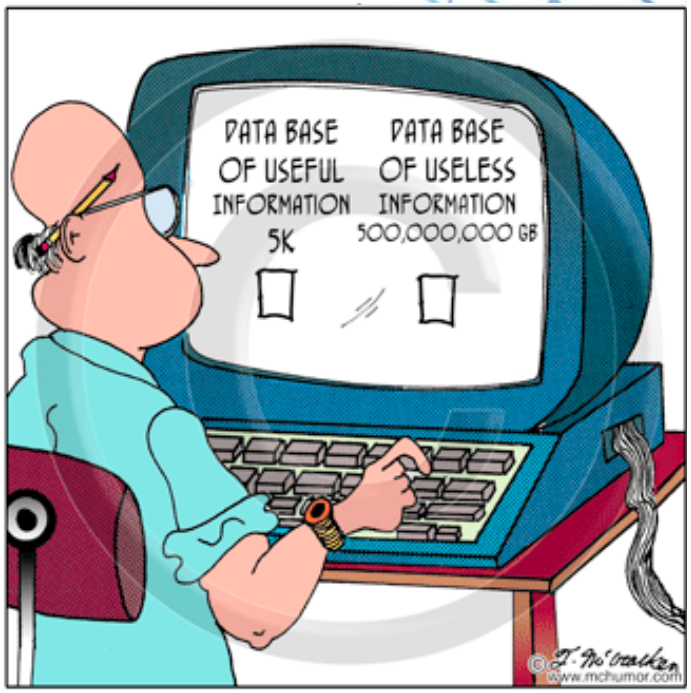

DELFT DATA SCIENCE

General Approach to Visual Analytics?

Visual Analytics:

Interactive hypothesis finding via Data

- Many challenges, but we will focus on:
- Dimensionality

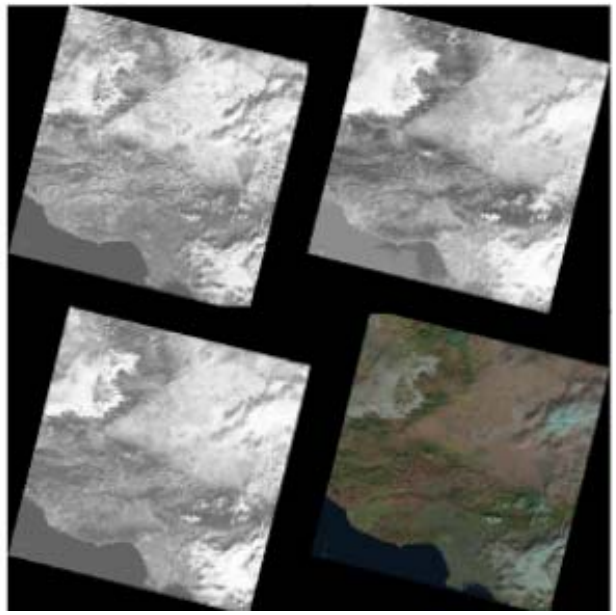

DELFT DATA SCIENCE

High-Dimensional Data

- Standard Solutions:

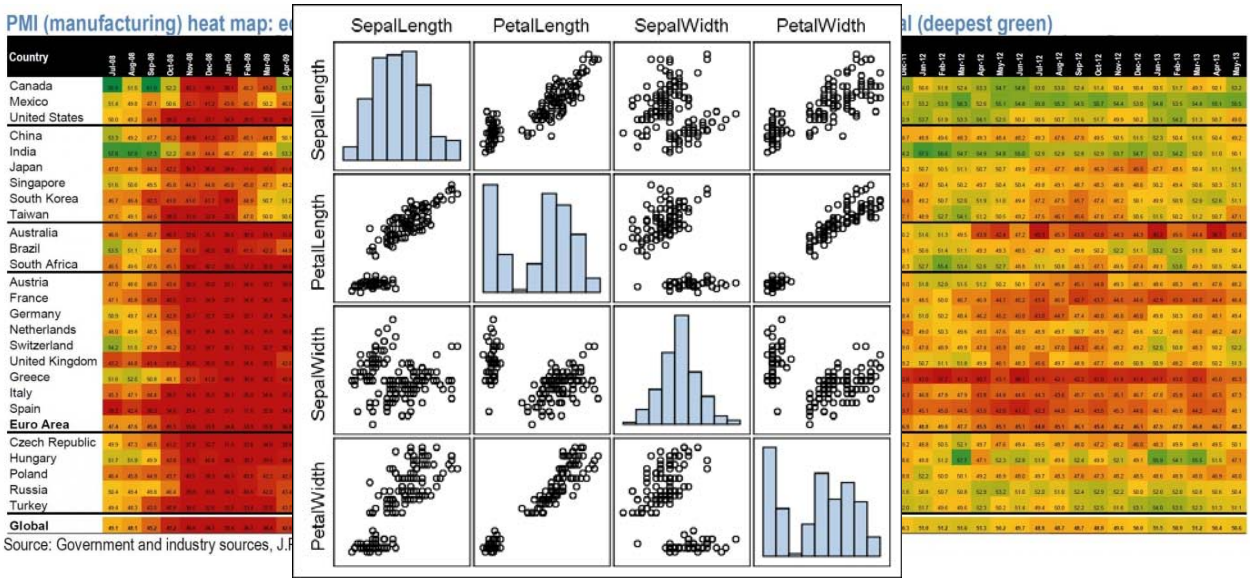

Parallel Coordinate Plots

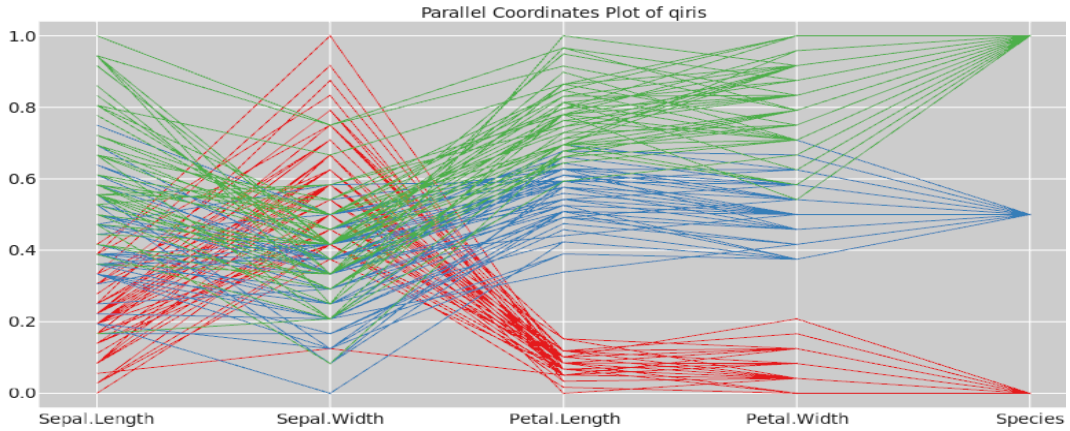

- More dimensions in one view
- Easier detection of patterns/outliers in data

DELFT DATA SCIENCE

Parallel Coordinates

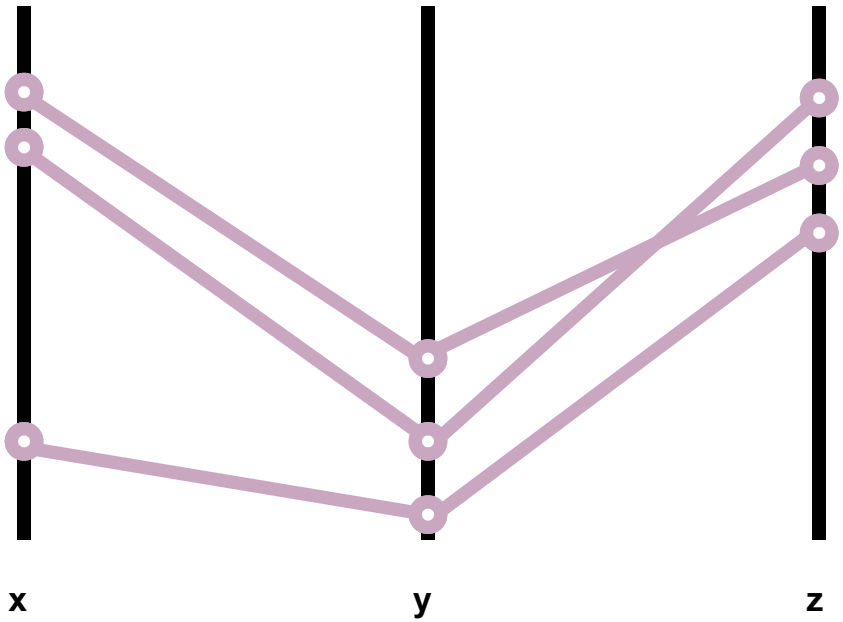

Parallel Coordinates

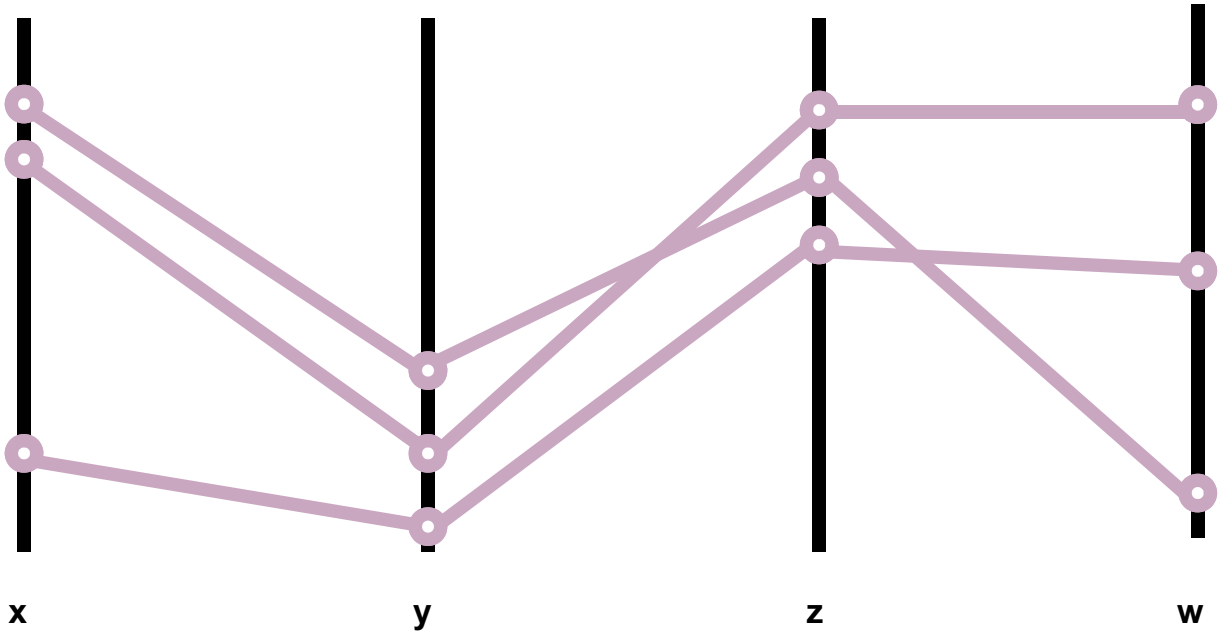

Parallel Coordinates

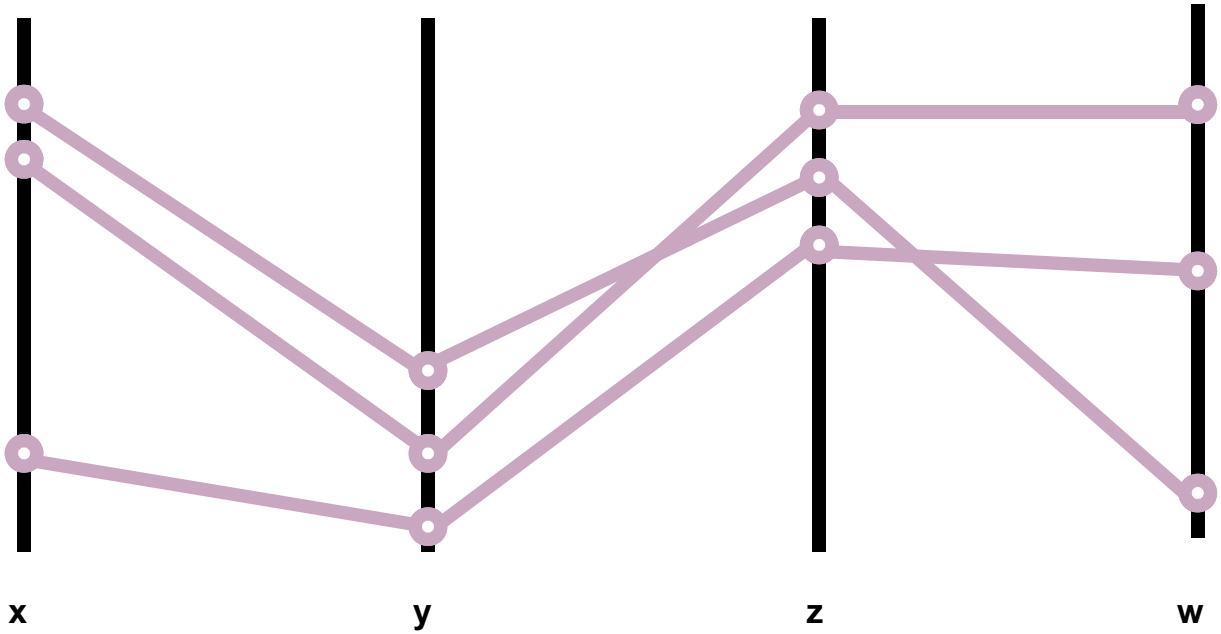

The Problem

- Clutter!

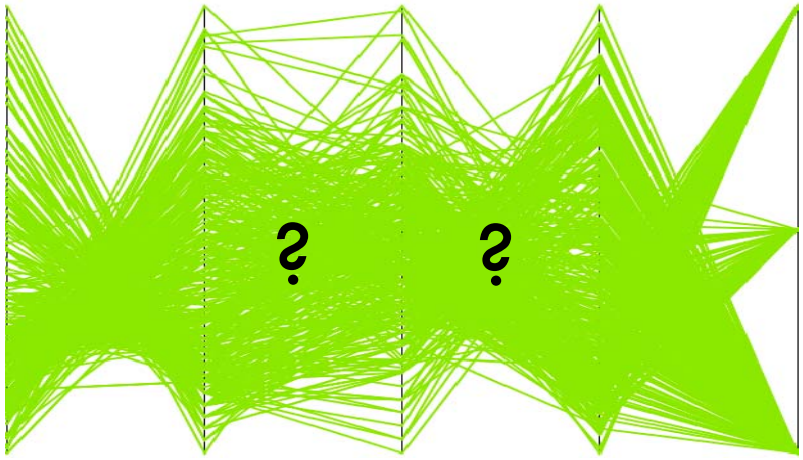

Parallel Coordinates

Drag a selection box on axis to **brush** data

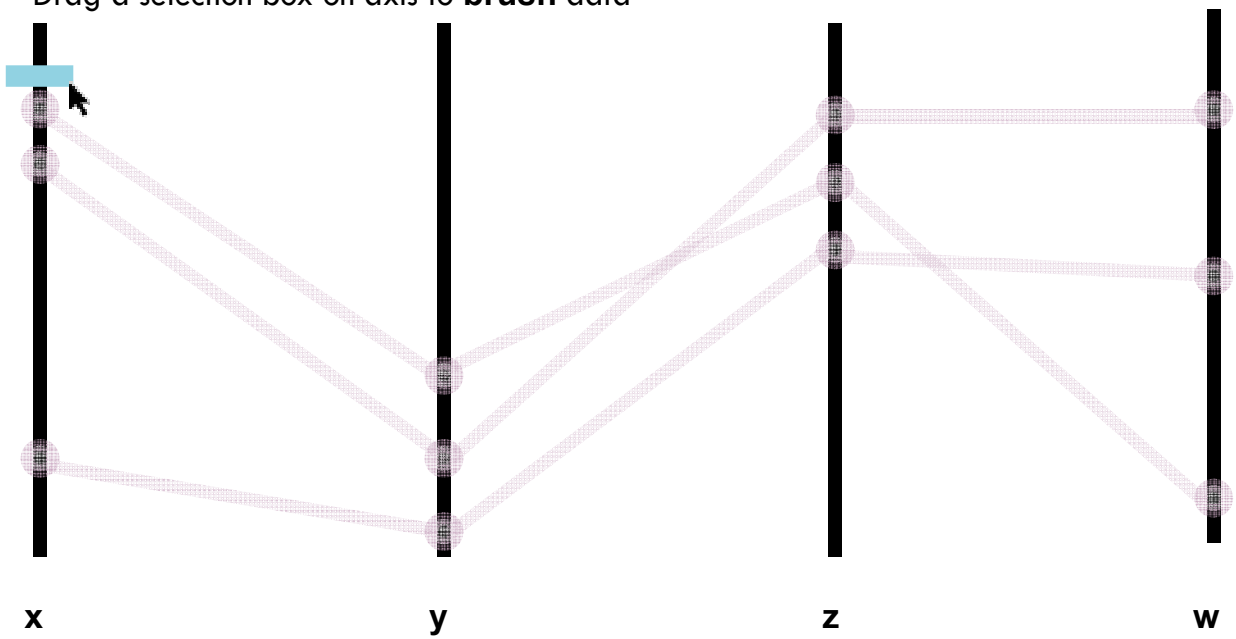

Parallel Coordinates

Drag a selection box on axis to **brush** data

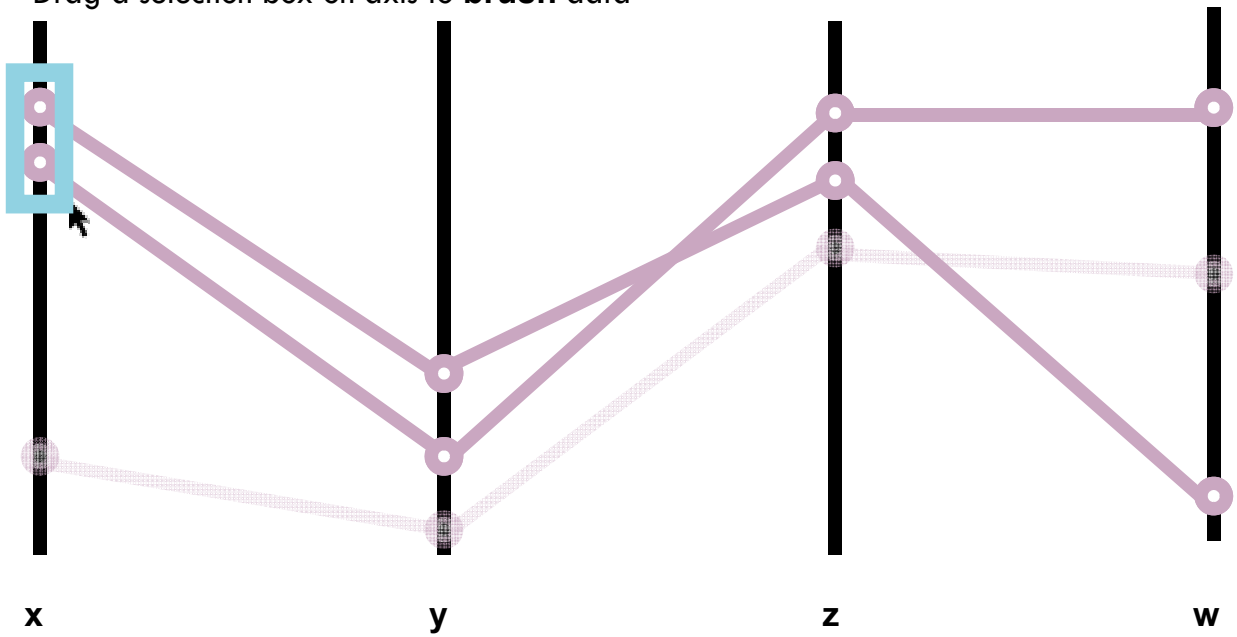

DELFT DATA SCIENCE

[Raidou, Eisemann, Breuwer, Eisemann, Vilanova - Vis 2015]

5-D Example

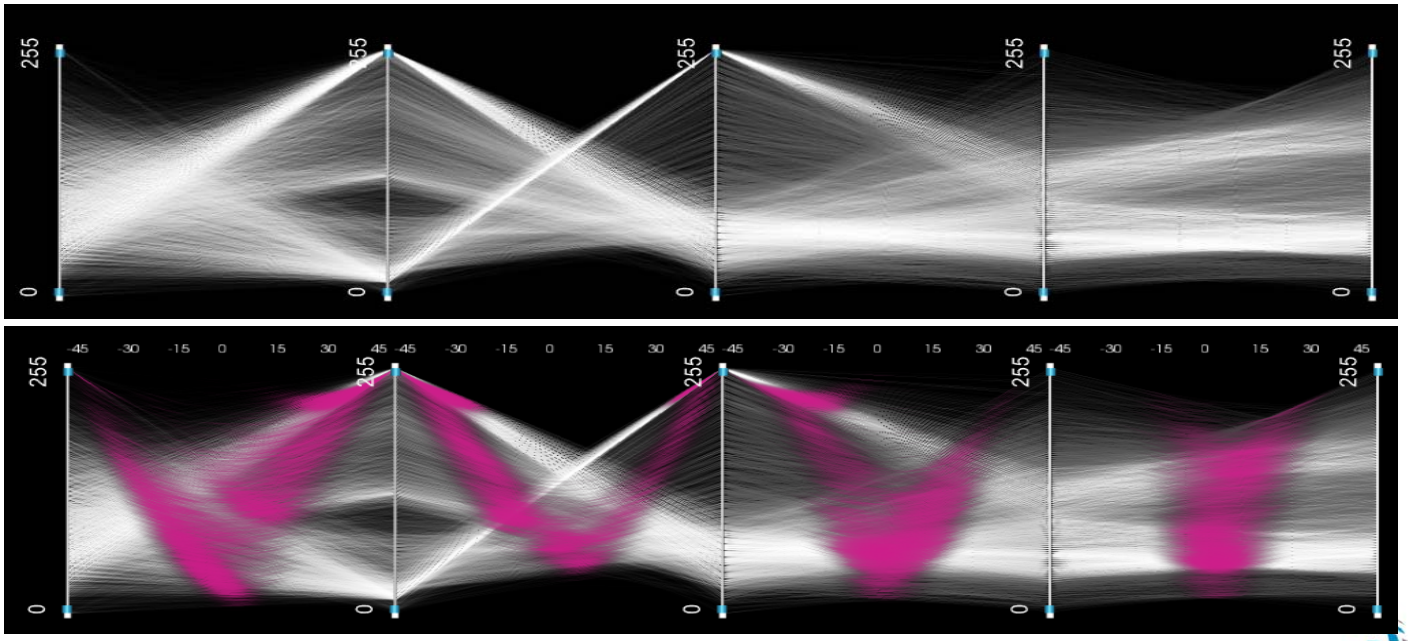

[Raidou, Eisemann, Breuwer, Eisemann, Vilanova - Vis 2015]

Orientation-Enhanced Parallel Coordinate Plots

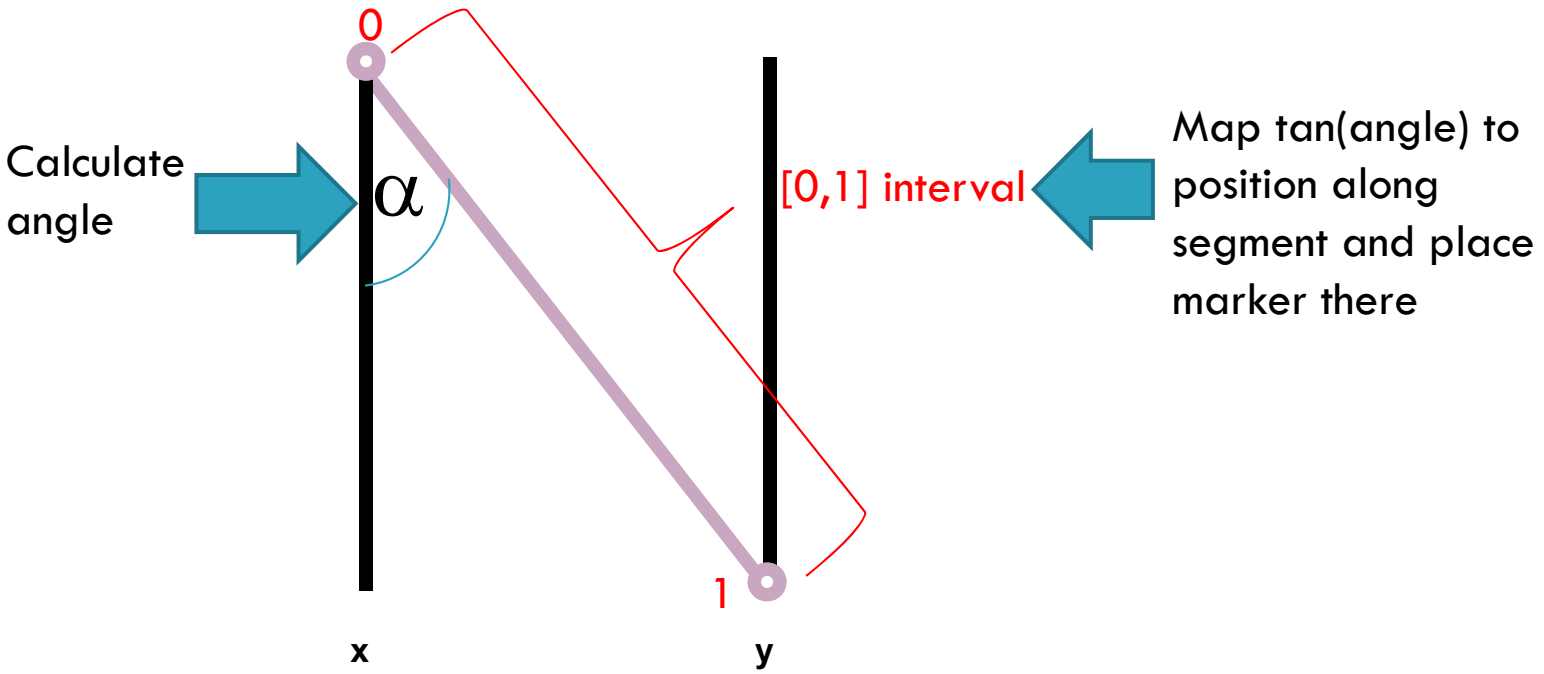

[Raidou, Eisemann, Breuwer, Eisemann, Vilanova - Vis 2015]

Orientation-Enhanced Parallel Coordinate Plots

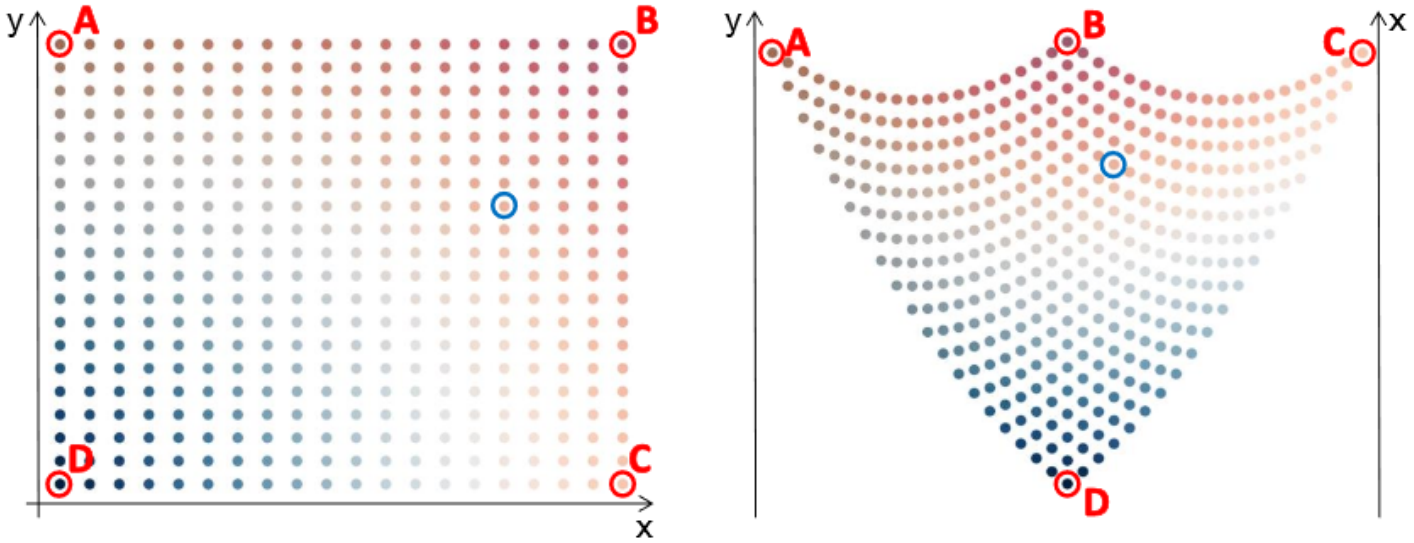

DELFT DATA SCIENCE

[Raidou, Eisemann, Breuwer, Eisemann, Vilanova - Vis 2015]

5-D Example

- Very effective to spot trends
- Supports interaction similar to the original PCP
- …but fixed dimensionality

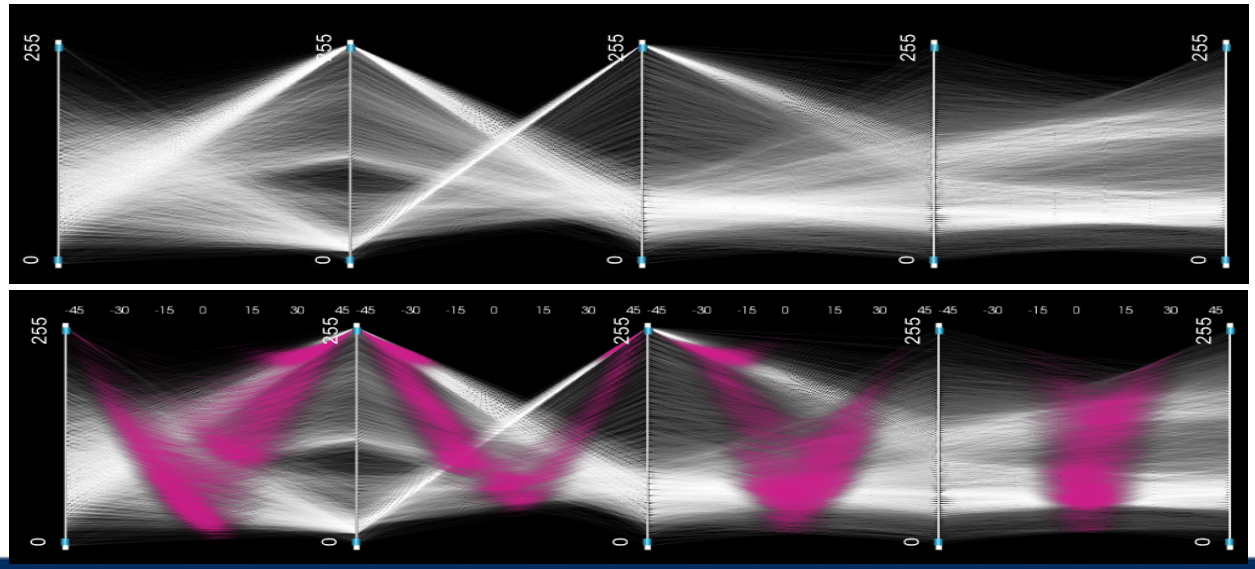

DELFT DATA SCIENCE

Koninklijk Instituut Van Ingenieurs

High-Dimensional Spaces

- A high dimensional space does not imply high dimensionality

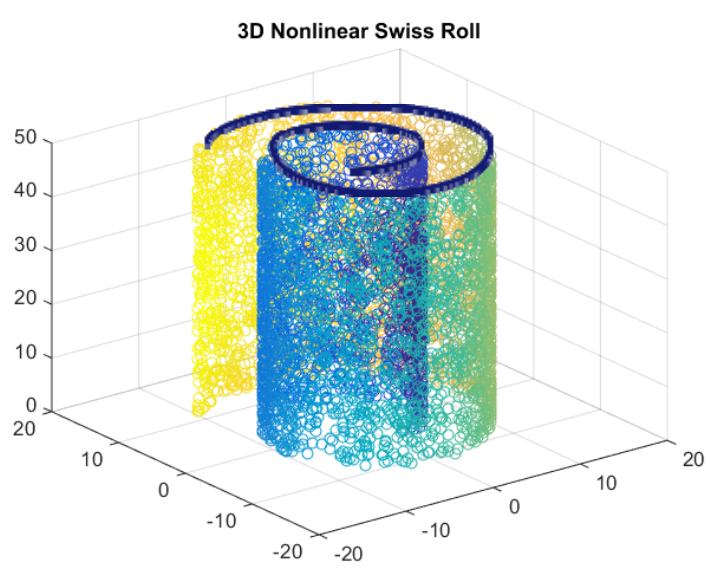

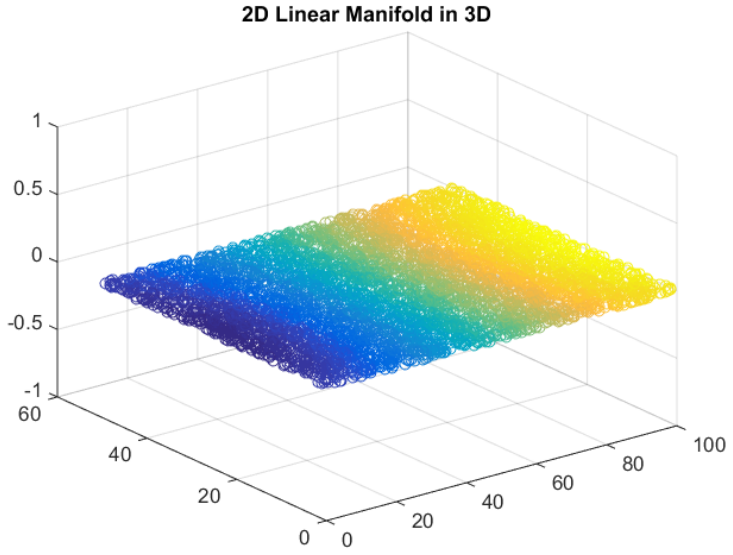

ELFT ATA SCIENCE

Multi-dimensional Embedding

- Reduce dimensions but maintain neighbors/distances

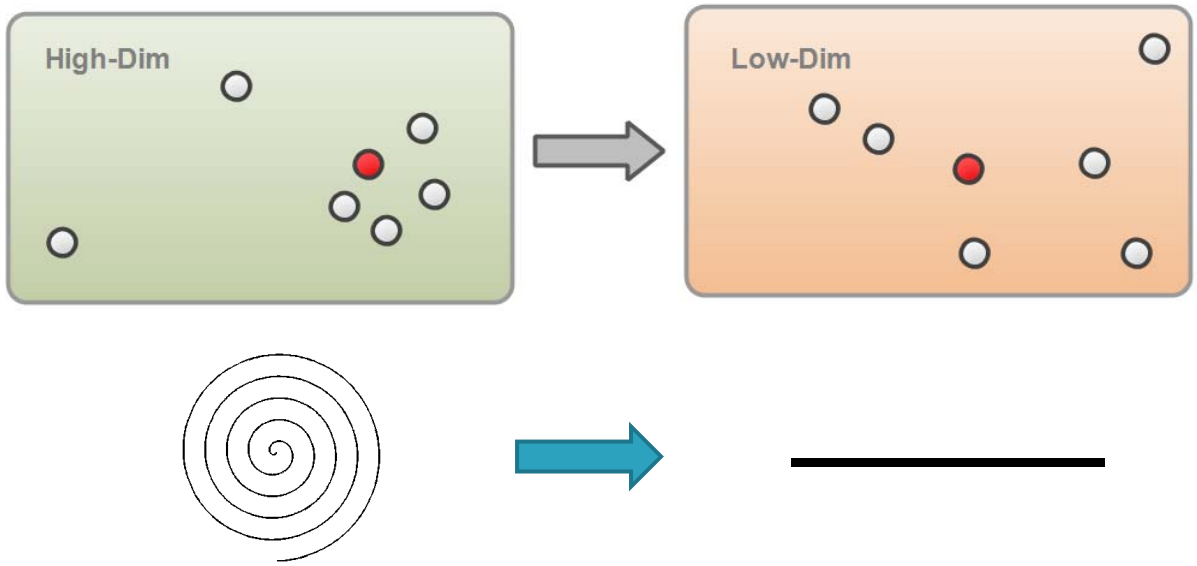

DELFT DATA SCIENCE

Example: Dimensionality Reduction

100K samples
with 4096 dim.

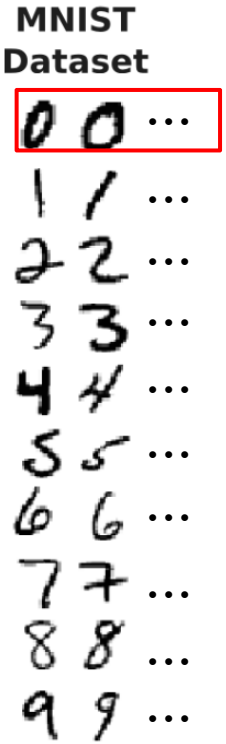

100K samples
embedded in 2D

T-SNE: Van der Maaten and Hinton – Journal of Machine Learning 2008

DELFT DATA SCIENCE

Challenges of Embedding

- Performance
  - Can be up to hours of calculation
- No Hierarchical information
  - Global embedding

Challenges of Embedding

- Performance
  - Can be up to hours of calculation

- No Hierarchical information
  - Global embedding

[Pezzotti, Leliveldt, van der Maaten, Hoelt, Eisemann, Vilanova – TVCG2016]

Extreme Acceleration

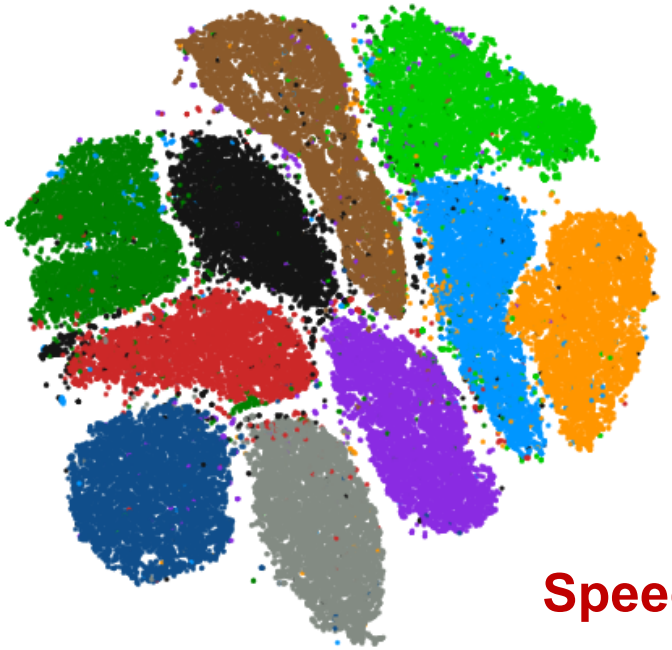

Speed up: 100x

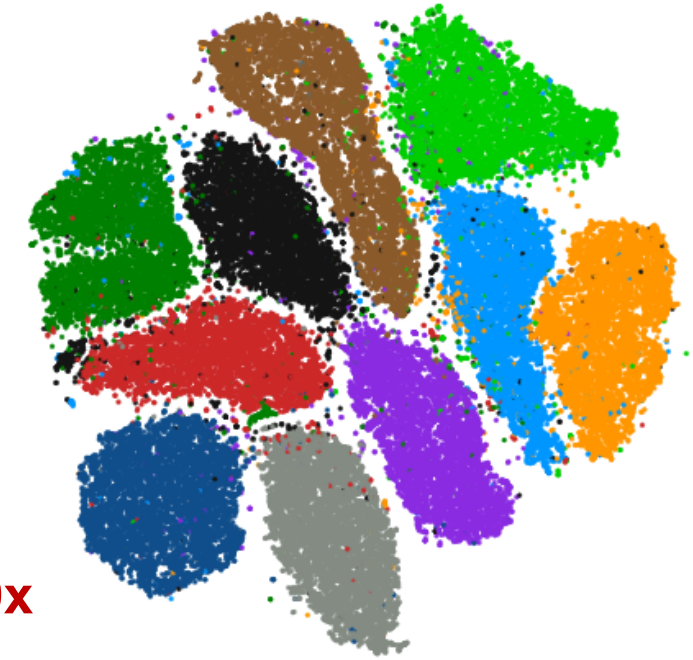

tSNE
Time: 3191.8 s

A-tSNE
Time: 30.1 s

DELFT DATA SCIENCE

Approximated similarities computation

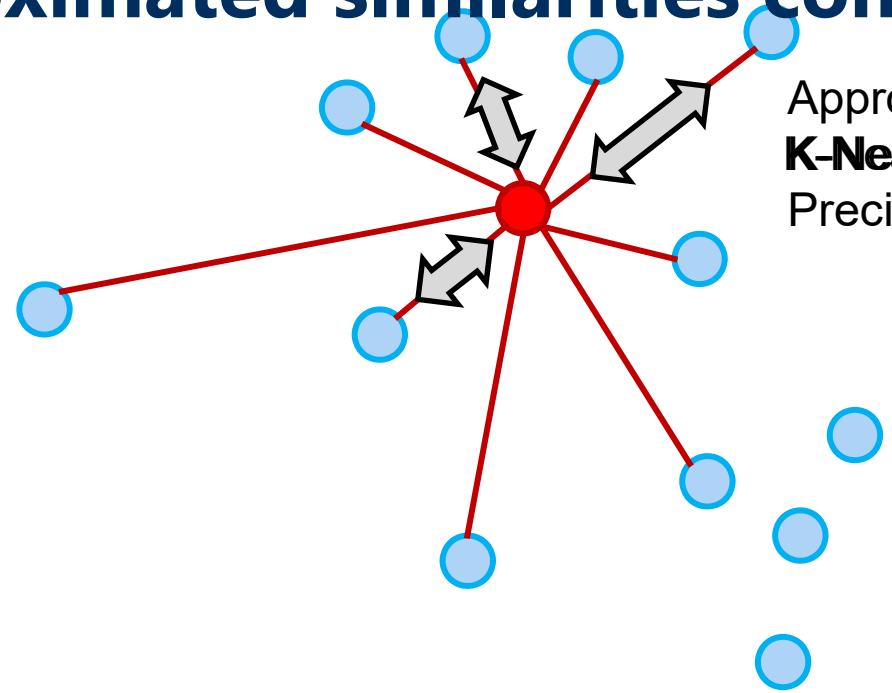

Approximated
K-Nearest-Neighborhood [1]
Precision: 50%

Fast Approximate Nearest Neighbors with Automatic Algorithm Configuration - Muja et al. - 2009

DELFT DATA SCIENCE

Approximated tSNE

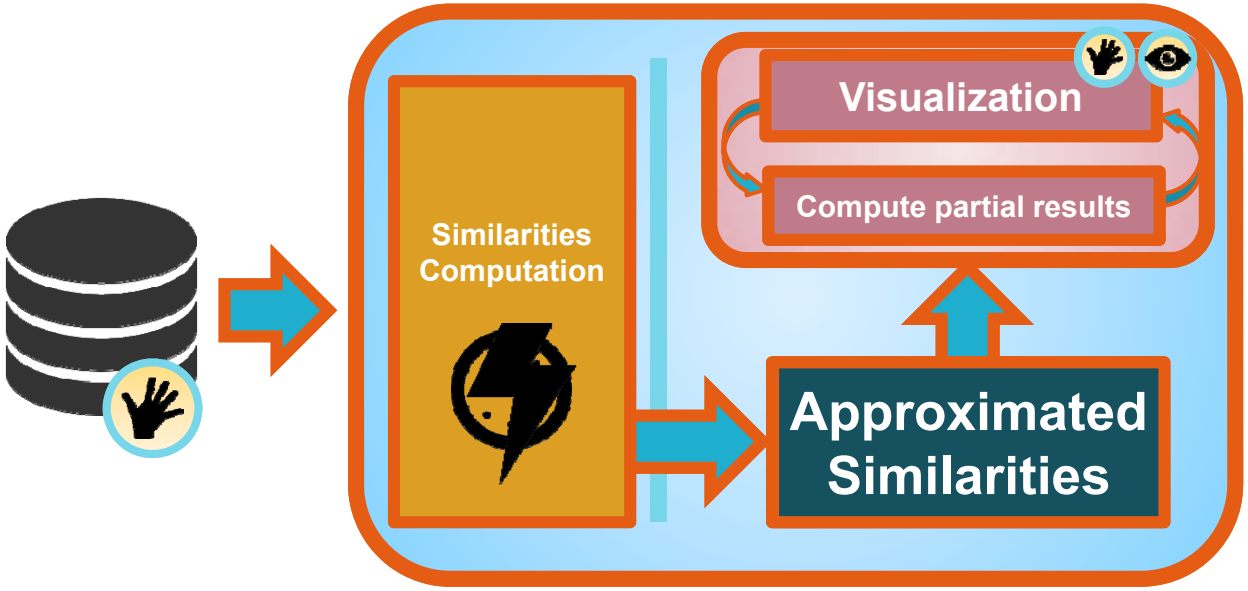

Approximated tSNE

[Pezzotti, Leliveldt, van der Maaten, Hoelt, Eisemann, Vilanova – TVCG2016]

Steerability & Approximation

- Density-based visualization
  - Interaction support
  - Fast feedback and convergence

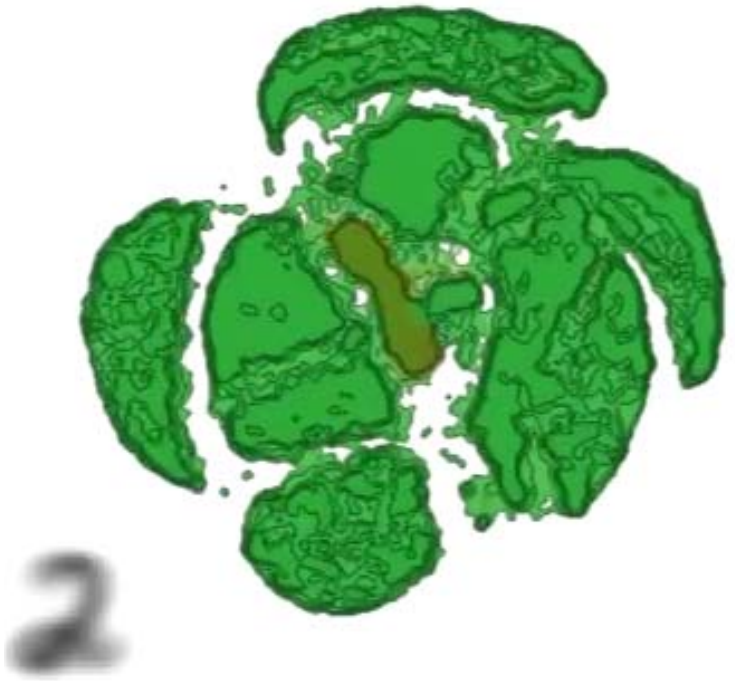

Preprocessing: 12 s (on purpose very low precision)

DELFT DATA SCIENCE

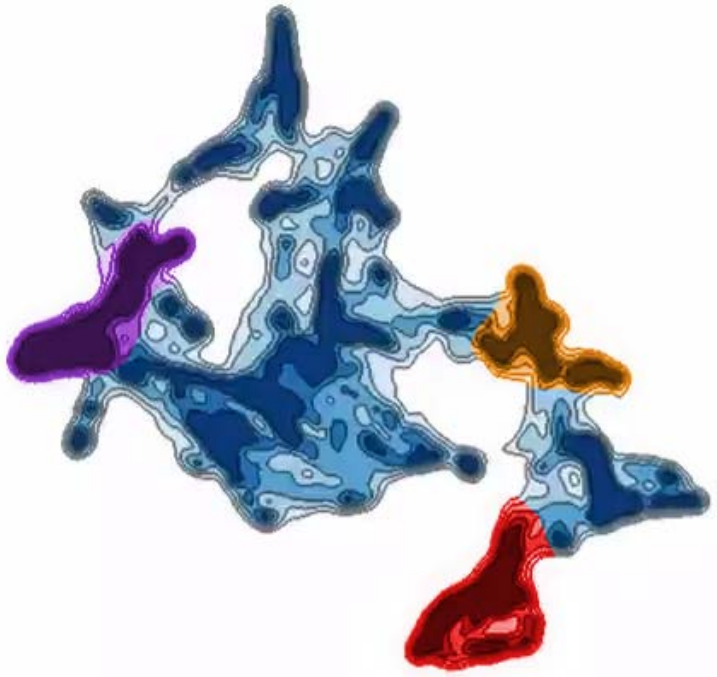

Sagittal

Coronal

Axial

3D Volume

61164 data points (Voxels)

4345 dimensions (Gene expression)

DELFT DATA SCIENCE

Case Study I : Gene expression

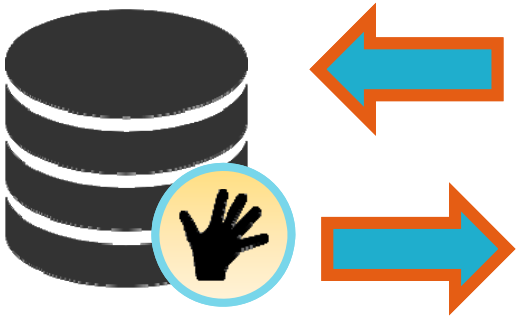

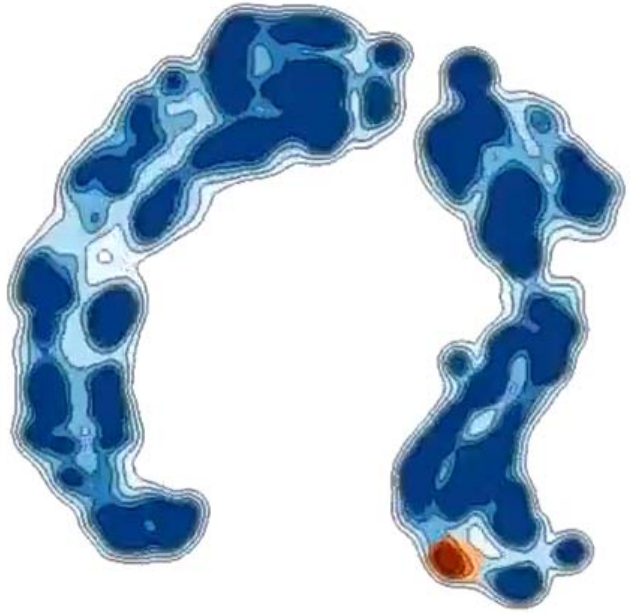

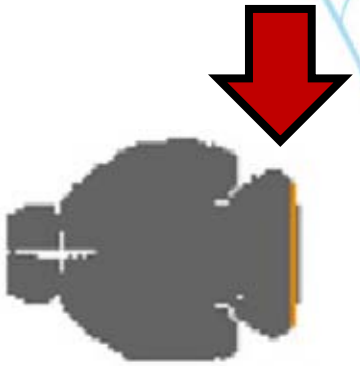

Speed up: 250x

A-tSNE 50 seconds – tSNE 3 hours and 50 minutes

DELFT DATA SCIENCE

Case Study II : High-dimensional data streams

Image courtesy of www.activ8all.com

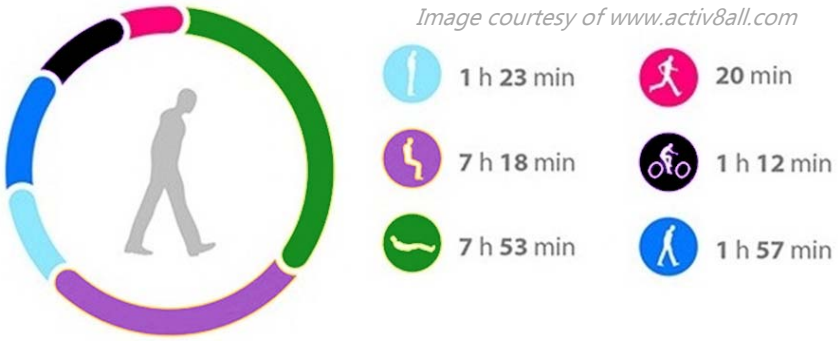

Chest - Ankle - Wrist

52 Dimensions every 100 ms

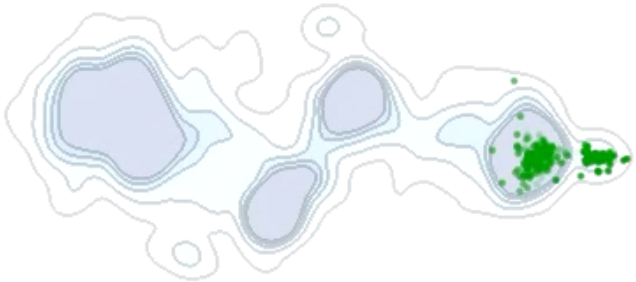

DELFT DATA SCIENCE

Challenges of Embedding

- Performance
  - Can be up to hours of calculation

- No Hierarchical information
  - Global embedding

How to structure data hierarchically without knowing about the hierarchy?

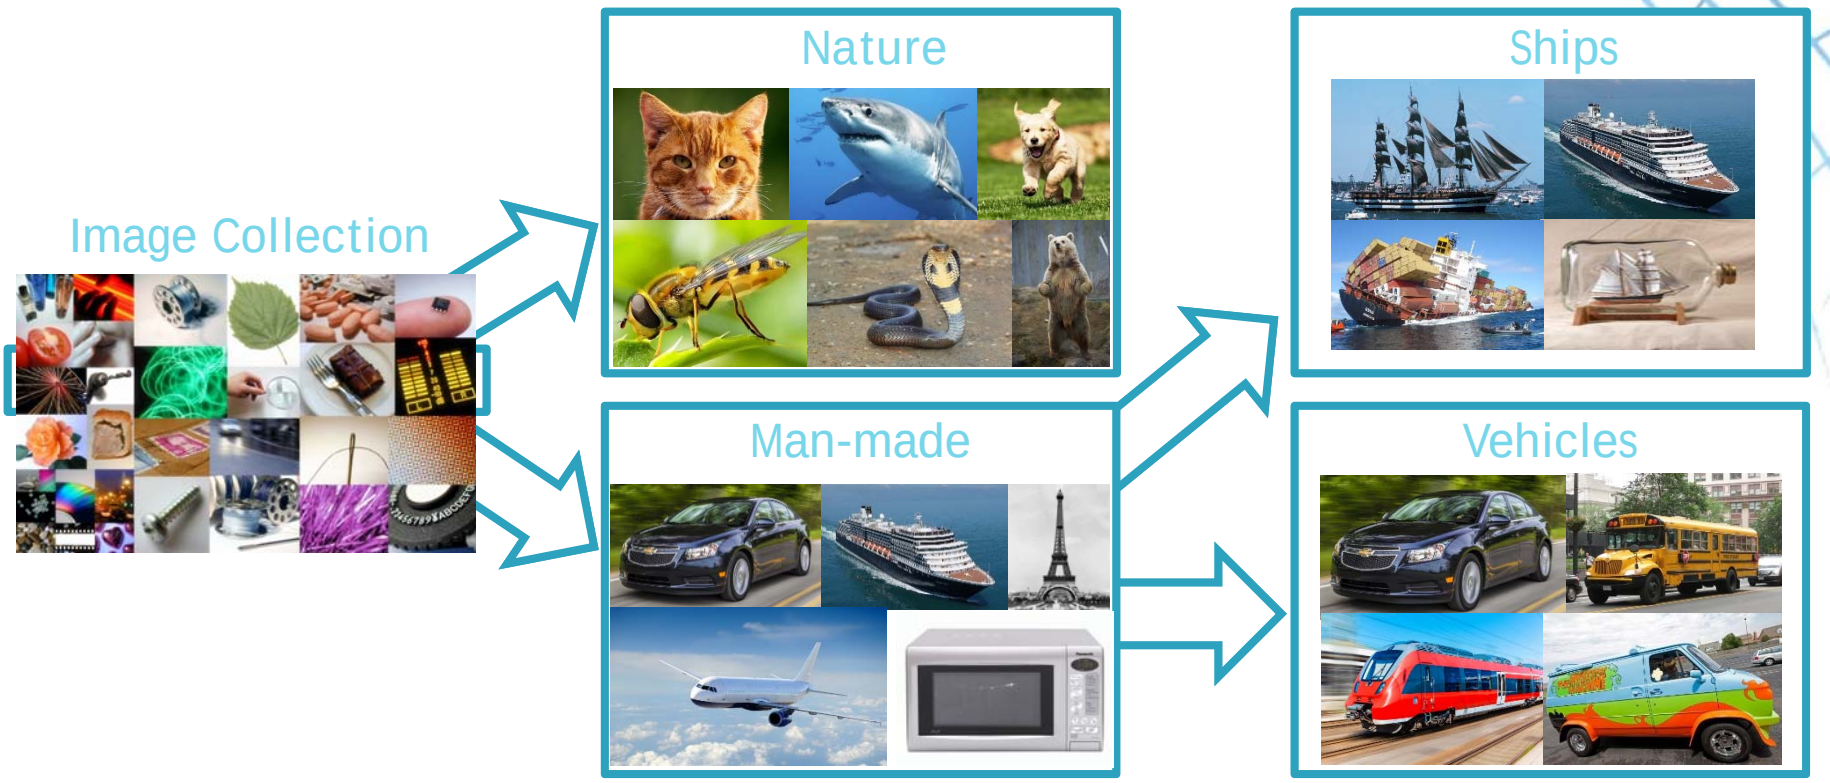

General Idea

- We want to find good representatives of each class
- Then subsample the data

Non-Linear DR with Landmarks

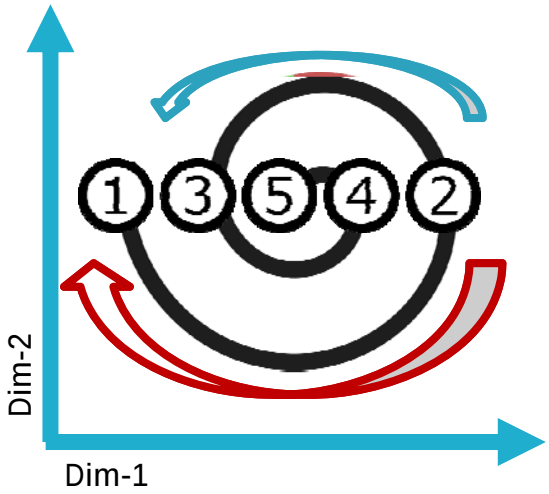

Non linear

[Landmark-SNE, Landmark-ISOMAP]

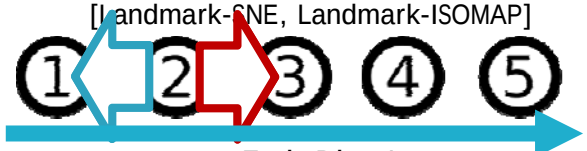

Hybrid techniques

[LSP, P-LSP, LAMP, LoCH, PekaIska]

DELFT DATA SCIENCE

[Pezzotti, Hoelt, Leliveldt, Eisemann, Vilanova – EuroVis 2016]

Hierarchical Stochastic Neighbor Embedding

- Multiscale Dimensionality Reduction
  - Non-linear DR
  - Landmark based
- Hierachical exploration of the data
  - Overview-first & Details-on-Demand
  - Filter & Drill-in

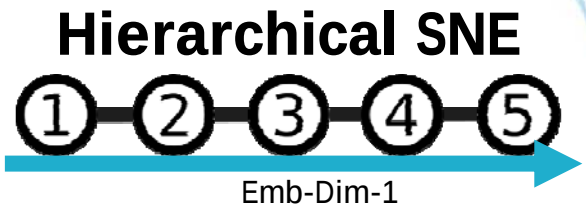

DELFT DATA SCIENCE

[Pezzotti, Hoelt, Leliveldt, Eisemann, Vilanova – EuroVis 2016]
[Pezzotti, Leliveldt, van der Maaten, Hoelt, Eisemann, Vilanova – TVCG2016]

Hierarchical Dimensionality Reduction

- Organize Data at Different Scales

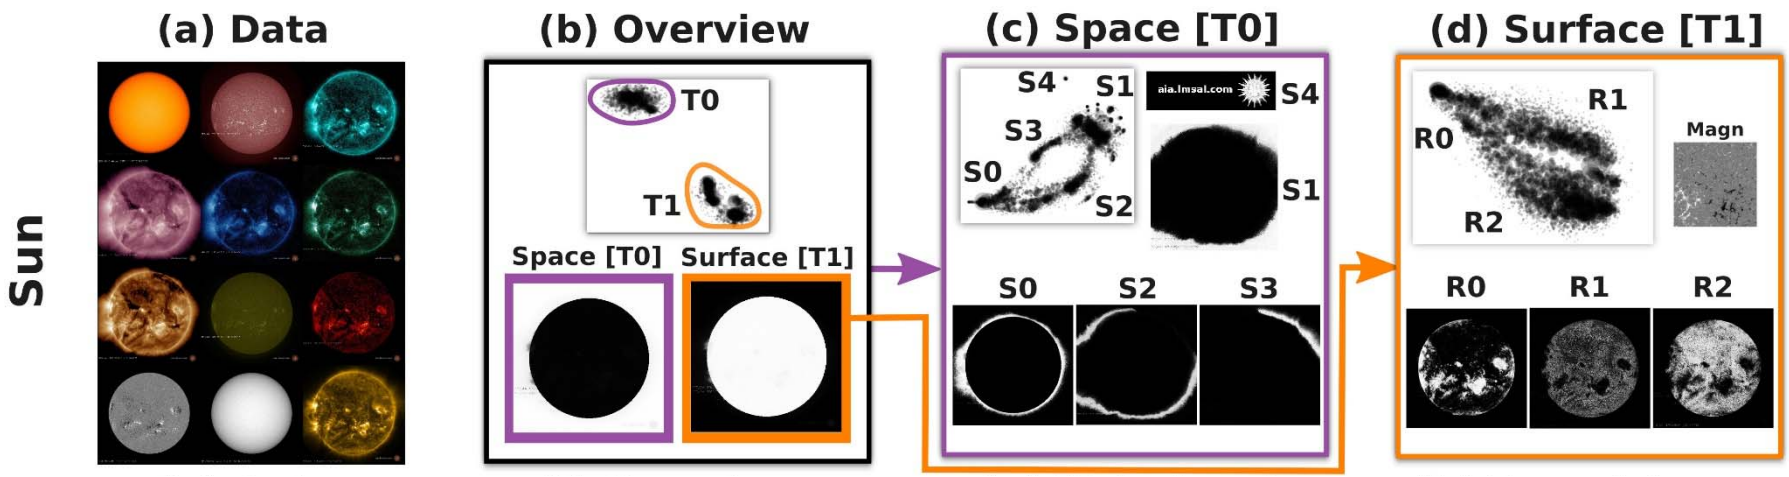

DELFT DATA SCIENCE

[Pezzotti, Hoelt, Leliveldt, Eisemann, Vilanova – EuroVis 2016]
[Pezzotti, Leliveldt, van der Maaten, Hoelt, Eisemann, Vilanova – TVCG2016]

Hierarchical Dimensionality Reduction

- Organize Data at Different Scales

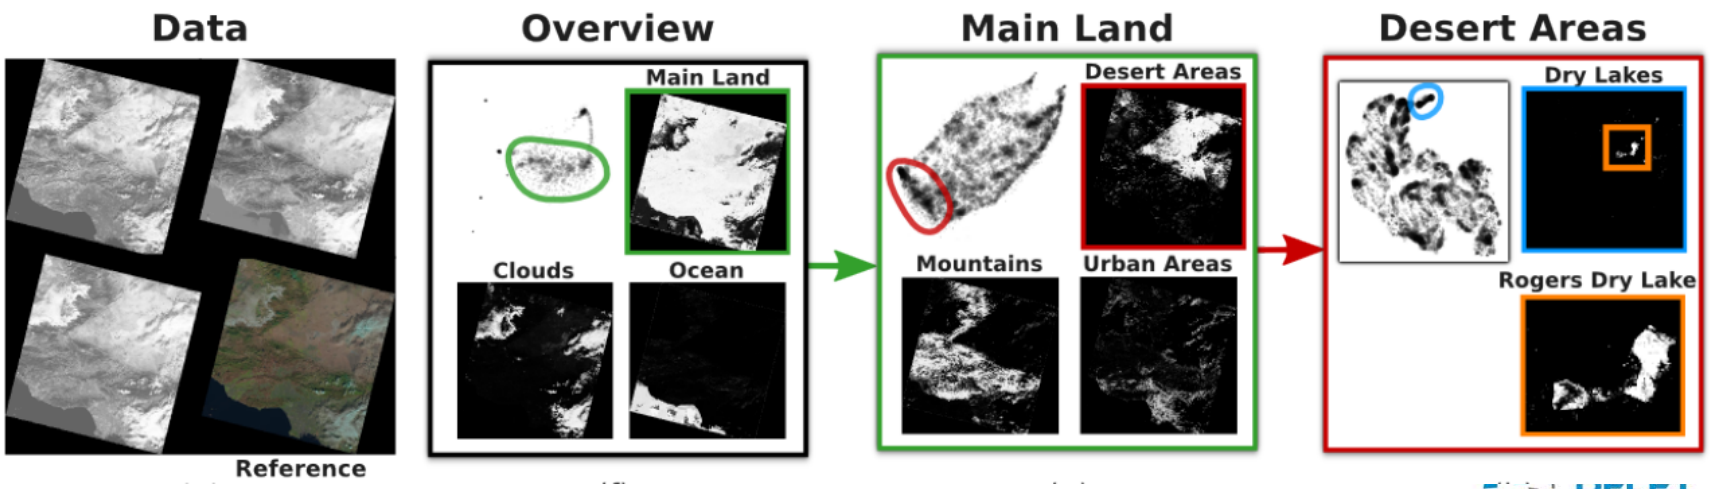

DELFT DATA SCIENCE

"Understanding" of relationships

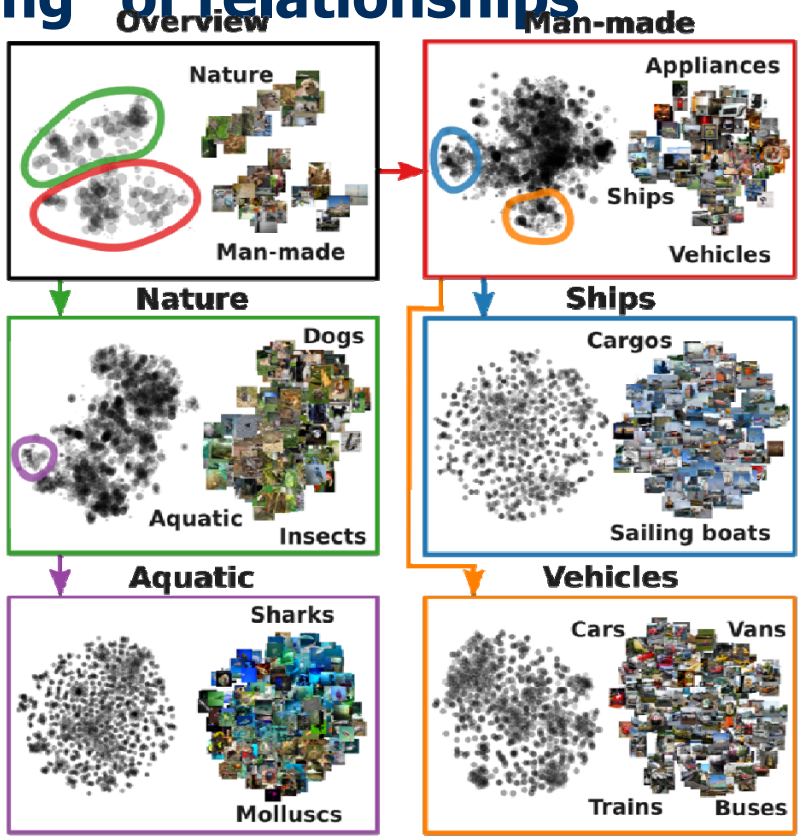

DELFT DATA SCIENCE

Effective Data Visualization Requires

- Large-Scale Rendering

- Visualization and Percepti

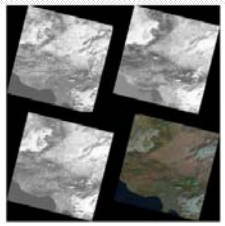

- Data Analysis

DELFT DATA SCIENCE

Data Development

- **Processing**
- **Analysis**
- **Interaction**
- **Visualization**
- **Guidance**

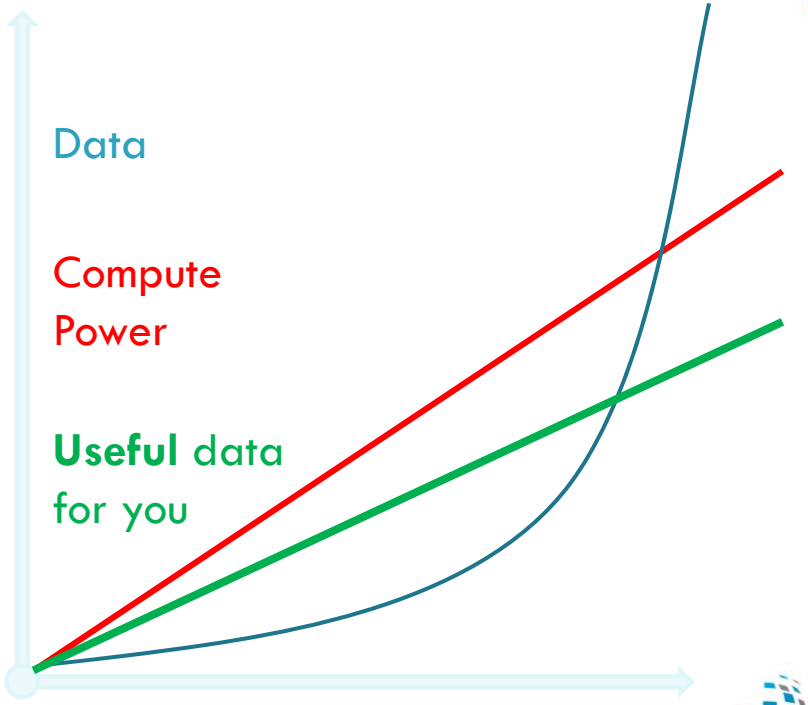

DELFT DATA SCIENCE

Large-Scale Rendering

- Ray Tracing
- Graphics Pipeline
- Specialized Methods for Different Data Types
  - Height-Field Data, Voxel Data, Data Management, Compression

Visualization and Perception

- Realistic Rendering
- Perceptual Methods
- Visualization & Interface

Data Analysis

- High-dimensional/Heterogeneous Data
- Dimensionality Reduction
- Visual Analytics

Effective Data Visualization Requires

- Large-Scale Rendering

- Visualization and Percepti

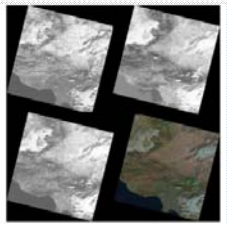

- Data Analysis

DELFT DATA SCIENCE

TU Delft

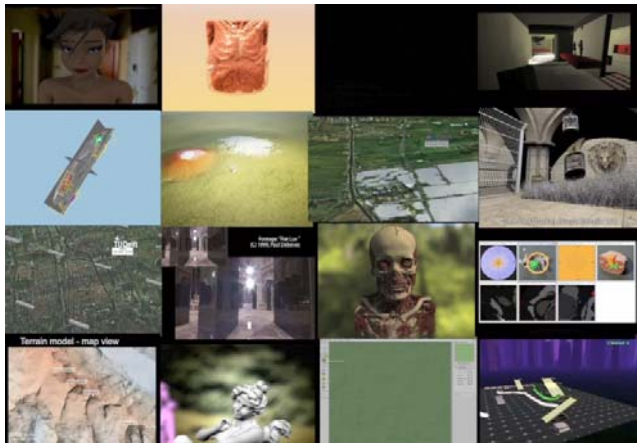

Thank you very much for your attention!

Elmar Eisemann
e.eisemann@tudelft.nl
http://graphics.tudelft.nl/~eisemann