Representations and Algorithms for Interactive Relighting

Nick Michiels

Promoter: Prof. Dr. Philippe Bekaert

Representations and Algorithms for Interactive Relighting

[Matterport, 2016]

Real Estate

[Hilsmann et al., Eurograhics 2013]

Clothing Industry

[Microsoft Hololens, 2016]

Augmented Reality

universiteit hasselt | EDM | iMinds

Representations and Algorithms for Interactive Relighting

immersive experience
view-dependent

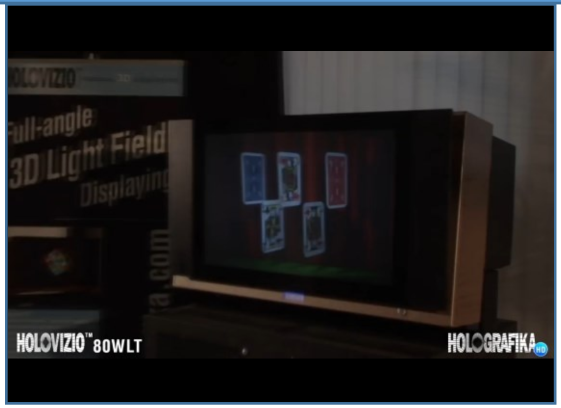

Representations and Algorithms for Interactive Relighting

universiteit hasselt | EDM | iMinds

Representations and Algorithms for Interactive Relighting

Representations and Algorithms for Interactive Relighting

Representations and Algorithms for Interactive Relighting

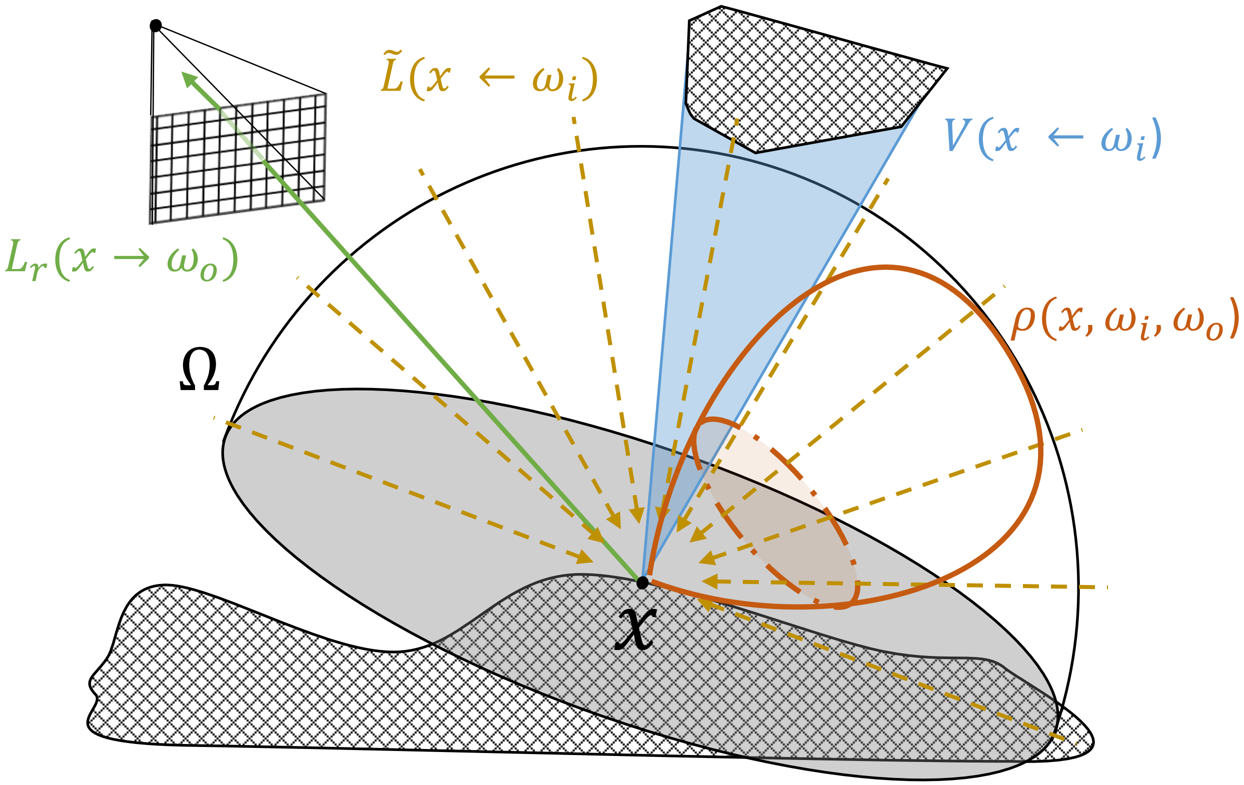

$$L_r(x \rightarrow \omega_o) = \int_\Omega V(x \leftarrow \omega_i) * \tilde{L}(x \leftarrow \omega_i) * \tilde{\rho}(x, \omega_i, \omega_o)\, d\omega_i$$

Representations and Algorithms for Interactive Relighting

$$L_r(x \to \omega_o) = \int_\Omega V(x \leftarrow \omega_i) * \tilde{L}(x \leftarrow \omega_i) * \tilde{\rho}(x, \omega_i, \omega_o)\ d\omega_i$$

Forward Rendering

Relighting

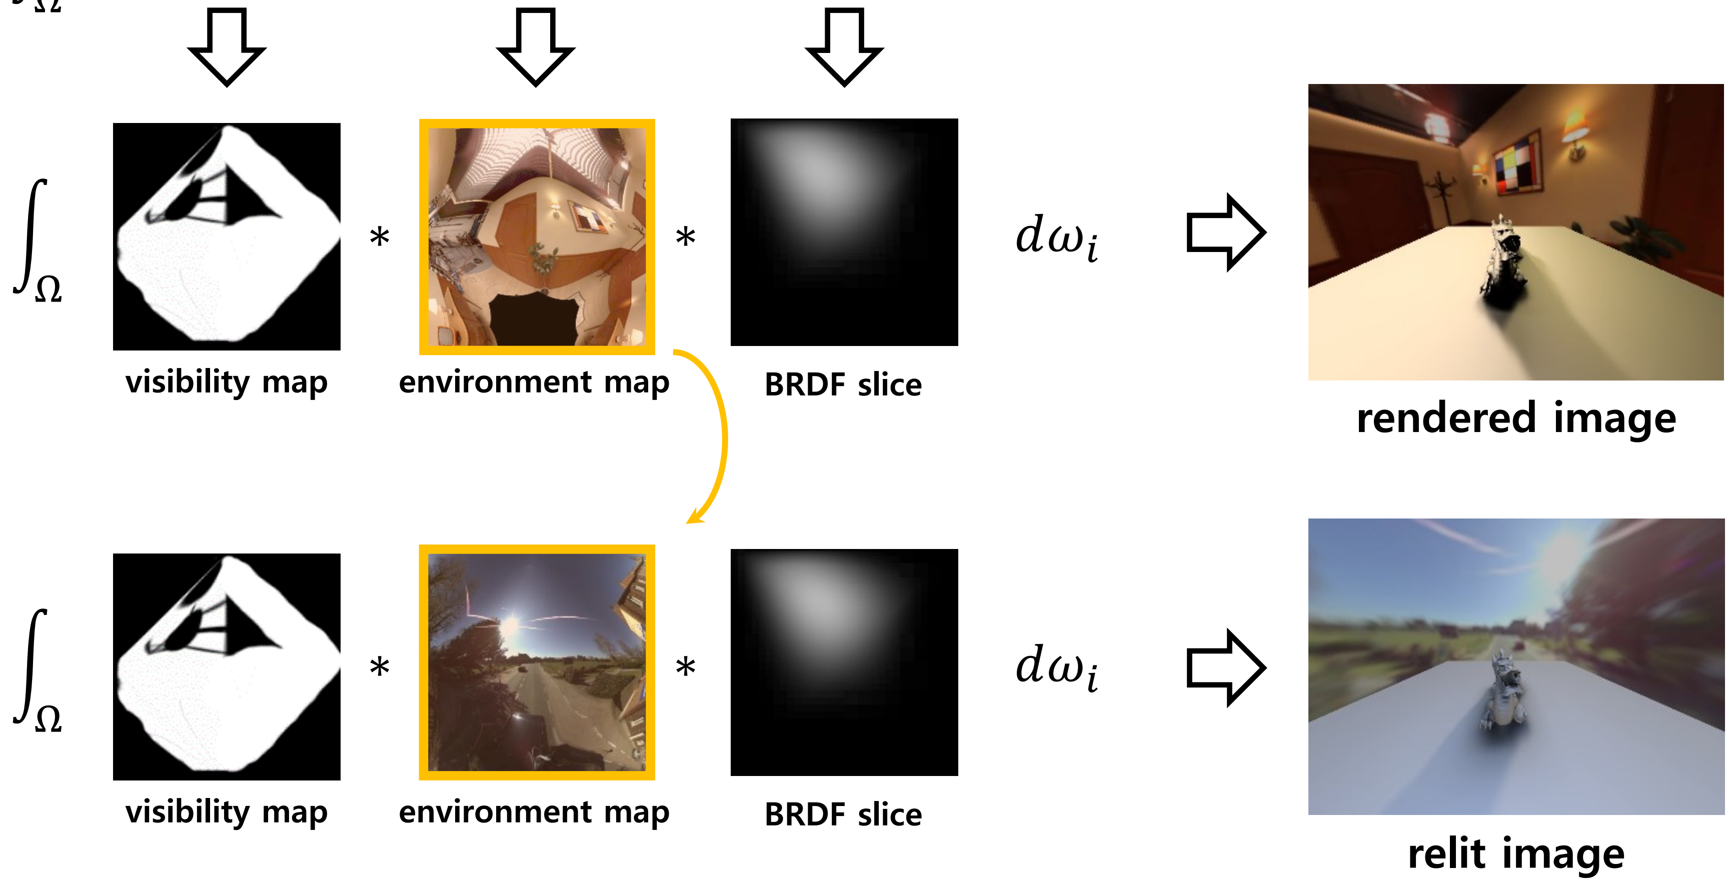

Representations and Algorithms for Interactive Relighting

1. Relighting of **Virtual Objects**

2. Relighting of **Real Objects**

Key Challenges

1. Texture-illumination ambiguity

2. Simulation of light propagation

Triple Product Integral

$$L_r(x \to \omega_o) = \int_\Omega V(x \leftarrow \omega) \ * \ \tilde{L}(x \leftarrow \omega) \ * \ \tilde{\rho}(x, \omega, \omega_o) \ d\omega$$

pixel domain $= \int_\Omega$ 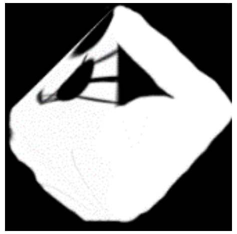 $*$ 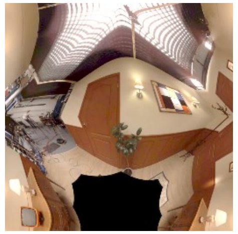 $*$ 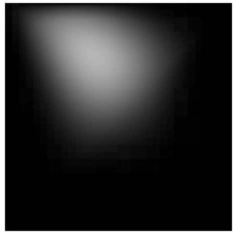 $d\omega$

new basis representation

$$= \int_\Omega \sum_i V_i \, \Psi_i(\omega) \ * \ \sum_j L_j \, \Psi_j(\omega) \ * \ \sum_k \tilde{\rho}_k \, \Psi_k(\omega) \ d\omega$$

triple product binding coefficients

$$= \sum_i \sum_j \sum_k V_i \tilde{L}_j \tilde{\rho}_k C_{ijk}$$

$$C_{ijk} = \int_\Omega \Psi_i(\omega) \Psi_j(\omega) \Psi_k(\omega) d\omega$$

universiteit hasselt | EDM | iMinds

Spherical Harmonics

- Original approach *[Sloan et al., 2002]*
- Equivalent of Fourier series on the sphere
- Set of orthogonal functions
- Linear combination of sine and cosine waves

+ Efficient

- Low-frequency lighting effects only

2D Haar Wavelets

- Haar Tripling Coefficient Theorem *[Ng et al., 2004]*
- Piecewise constant functions
- Orthonormal basis

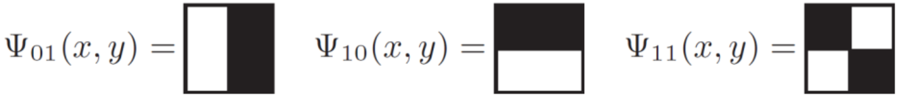

+ Few coefficients
+ All-frequency lighting

Proposed Representations

1. High-Order Wavelets

2. Spherical Radial Basis Functions

Why high-order wavelets?

- Representation should be tailored to the signal
- Smooth high-order wavelets (e.g. Daubechies-4) require an order of magnitude less coefficients to represent a smooth signal

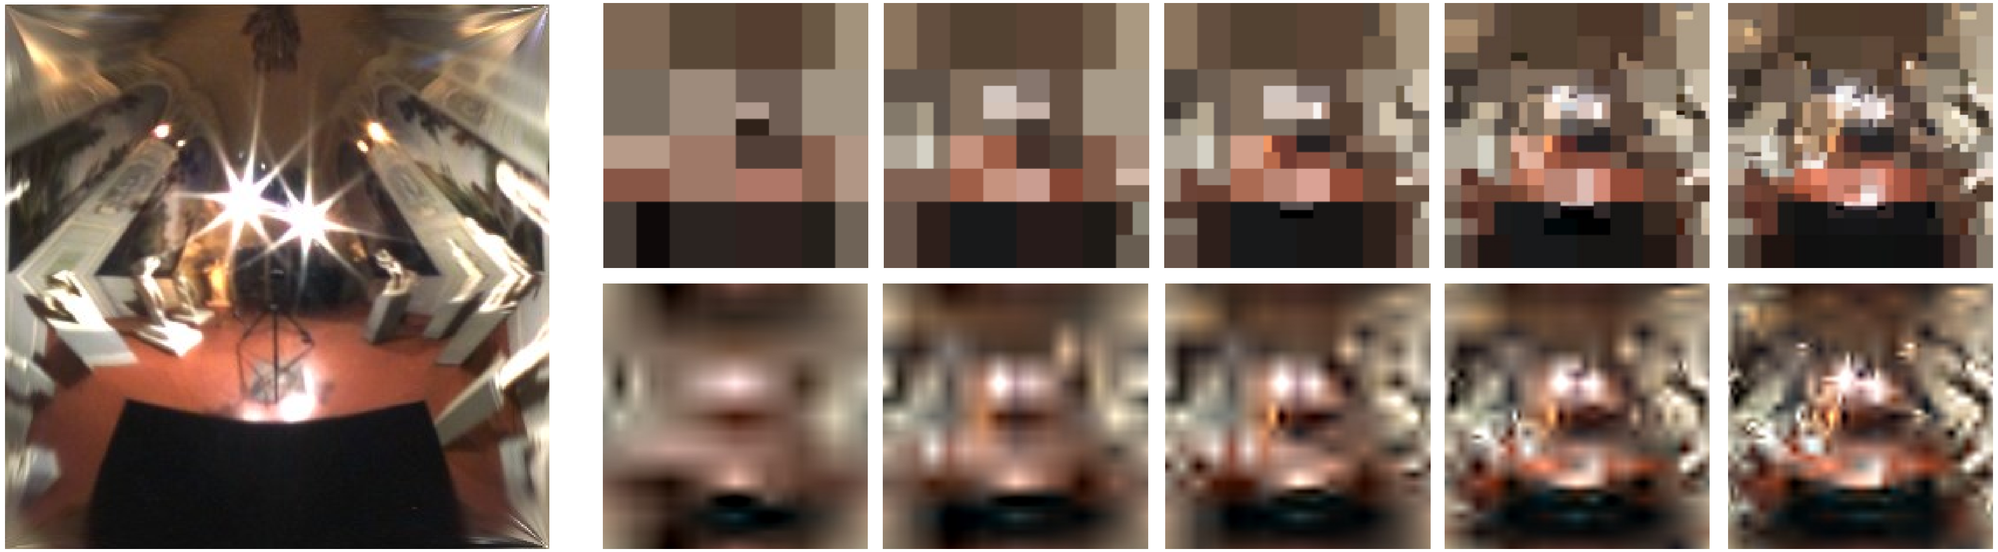

Haar

Daubechies

environment lighting | 16 | 32 | 64 | 128 | 256

inceasing number of coefficients

Why is it so difficult to use high-order wavelets?

$$C_{ijk} = \int_{\Omega} \Psi_i(\omega)\Psi_j(\omega)\Psi_k(\omega)d\omega$$

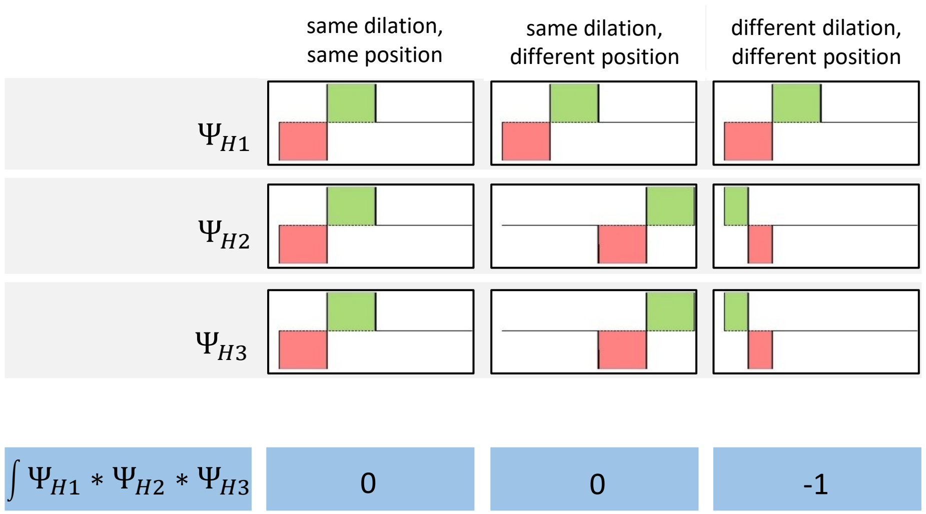

Haar wavelets

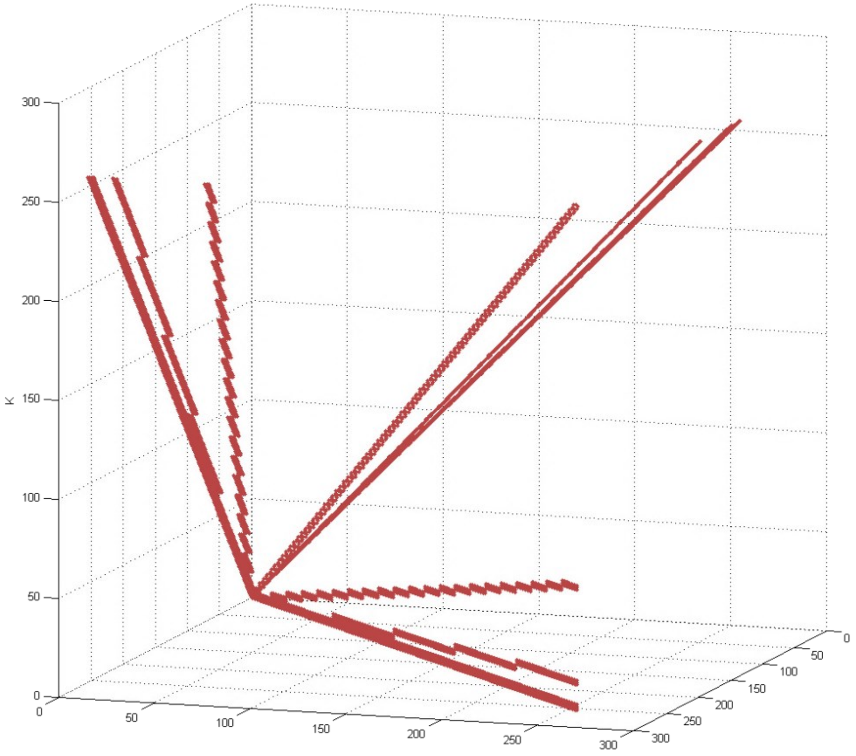

tensor of binding coefficients

Why is it so difficult to use high-order wavelets?

$$C_{ijk} = \int_{\Omega} \Psi_i(\omega)\Psi_j(\omega)\Psi_k(\omega)d\omega$$

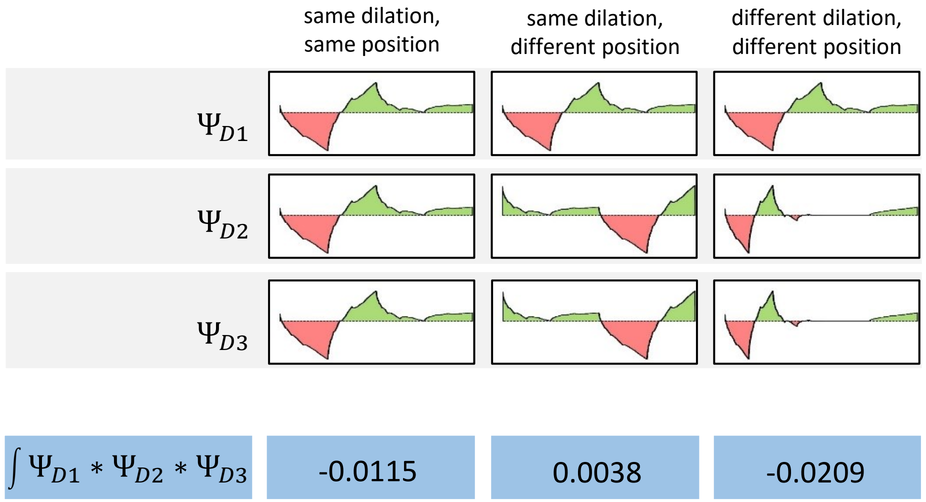

Daubechies wavelets

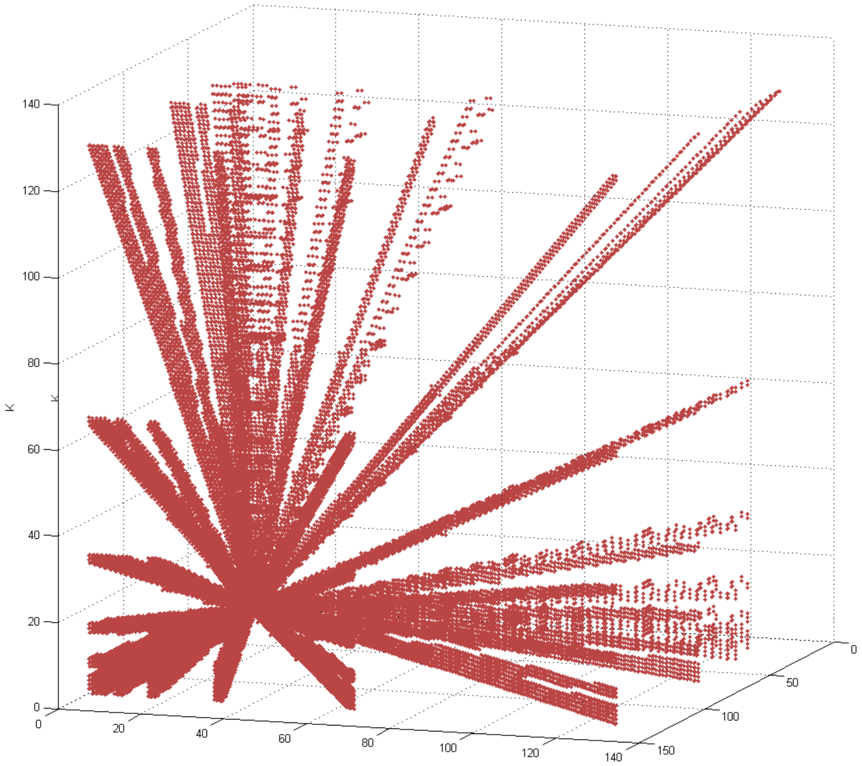

tensor of binding coefficients

General Wavelet Tripling Coefficient Theorem

- Naive approach
- Hierarchical approach
- Symmetry
- Wavelet sliding
- Vanishing moments

Test example

- D: dimensionality of the integral (double, triple, quadruple)
- S_i: signal of $\mathrm{r}\ \mathrm{x}\ \mathrm{r}$ resolution $(i = 1, \cdots, D)$
- N: number of dilations and translations of the basis function for S_i

General Wavelet Tripling Coefficient Theorem

- Naive approach
  - Iterate over all binding coefficients

  - $O(N^D r^2)$
    - **$4,7 \times 10^{21}$ operations**
    - $r = 512, N = 262144, D = 3$

General Wavelet Tripling Coefficient Theorem

- Hierarchical approach
  - Exploiting support of wavelet
  - Non-overlapping cases can be skipped

  - $O(NC\,(\log N)^{D-1}\,r^2)$
    - C relates to the enlargement of support
    - **$3,2 \times 10^{13}$ operations**
    - $r = 512, N = 262144, D = 3$

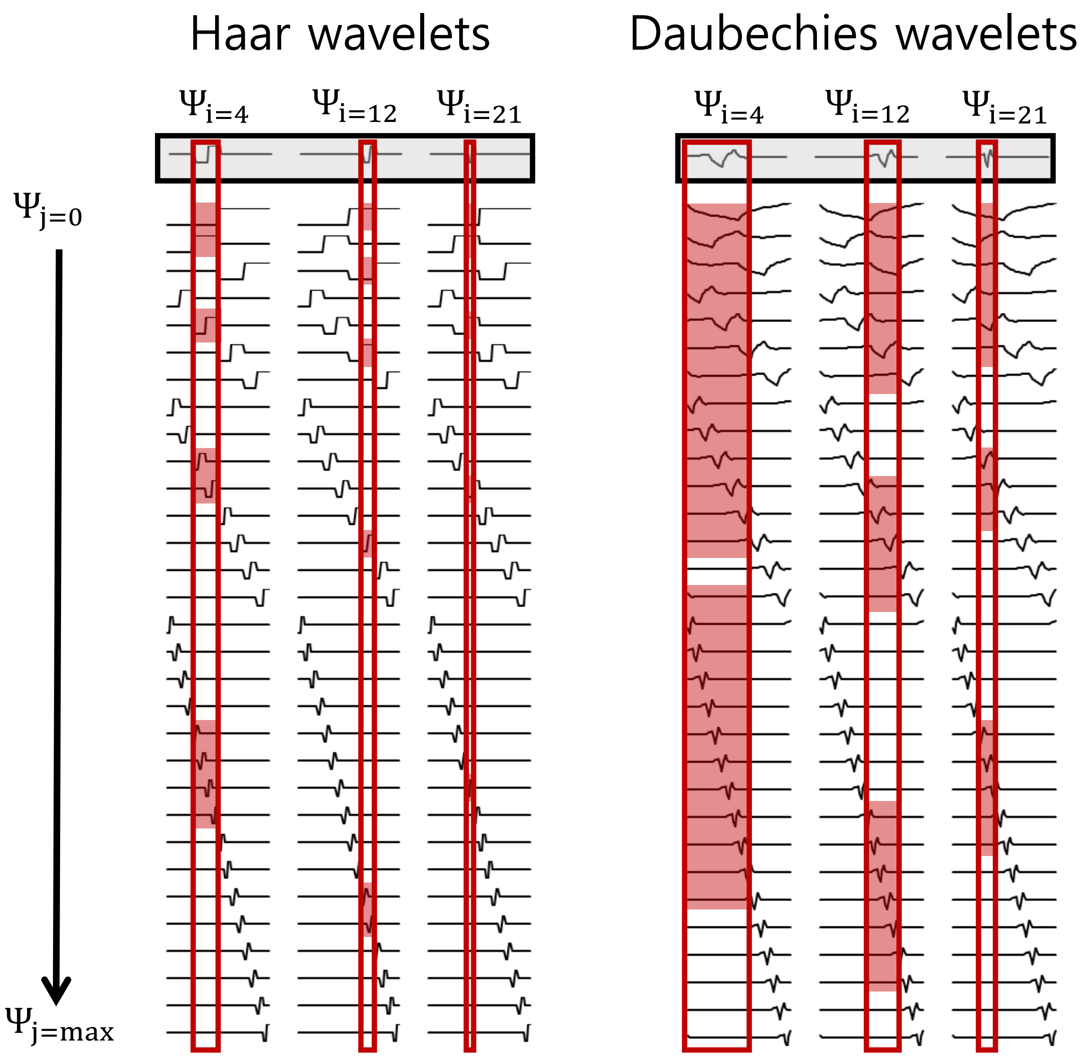

universiteit hasselt | EDM | iMinds

General Wavelet Tripling Coefficient Theorem

- Symmetry
  - Homogeneous products
  - $\int \Psi_i \Psi_j = \int \Psi_j \Psi_i$

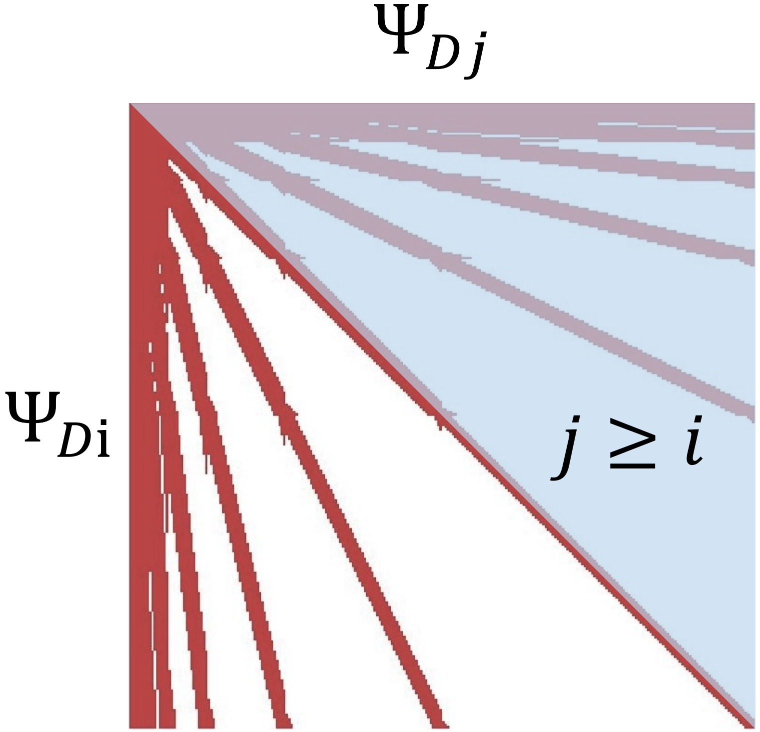

General Wavelet Tripling Coefficient Theorem

- Wavelet sliding
  - Reuse duplicate branches in the tensor

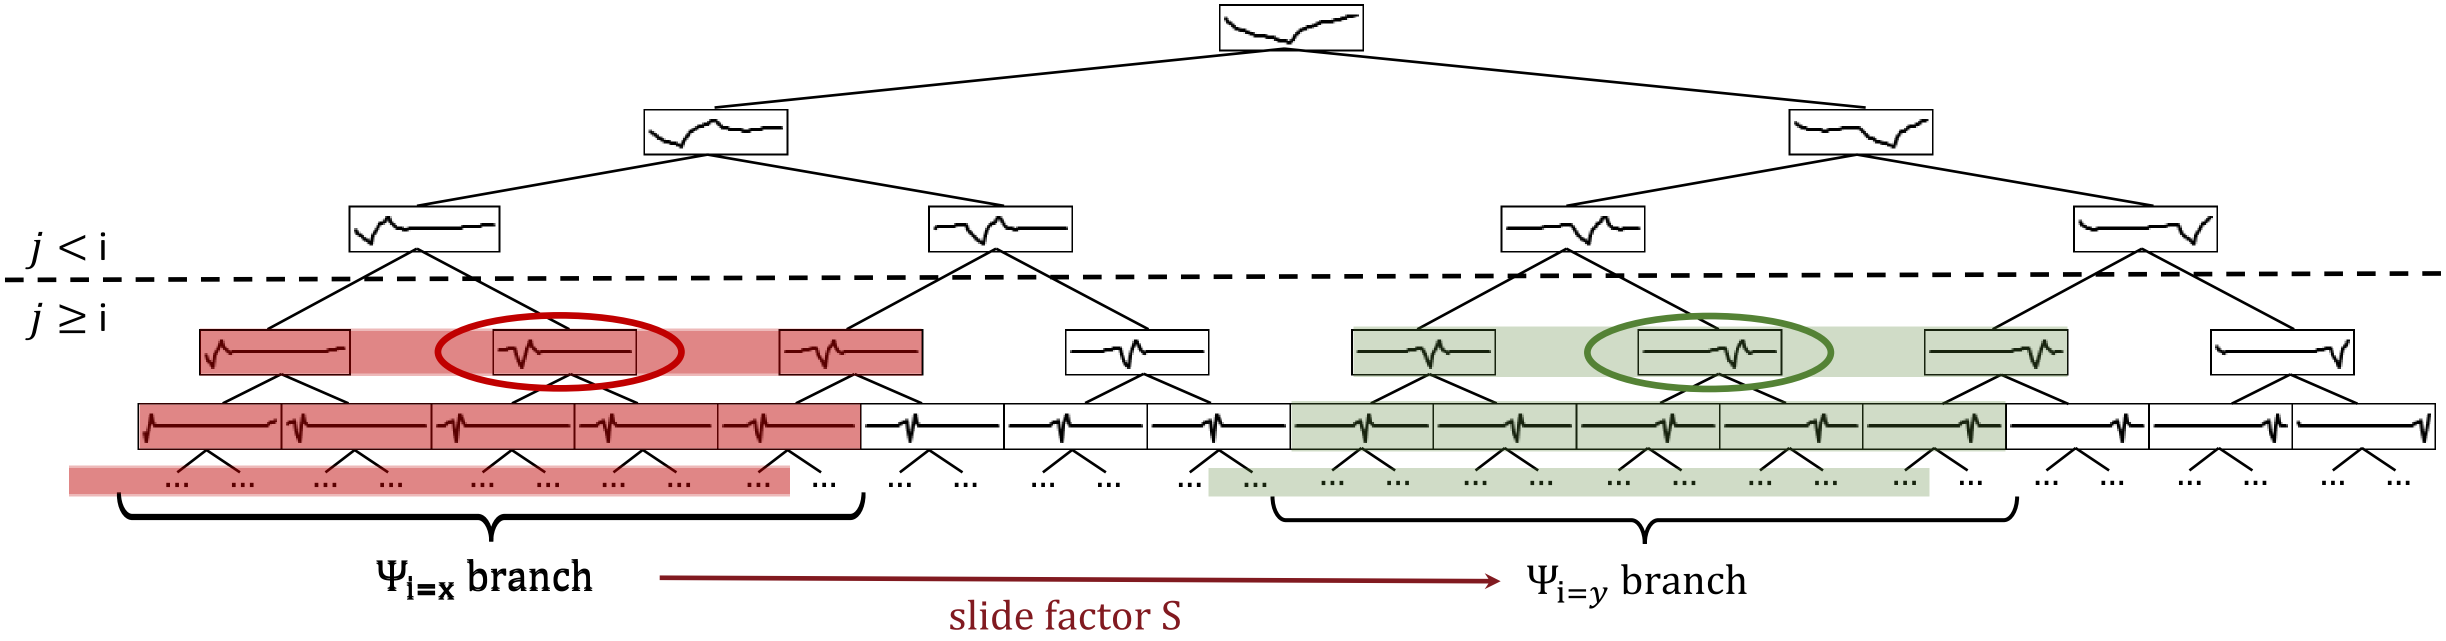

$$\text{slide factor } S = (y - x) \times \text{support}(\Psi_{i=x})$$

General Wavelet Tripling Coefficient Theorem

- Wavelet sliding
  - Reuse duplicate branches in the tensor
  - Calculate products for $\log N$ branches instead of N branches

  - $O(C\,(\log N)^D\,r^2)$
    - C relates to the enlargement of support
    - **$1,5 \times 10^9$ operations**
    - $r = 512, N = 262144, D = 3$

universiteit hasselt | EDM | iMinds

General Wavelet Tripling Coefficient Theorem

- Vanishing moments
  - High-order wavelets provide more vanishing moments
  - Causes zero integrals for certain translations within their support
  - Increased sparsity in tensor of binding coefficients
  - Identified and incorporated in wavelet sliding algorithm

Sparseness – 8×8

Ψ_i	Ψ_j	Ψ_k	resolution	non-zero coeffs	total coeffs	sparseness
Haar-2	Haar-2	Haar-2	8×8	1288	2.6E+5	**0.4913%**
Daub-4	Daub-4	Daub-4	8×8	99088	2.6E+5	**37.7991%**
Daub-6	Daub-6	Daub-6	8×8	214252	2.6E+5	**81.7307%**
Coiflet-5	Coiflet-5	Coiflet-5	8×8	186706	2.6E+5	**71.2227%**
Haar-2	Daub-4	Daub-4	8×8	31960	2.6E+5	**12.1918%**
Haar-2	Coiflet-5	Coiflet-5	8×8	59902	2.6E+5	**22.8508%**

Sparseness - 16 × 16

Ψ_i	Ψ_j	Ψ_k	resolution	non-zero coeffs	total coeffs	sparseness
Haar-2	Haar-2	Haar-2	16 × 16	1288	1.6E+7	**0.0443%**
Daub-4	Daub-4	Daub-4	16 × 16	99088	1.6E+7	**8.2263%**
Daub-6	Daub-6	Daub-6	16 × 16	214252	1.6E+7	**25.9928%**
Coiflet-5	Coiflet-5	Coiflet-5	16 × 16	186706	1.6E+7	**20.8433%**
Haar-2	Daub-4	Daub-4	16 × 16	31960	1.6E+7	**1.6133%**
Haar-2	Coiflet-5	Coiflet-5	16 × 16	59902	1.6E+7	**3.7315%**

Sparseness - 32 × 32

Ψ_i	Ψ_j	Ψ_k	resolution	non-zero coeffs	total coeffs	sparseness
Haar-2	Haar-2	Haar-2	32 × 32	38920	1.1E+9	**0.0036%**
Daub-4	Daub-4	Daub-4	32 × 32	9720040	1.1E+9	**0.9052%**
Daub-6	Daub-6	Daub-6	32 × 32	29582032	1.1E+9	**2.7550%**
Coiflet-5	Coiflet-5	Coiflet-5	32 × 32	16408816	1.1E+9	**1.5282%**
Haar-2	Daub-4	Daub-4	32 × 32	1555120	1.1E+9	**0.1448%**
Haar-2	Coiflet-5	Coiflet-5	32 × 32	2715112	1.1E+9	**0.2529%**

Sparseness - 64×64

Ψ_i	Ψ_j	Ψ_k	resolution	non-zero coeffs	total coeffs	sparseness
Haar-2	Haar-2	Haar-2	64×64	192520	6.9E+10	**0.0003%**
Daub-4	Daub-4	Daub-4	64×64	47918464	6.9E+10	**0.0697%**
Daub-6	Daub-6	Daub-6	64×64	145473456	6.9E+10	**0.2117%**
Coiflet-5	Coiflet-5	Coiflet-5	64×64	48918464	6.9E+10	**0.0712%**
Haar-2	Daub-4	Daub-4	64×64	7327168	6.9E+10	**0.0107%**
Haar-2	Coiflet-5	Coiflet-5	64×64	8699044	6.9E+10	**0.0127%**

Render Application

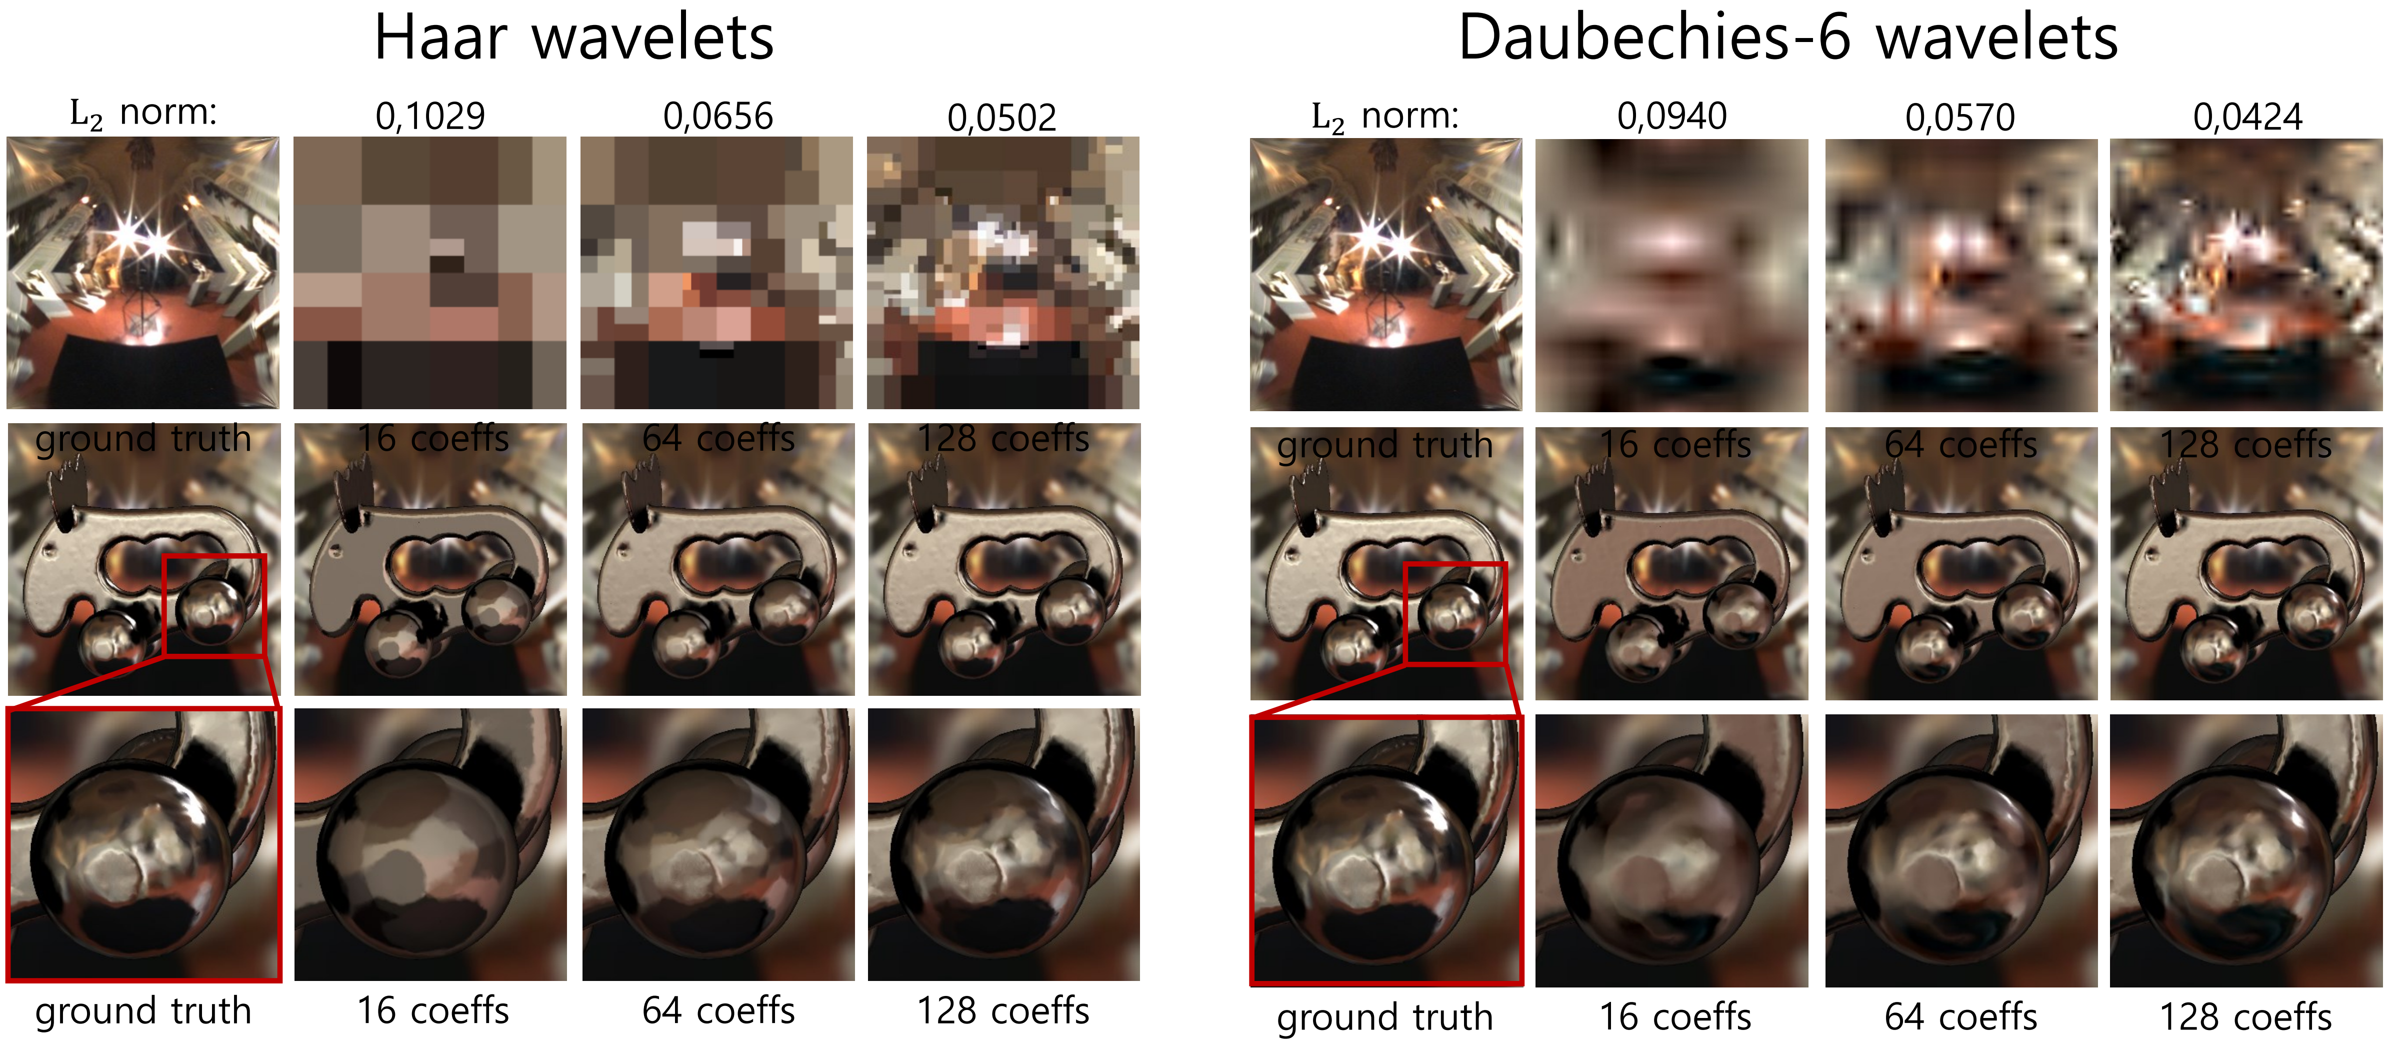

What Is Wrong with Wavelets?

- 2D Haar wavelets *[Ng et al, 2004]*
- Spherical Haar wavelets *[Put et al., 2014]*
- High-order wavelets *[Michiels et al., 2014]*

+ Few coefficients
+ All-frequency
- Preprocessing
- No efficient rotation operator *[Wang et al., 2006]*

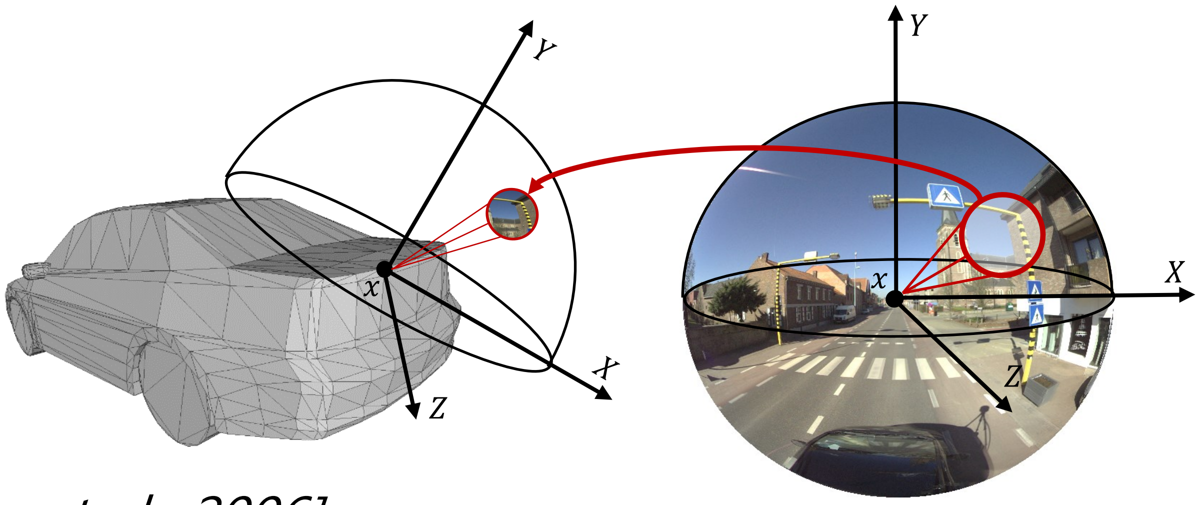

Contributions

1. High-Order Wavelets

2. Spherical Radial Basis Functions

Spherical Radial Basis Functions (SRBFs)

- Radial Basis Functions, defined on the sphere
  - $\text{Poisson}(\omega, \boldsymbol{c}, \lambda, \mu) = \mu \frac{1-\lambda^2}{(1-2\lambda(\omega\cdot\boldsymbol{c})+\lambda^2)^{3/2}}$
  - $\text{Multiquadratic}(\omega, \boldsymbol{c}, \lambda, \mu) = \mu\sqrt{1 + (\lambda(\omega \cdot \boldsymbol{c}))^2}$
  - $\boldsymbol{Gaussian(\omega, c, \lambda, \mu) = \mu e^{\lambda(\omega\cdot c-1)}}$

+ All-frequency
+ Decent compression performance
+ Efficient rotation operator
+ Analytic evaluation of the binding coefficients

Triple product rendering with SRBFs

	all-frequency	dynamic lighting	dynamic geometry	dynamic brdf
Sloan et al., 2003	×	×	×	×
Ng et al., 2004	✓	×	×	×
Tsai and Shih, 2006	✓	rotation only	×	×
Haber et al., 2009	✓	×	×	×
Wang et al., 2009	✓	rotation only	×	✓
Lam et al., 2010	✓	rotation only	×	×
Iwasaki et al., 2012	✓	rotation only	low-poly	✓
our method	✓	✓	✓	✓

- Current techniques constrain one or several factors

+ Our approach is able to dynamically construct and update all three factors

Dynamic Materials

SRBFs

+ Directly approximated with Gaussian lobes
+ Phong, Cook-Torrance, Ward, Blinn-Phong, …

[Wang et al., 2009]

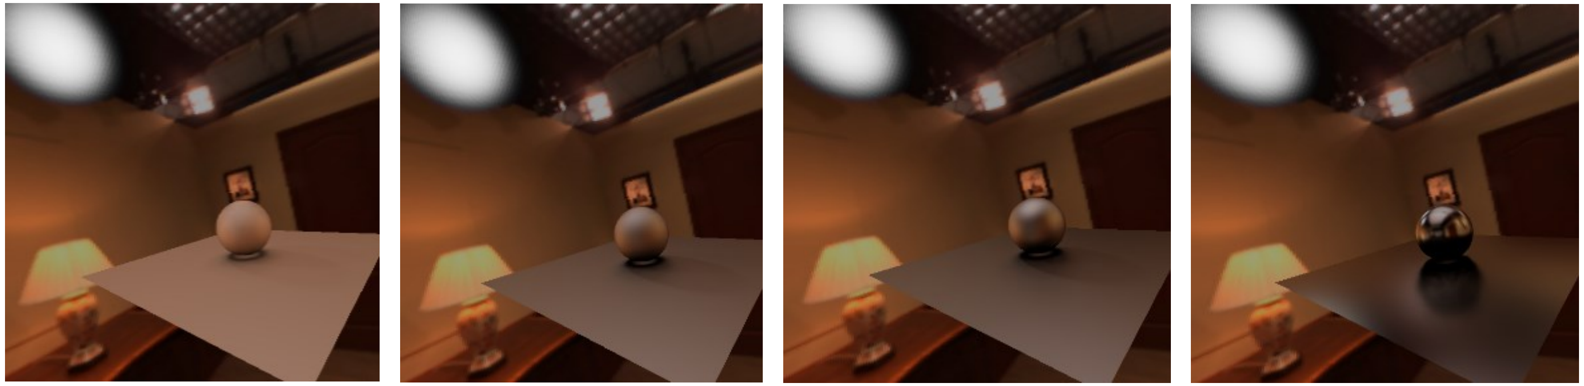

Lambertian *Diffuse Phong* *Glossy Phong* *Specular Phong*

Previous approaches

- Sampled in pixel domain for each BRDF slice (x, ω_i, ω_o)
- On-the-fly transformation to wavelets

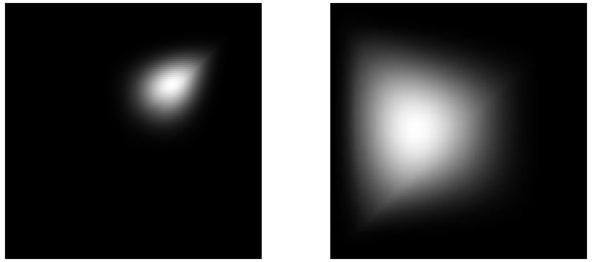

[Haber et al., 2009]

Dynamic Visibility

Our approach (SRBFs)

+ Combination of PRT and voxelization
+ One-pass voxelization *[Crassin and Green, 2012]*
+ Mapping of visibility SRBFs to voxel cones
+ Entirely on GPU

Previous approaches

- Rely on precomputation
- Limited to static scenes
- Ray tracing in pixel domain *[Haber et al., 2009]*
- Approximated with Spherical Signed Distance Functions *[Wang et al., 2009]*
- Projecting bounding volumes on hemisphere *[Iwasaki et al., 2012]*

Dynamic Visibility – Cone tracing

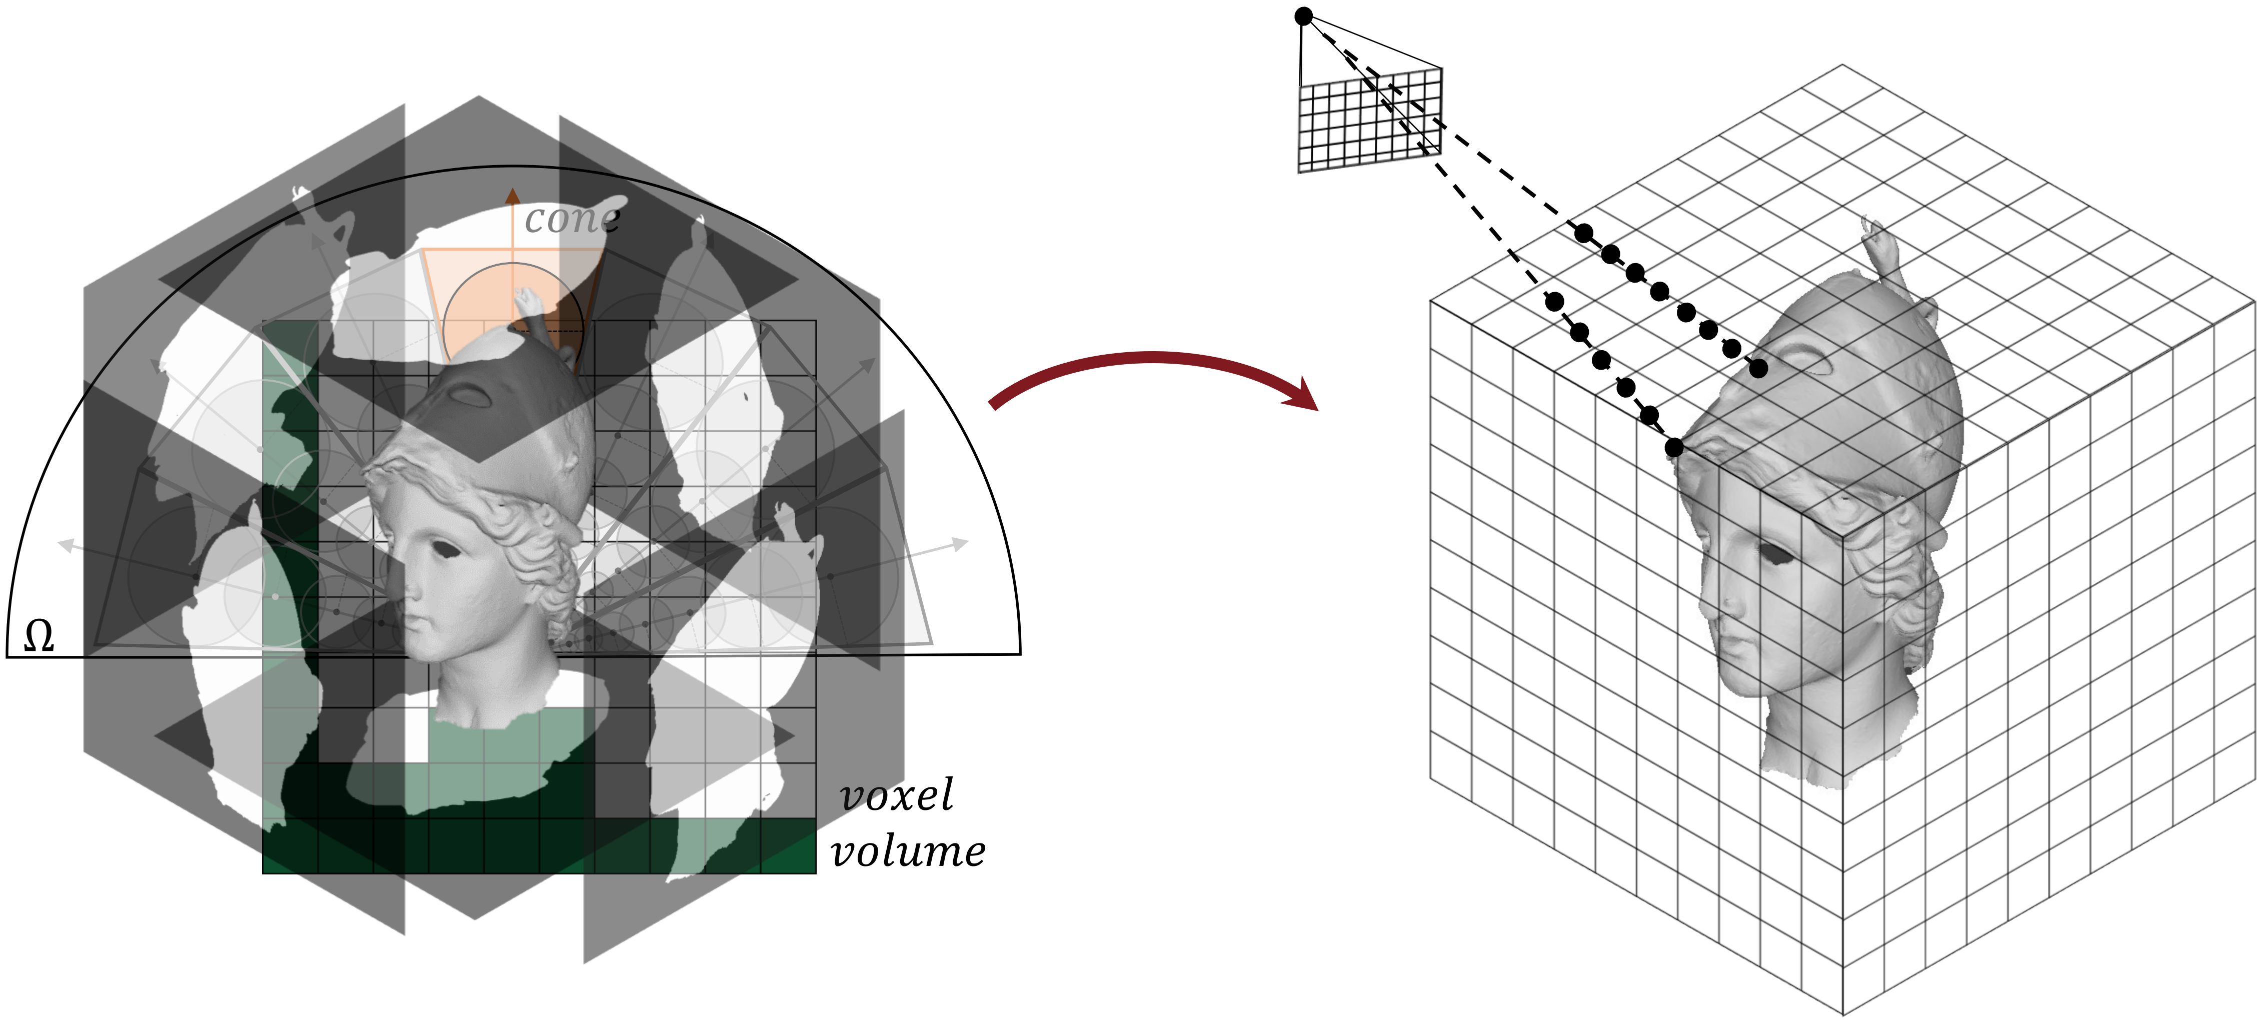

Dynamic Visibility – Mapping SRBF to Visibility Cone

- Sampling of visibility in the SRBF lobe
- Mapping SRBF to a corresponding visibility cone

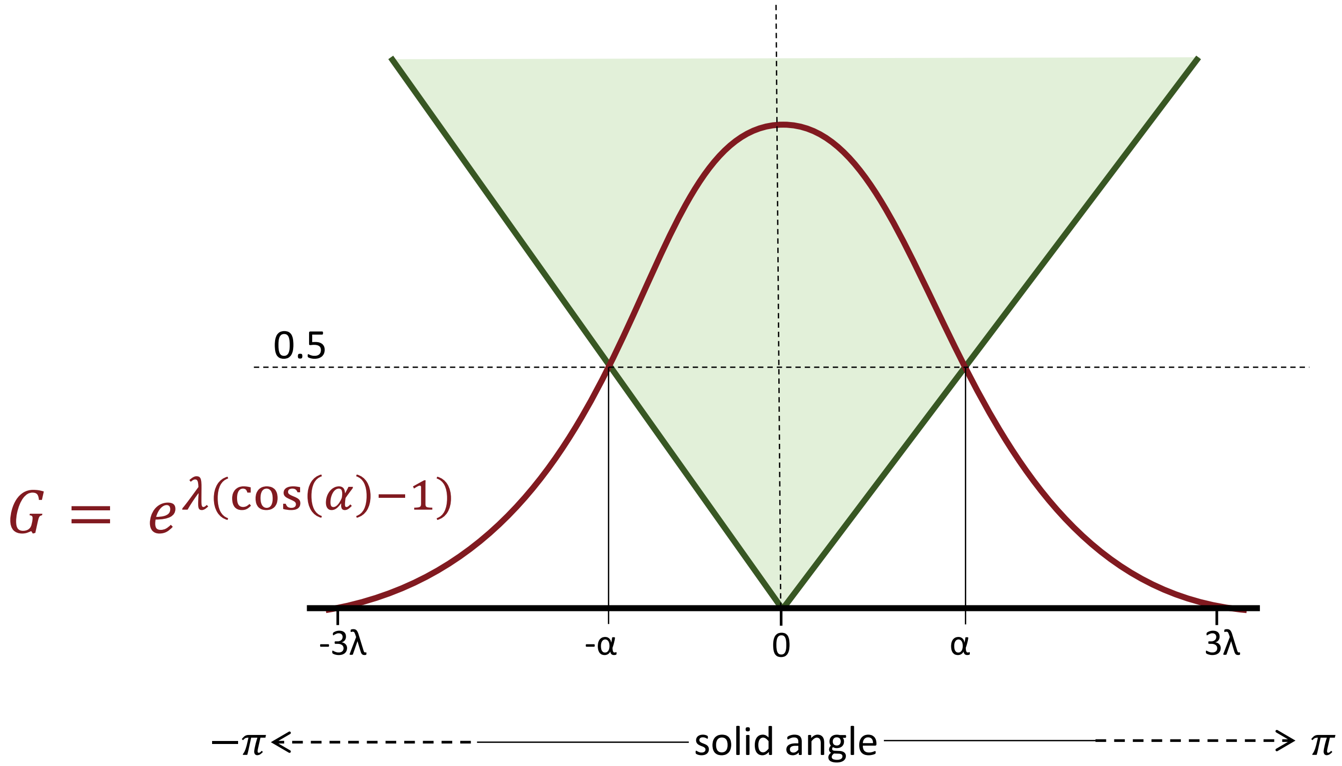

$$G = e^{\lambda(\cos(\alpha)-1)} = 0.5$$

$$\Leftrightarrow \alpha = \cos^{-1}\left(\frac{\log(0.5)}{\lambda} + 1\right)$$

Dynamic Visibility – Subsampling Scheme

- Subsampling of SRBF lobe is essential
  - Avoid integration of high-frequency visibility detail over larger area of the hemisphere
  - Circle packing: maximize density of subsampled cones
  - Adaptive subsampling based on BRDF lobe

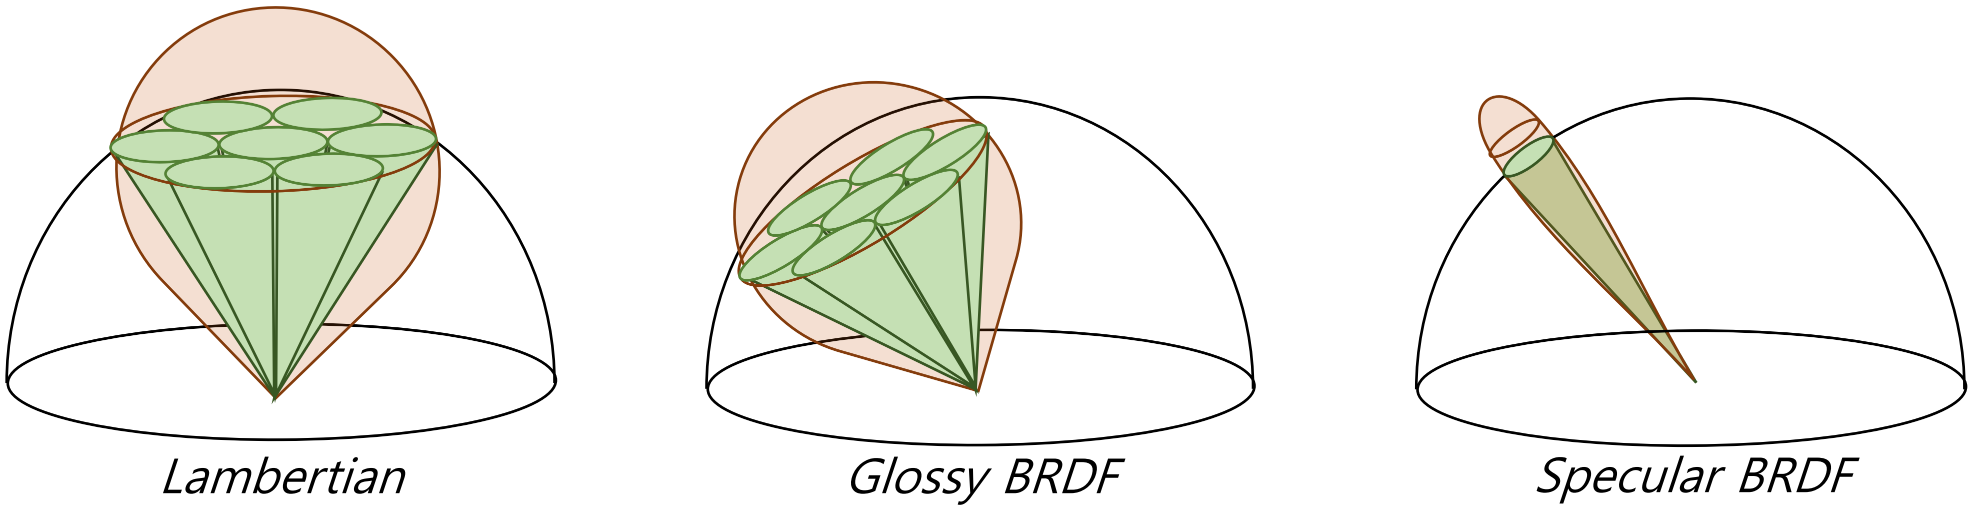

Dynamic Lighting

Our approach (SRBFs)

+ HDR omnidirectional photo/video
+ SRBF fitting
+ Multi-scale algorithm using fixed grid
+ SRBF centers defined by Healpix distribution
+ Entirely on GPU

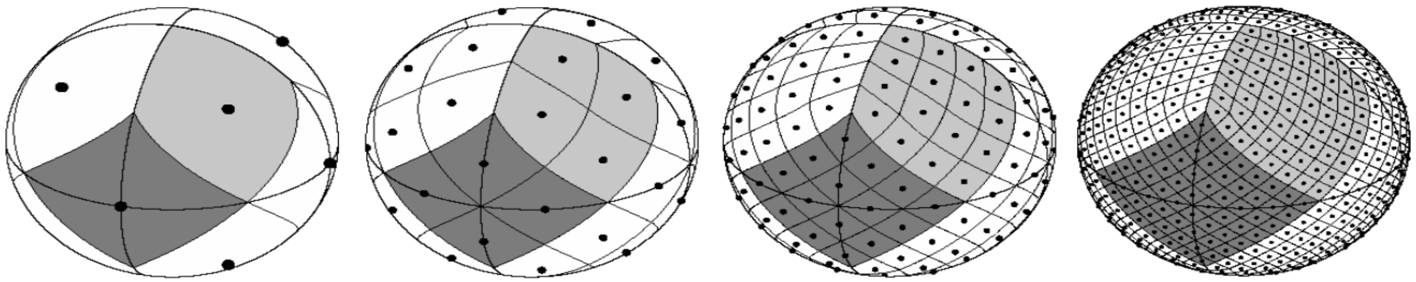

Previous approaches

- Optimization *[Tsai and Shih, 2006]*
  + *Good compression*
  - *Slow*
- Least-square projection *[Lam et al., 2010]*

Dynamic Lighting

- Multi-scale residual transform

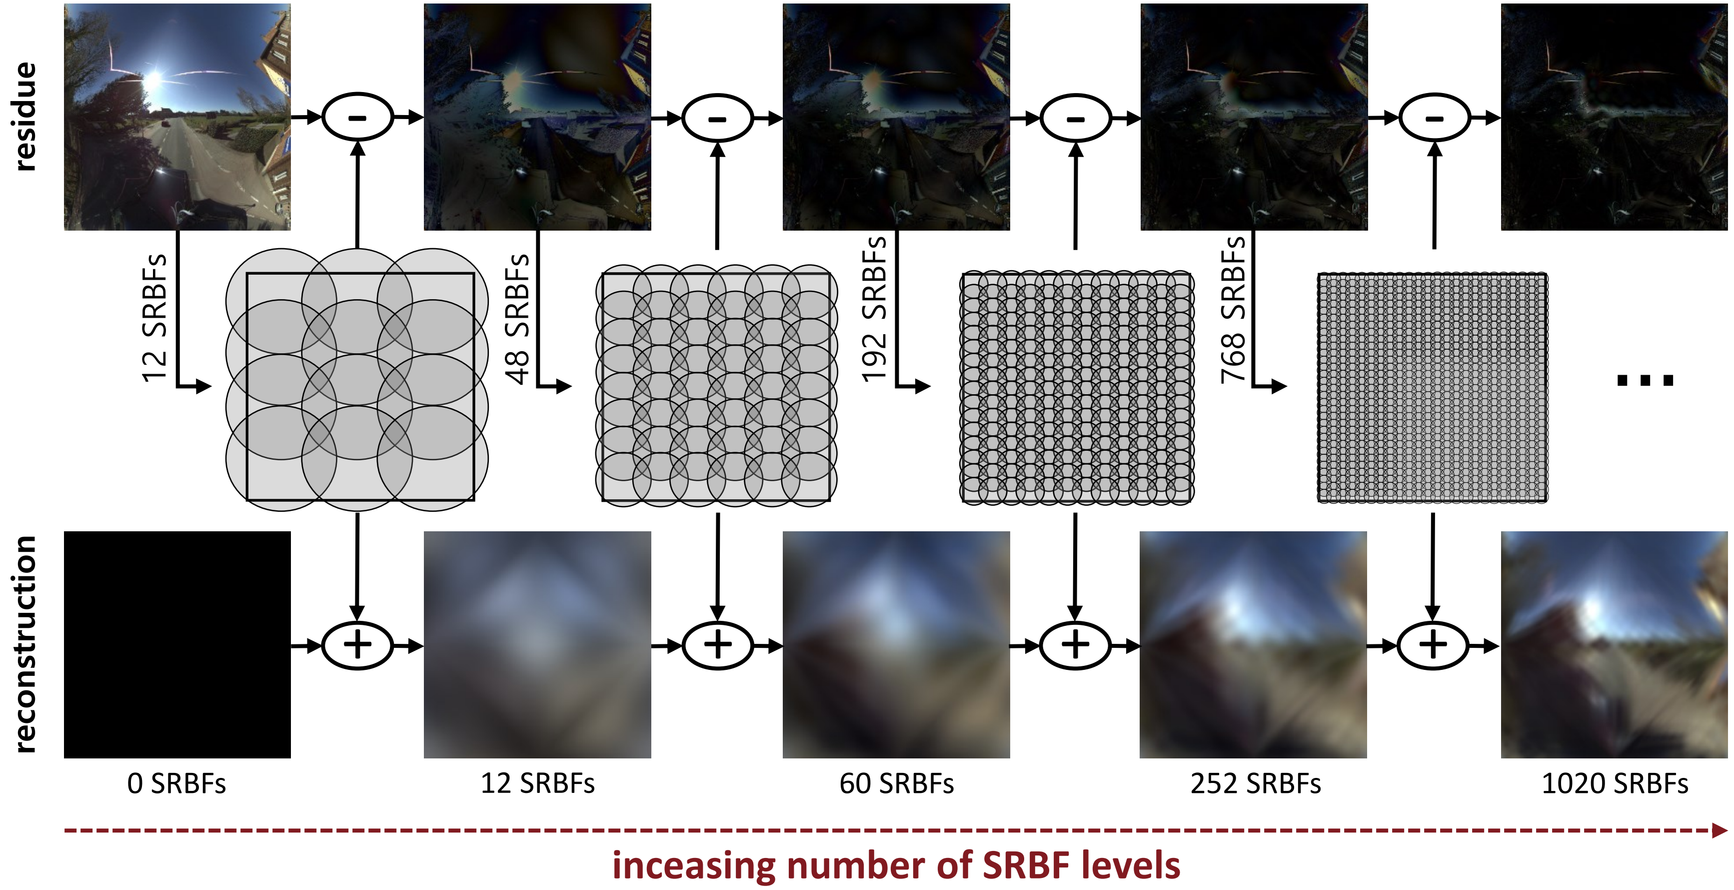

Overlapping Lighting SRBFs

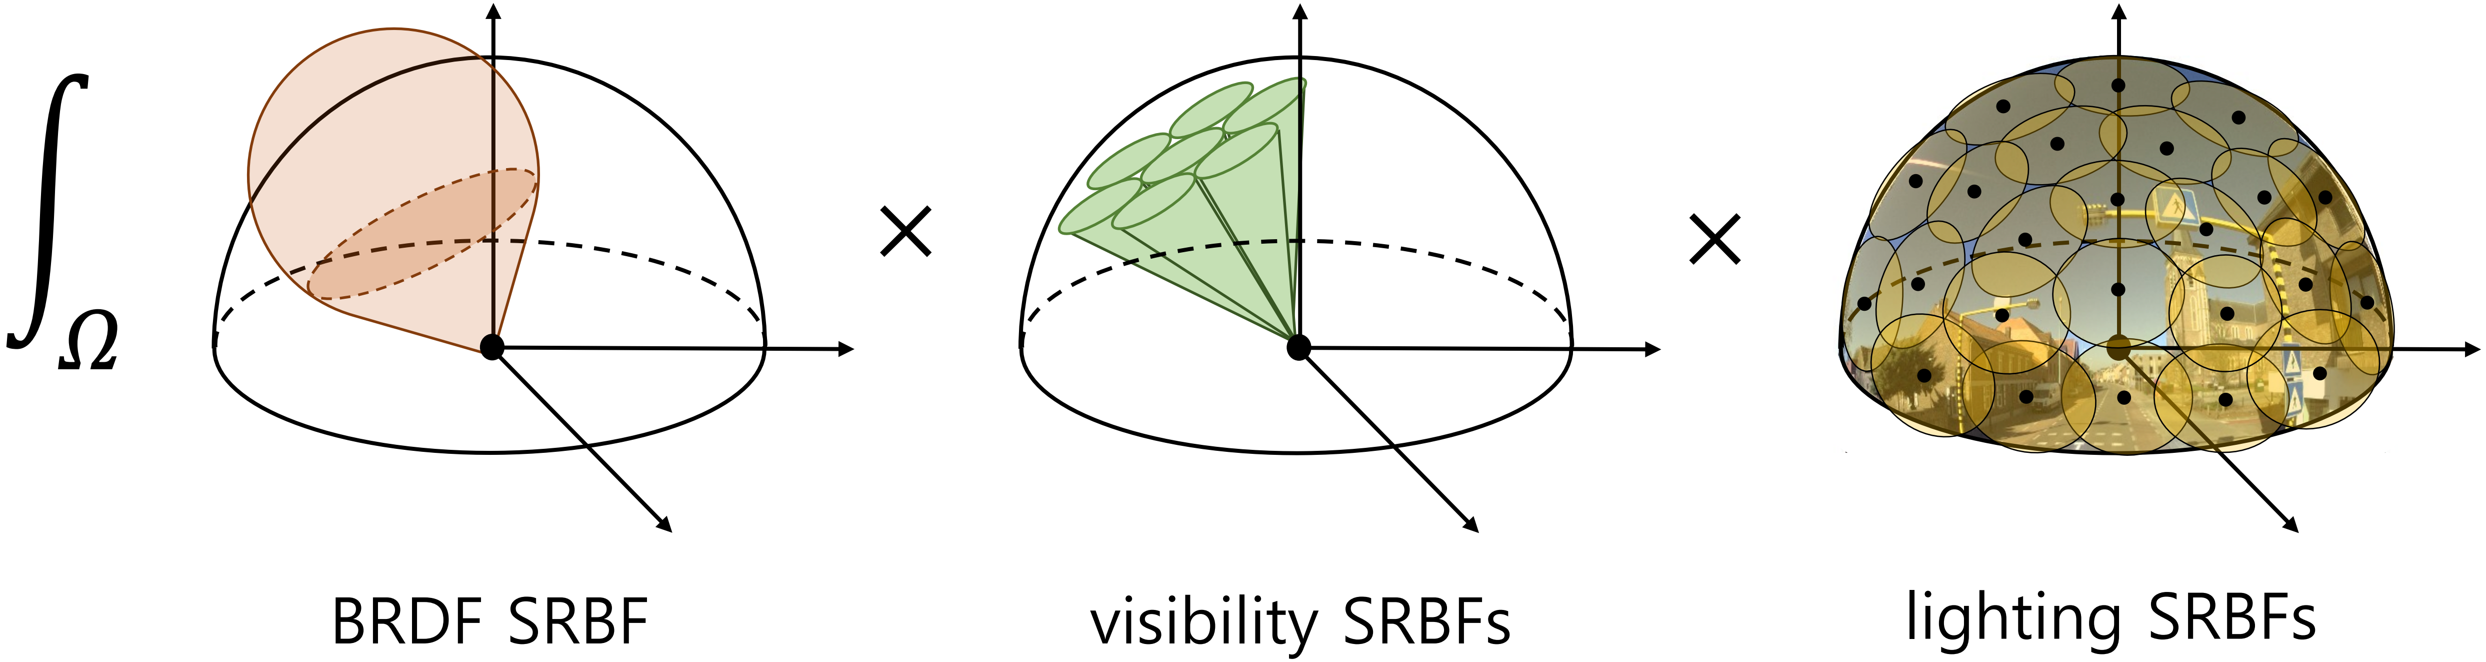

Overlapping Lighting SRBFs

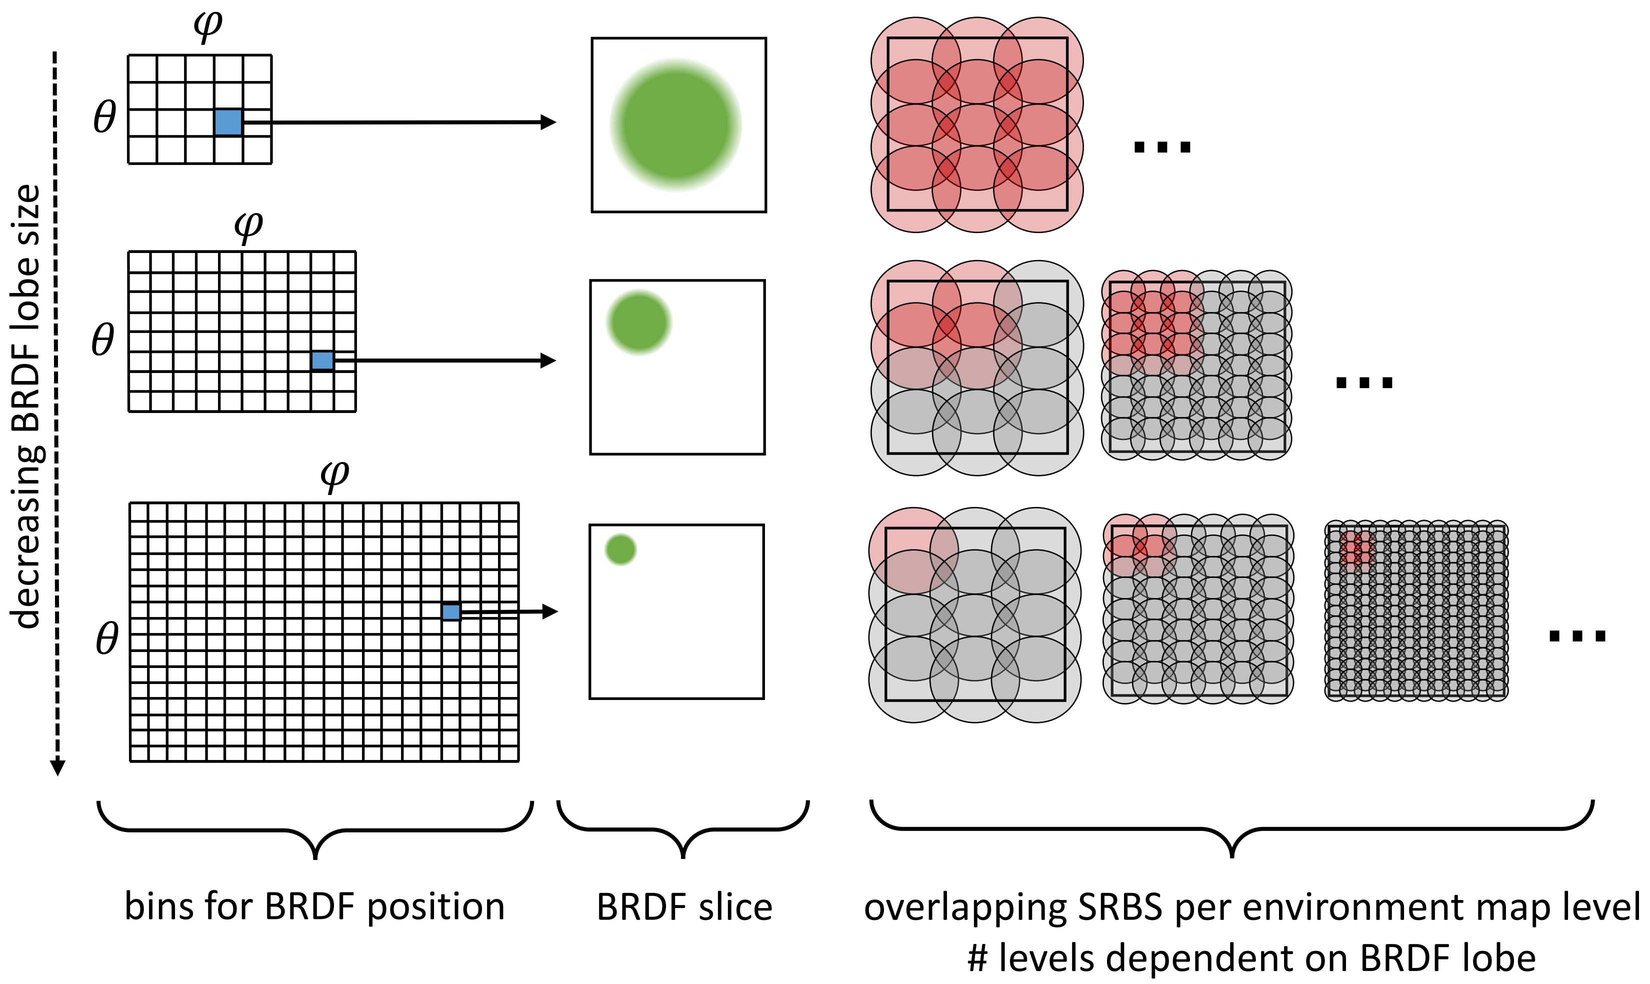

Peak Detection

- Problem: bright area light sources
  - Requires too fine-grained subsampling of visibility cones
+ Solution: peak detection
  - Treated as a special case
  - Thresholding / connected components / fitting


environment map


connected components

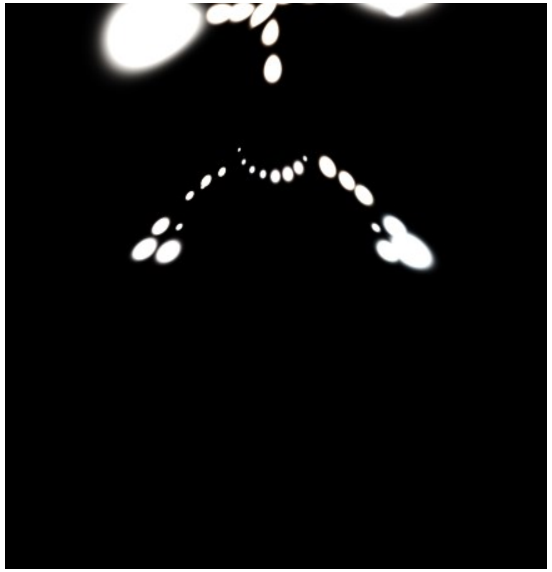
SRBF fitting

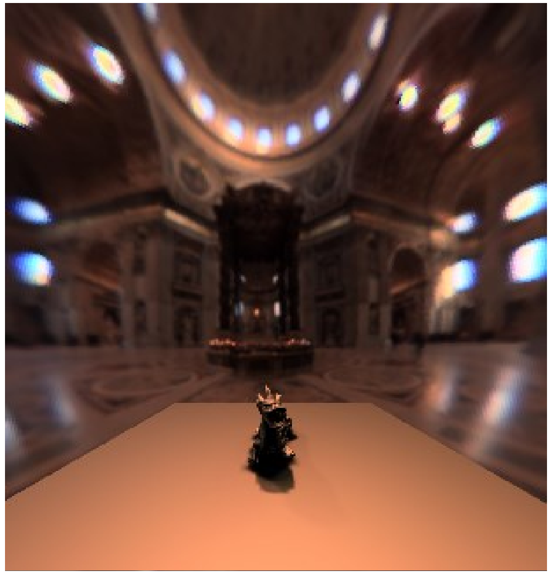
rendering

Results

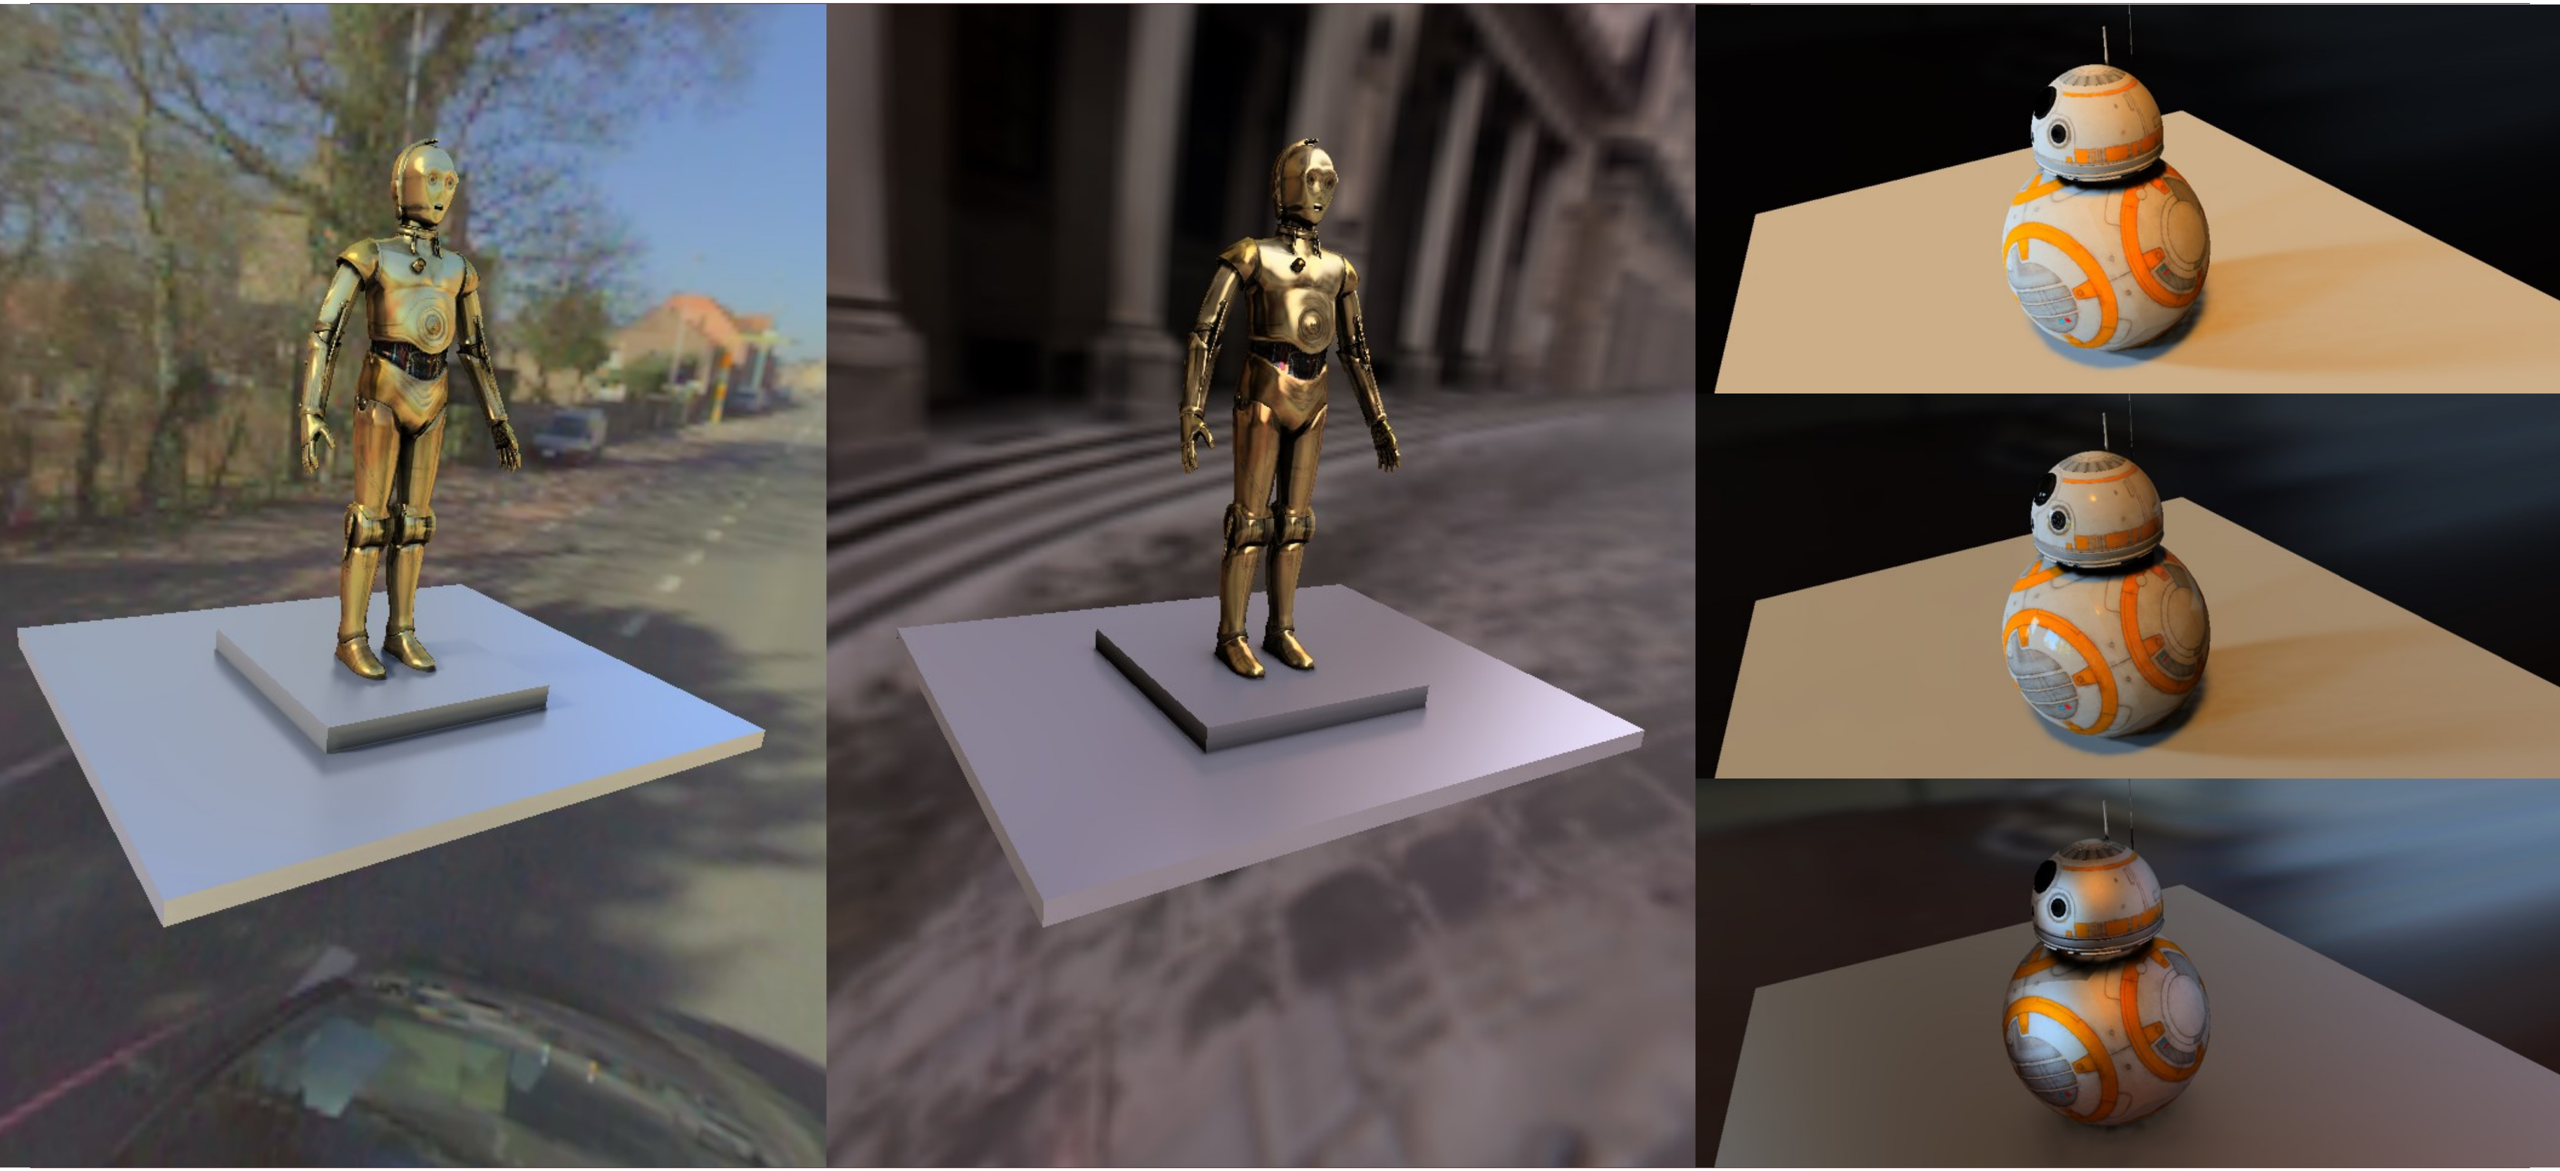

Results

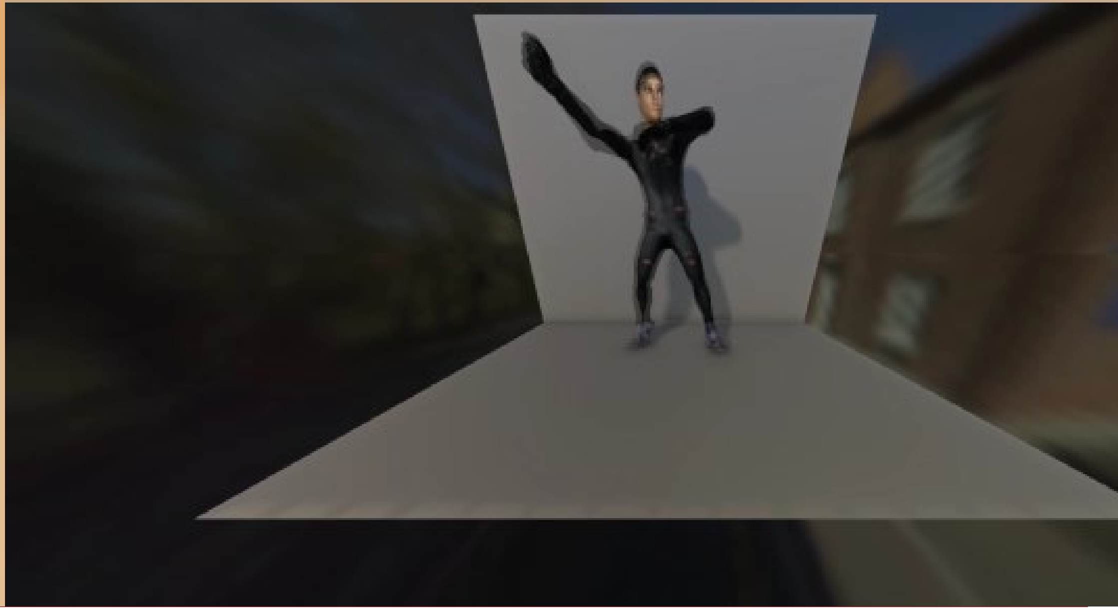

Results

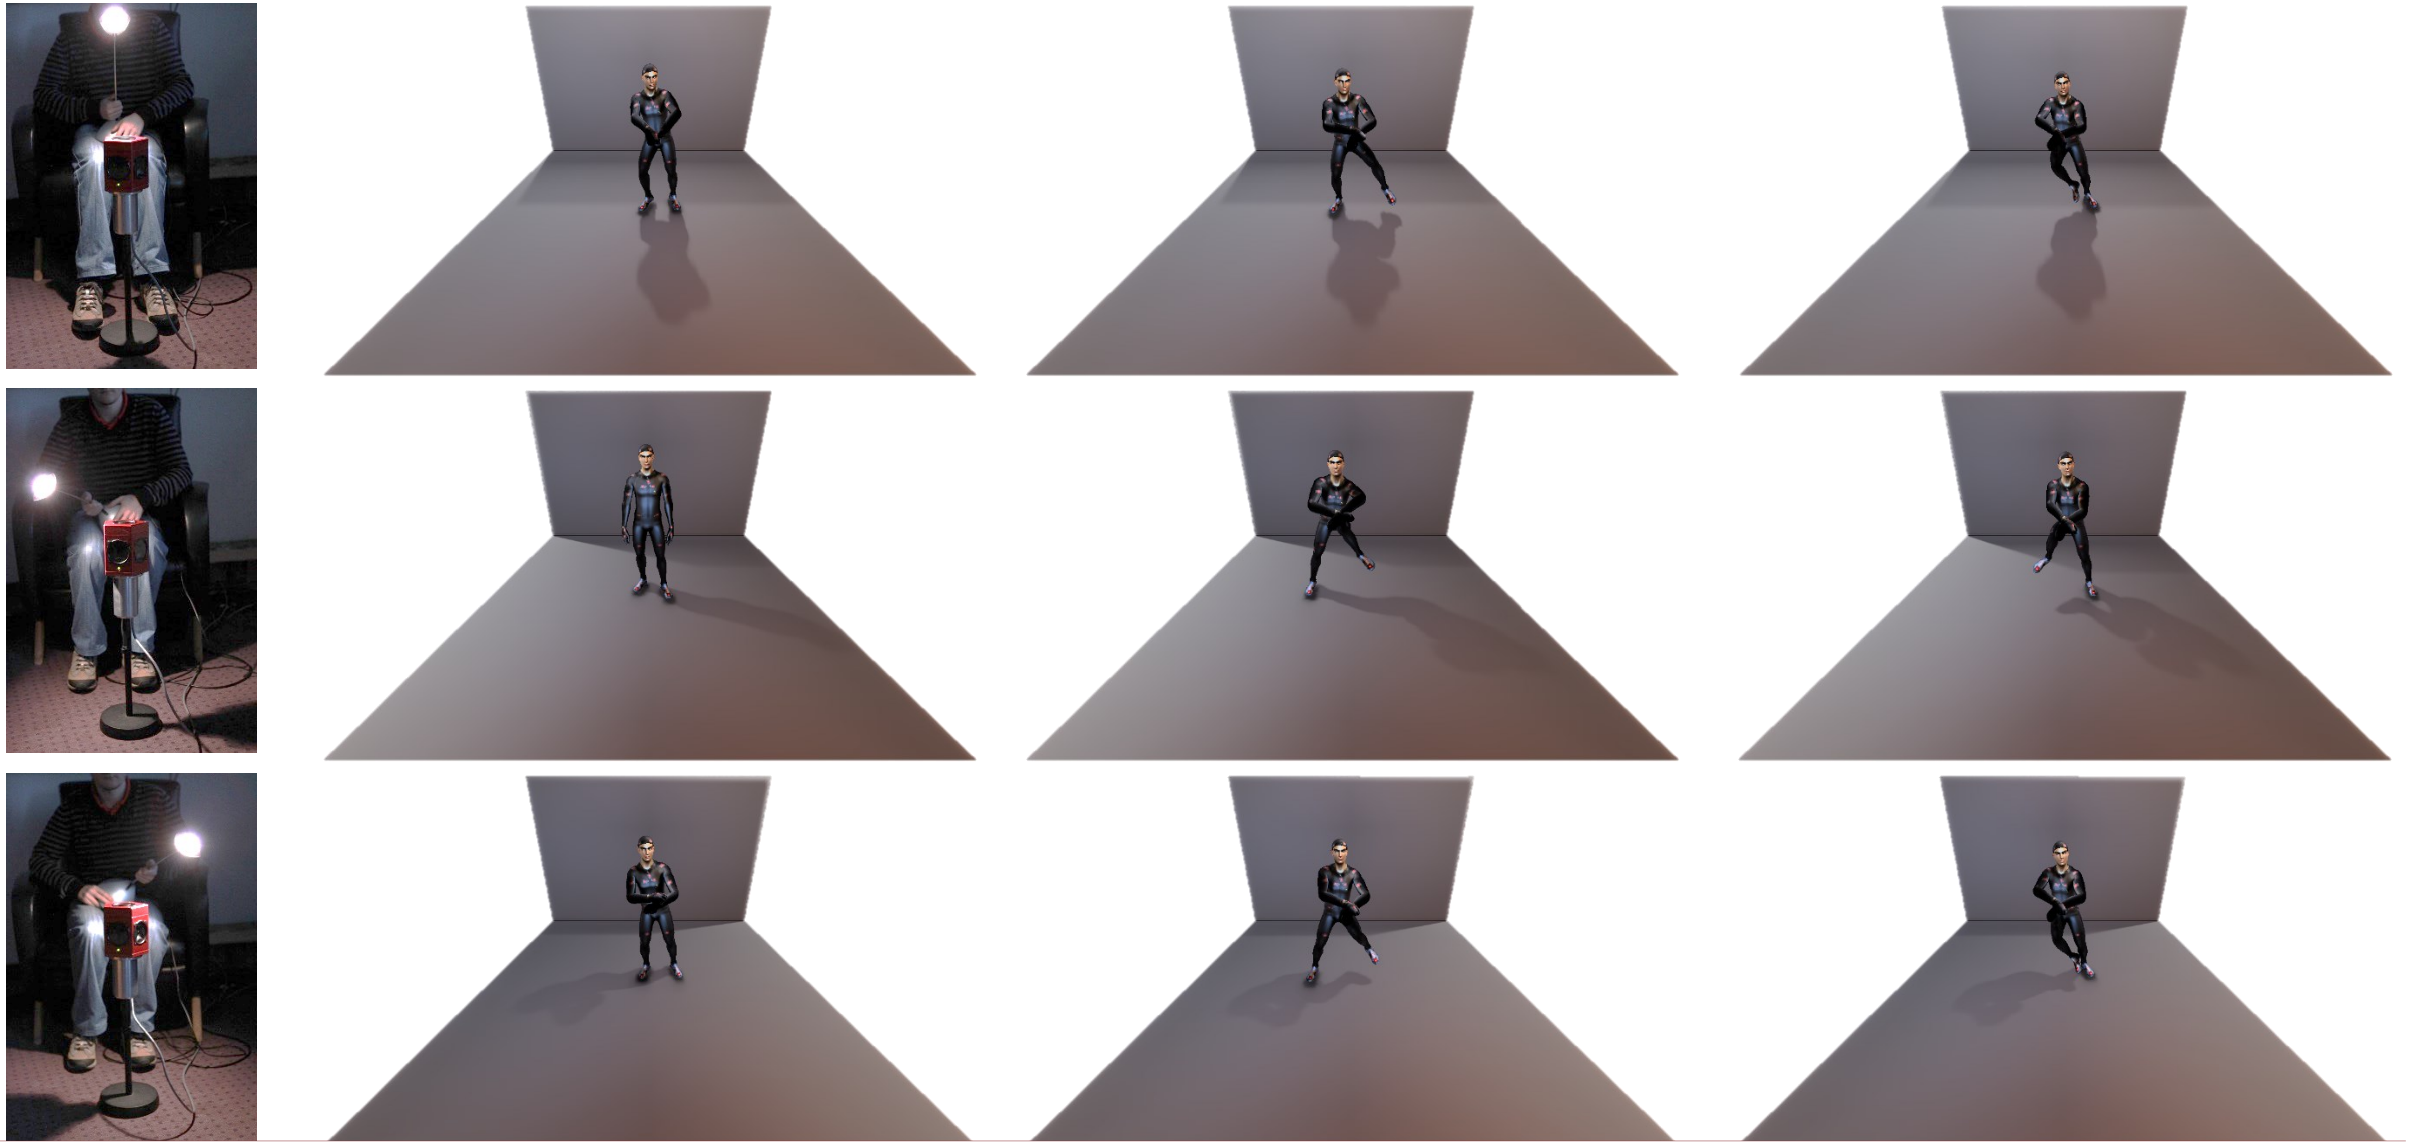

Relighting Use Cases

1. Relighting of **Virtual Objects**

2. Relighting of **Real Objects**

universiteit hasselt | EDM | iMinds

Relighting of **Virtual Objects**

Relighting of **Virtual Objects**

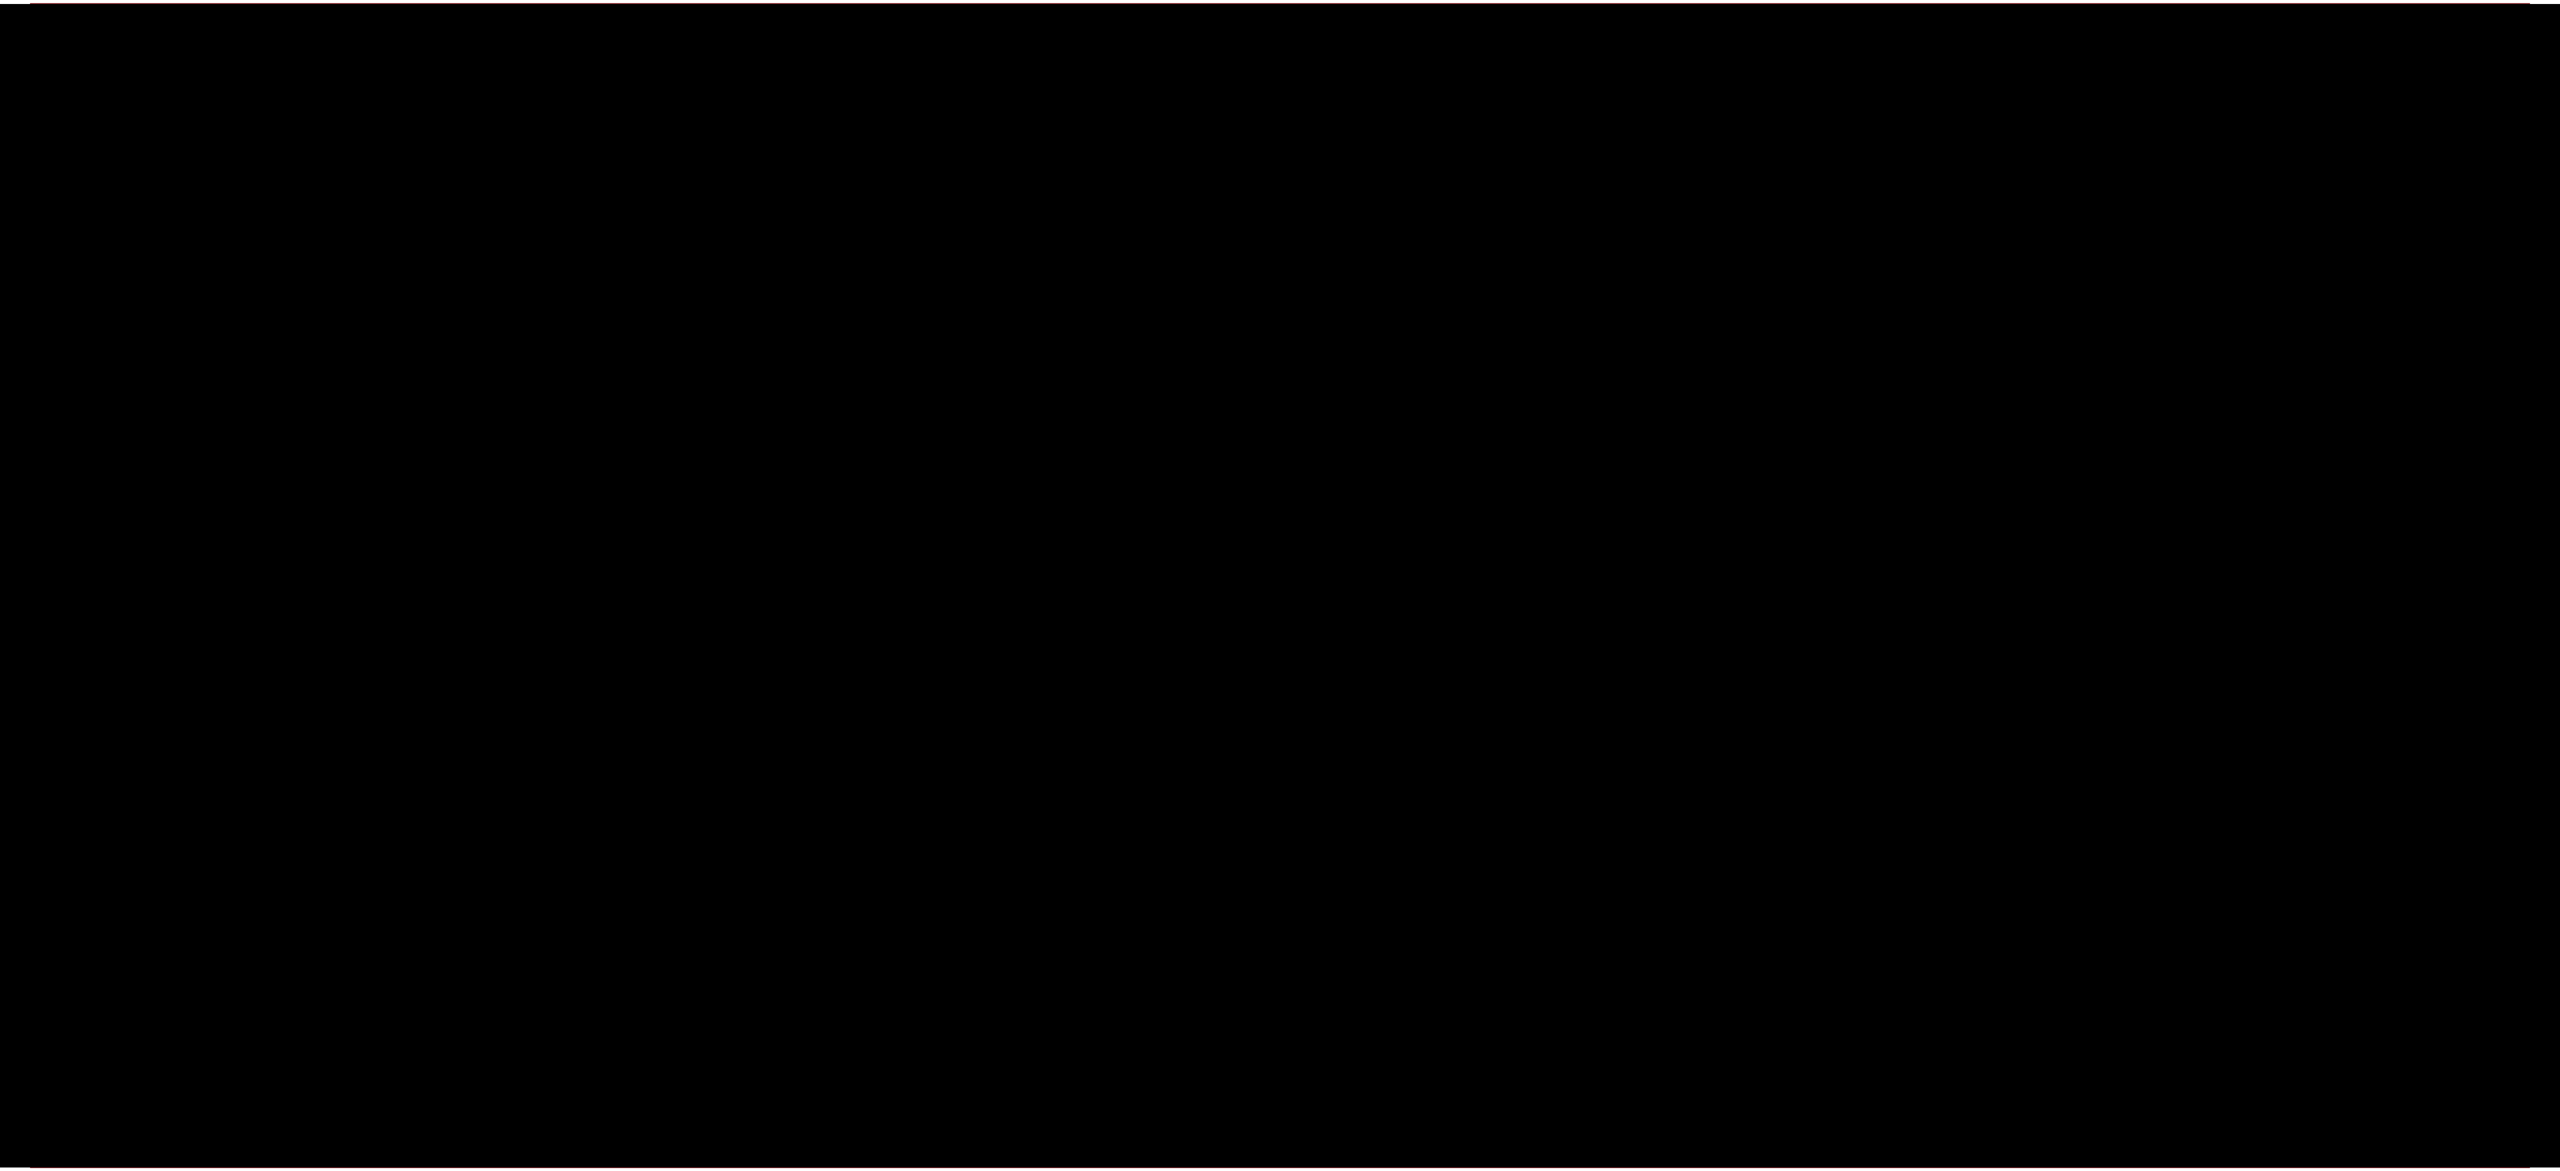

Relighting Use Cases

1. Relighting of **Virtual Objects**

2. Relighting of **Real Objects**

Relighting of **Real Objects**

- Inverse rendering using wavelets *[Haber et al., 2009]*
  - Hierarchical refinement using smooth high-order wavelets

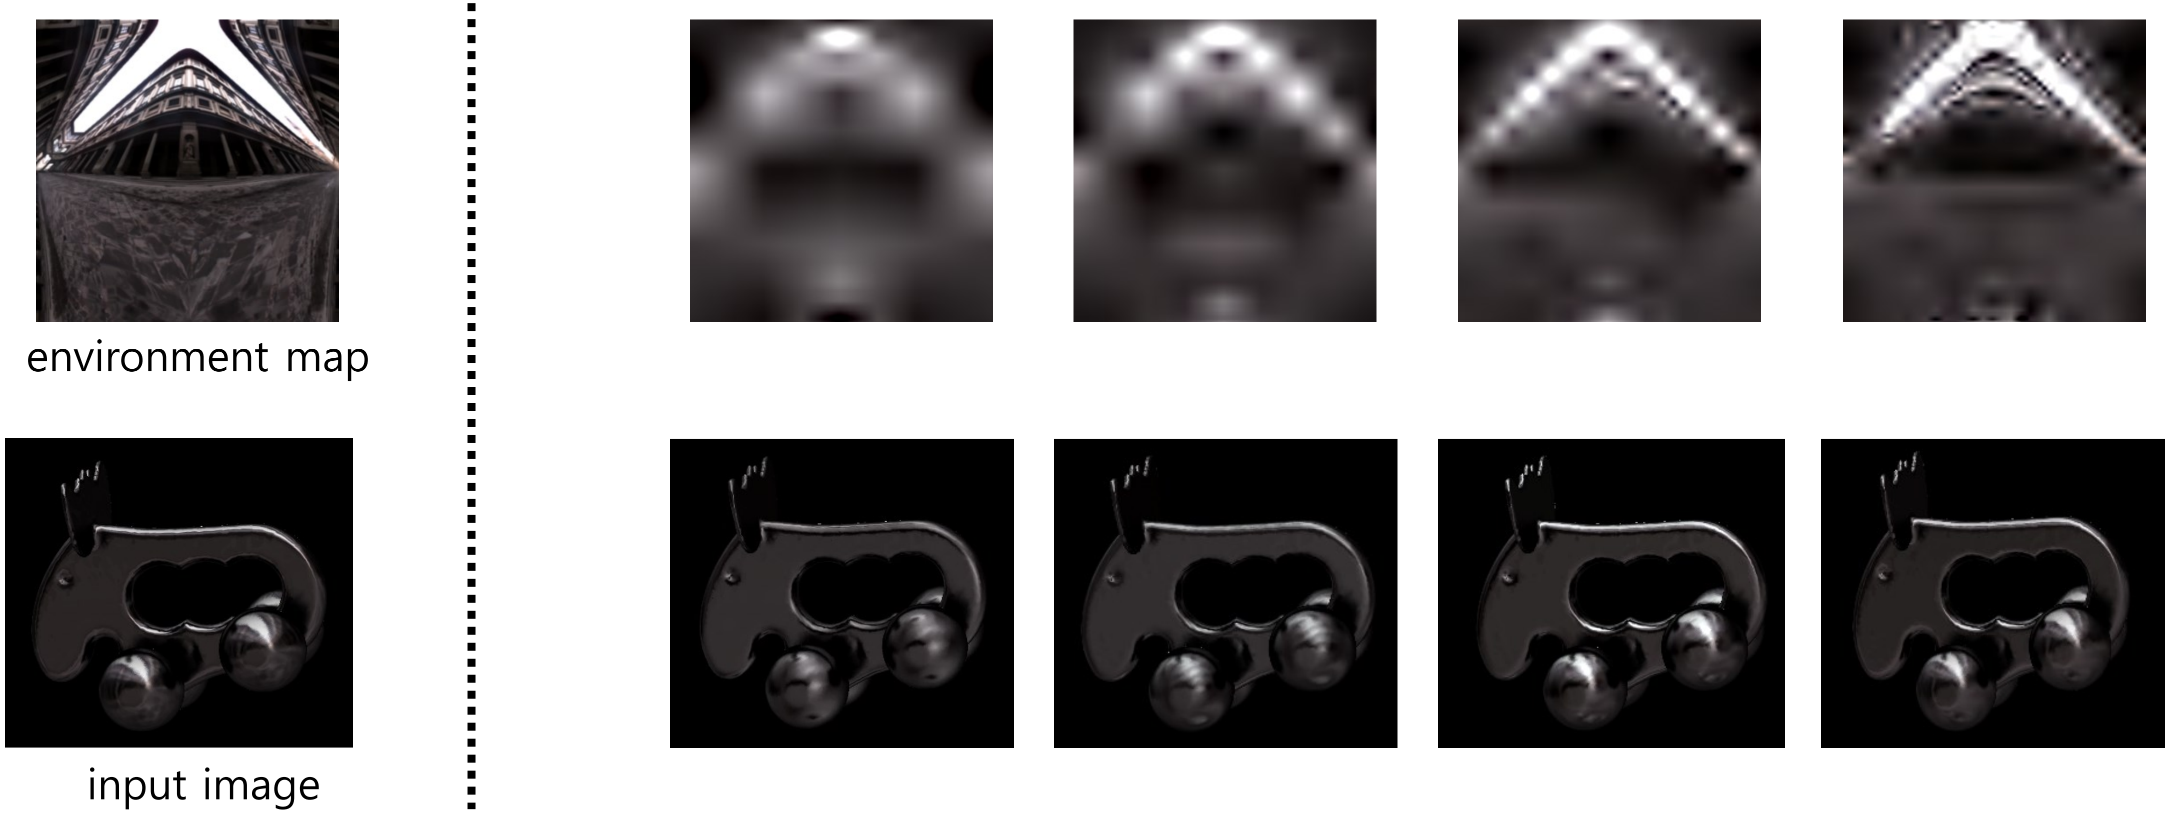

Relighting of **Real Objects**

- Inverse rendering using wavelets *[Haber et al., 2009]*
  - Hierarchical refinement using high-order wavelets
  - Temporal information

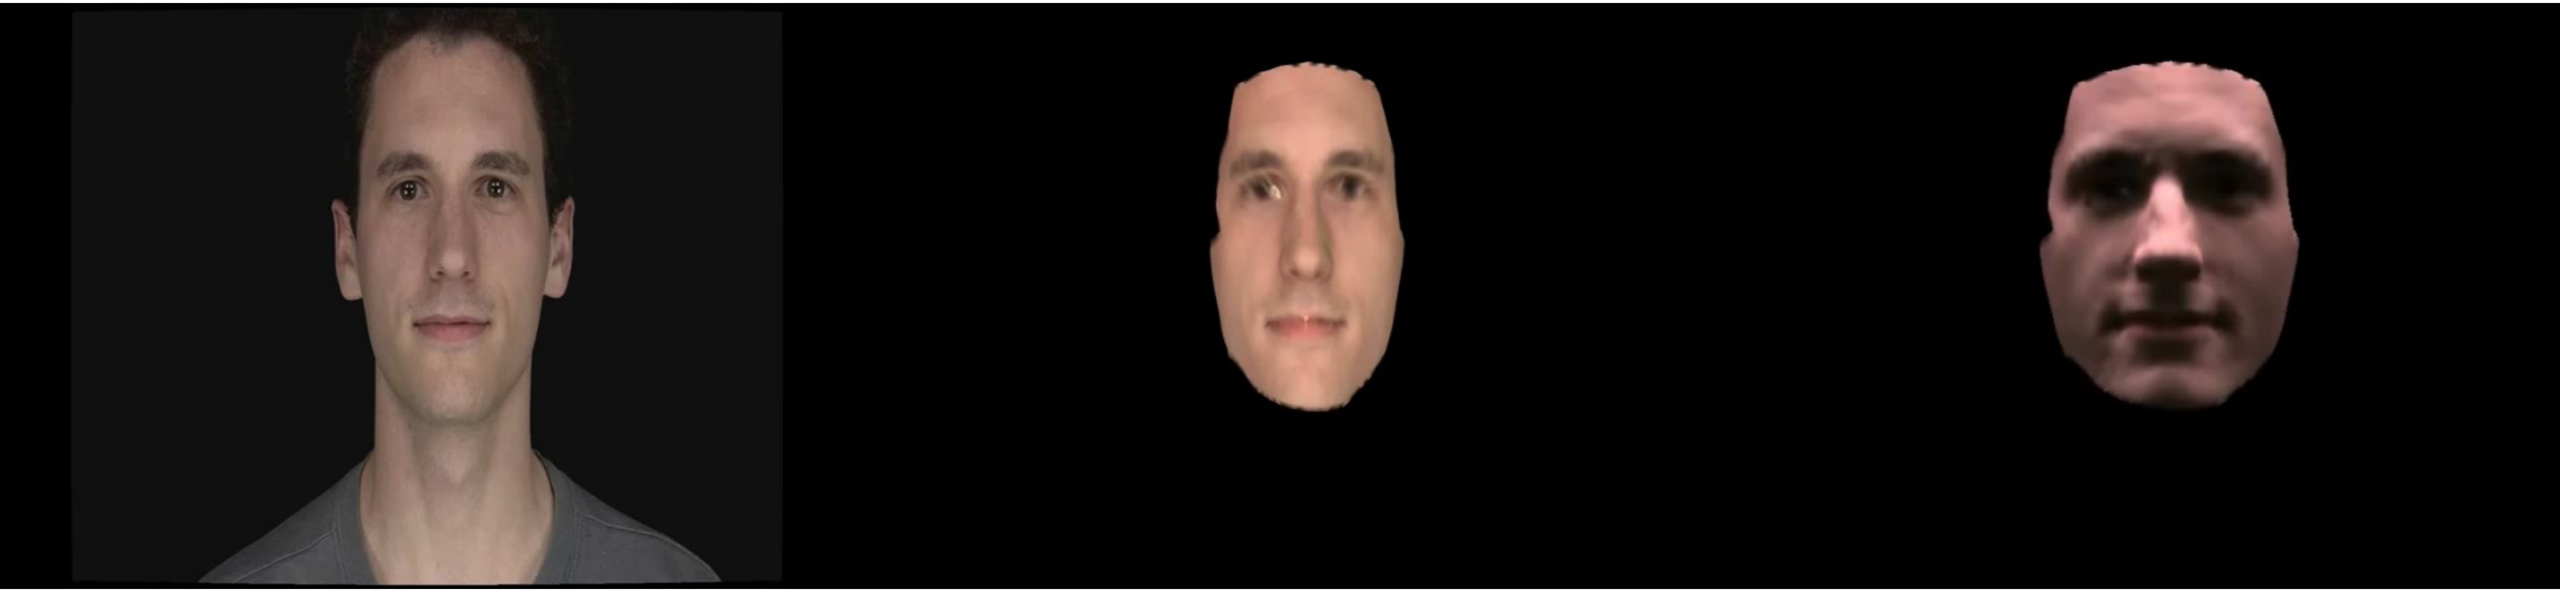

input sequence

reconstructed sequence

relit with Uffizi environment map

Relighting of **Real Objects**

- Inverse rendering using wavelets *[Haber et al., 2009]*
  - Hierarchical refinement using high-order wavelets
  - Temporal information
  - Near-field lighting

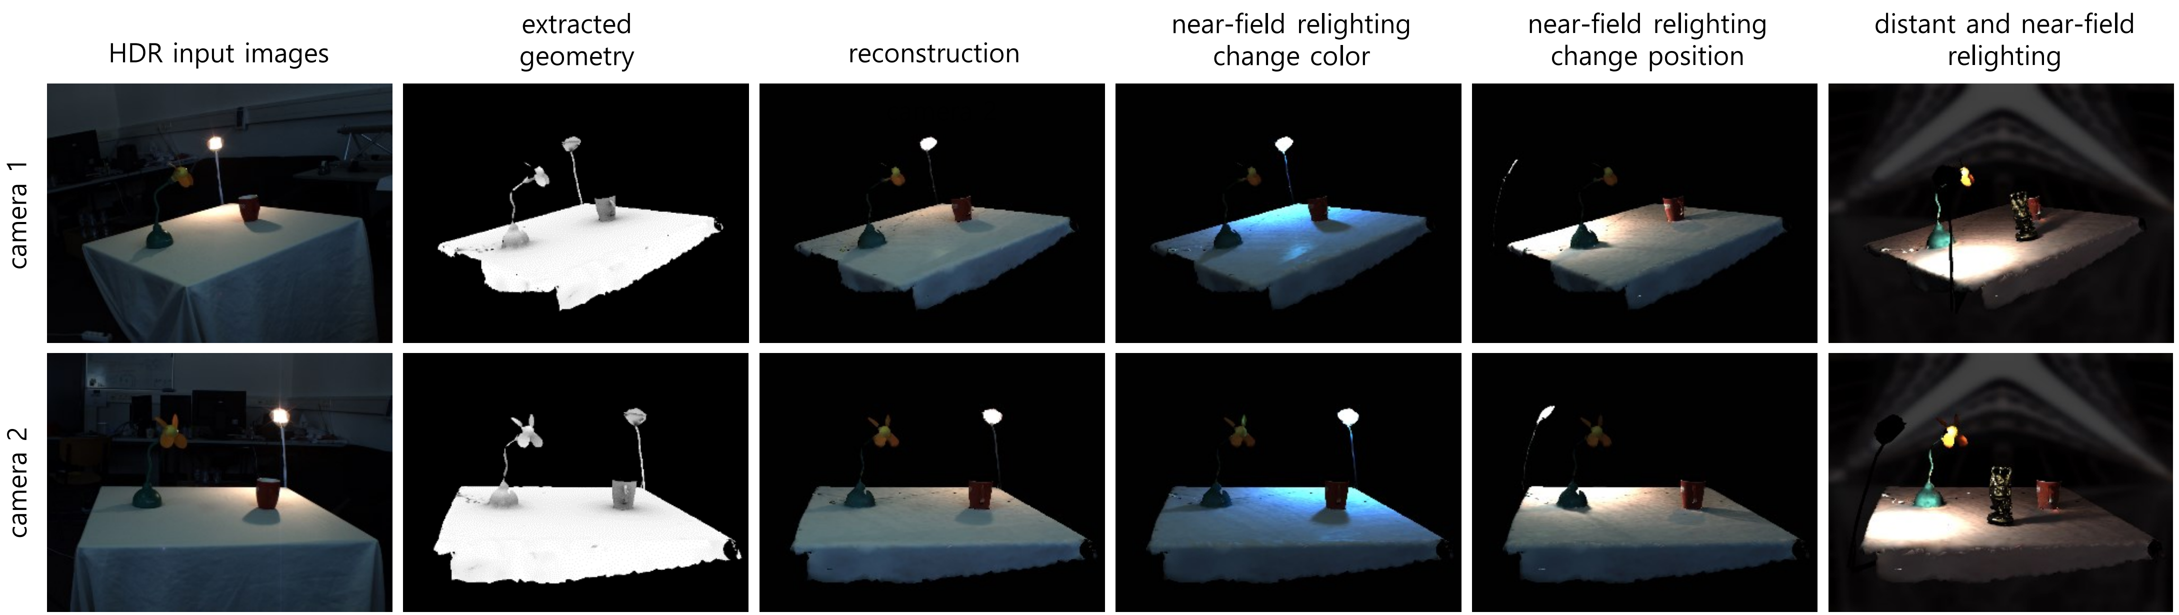

universiteit hasselt | EDM | iMinds

Relighting of **Real Objects**

- Inverse rendering using wavelets *[Haber et al., 2009]*
  - Hierarchical refinement using high-order wavelets
  - Temporal information
  - Near-field lighting
- Inverse rendering using SRBFs *[Haber et al., 2009]*

Relighting of Real Objects - **Inverse Rendering using SRBFs**

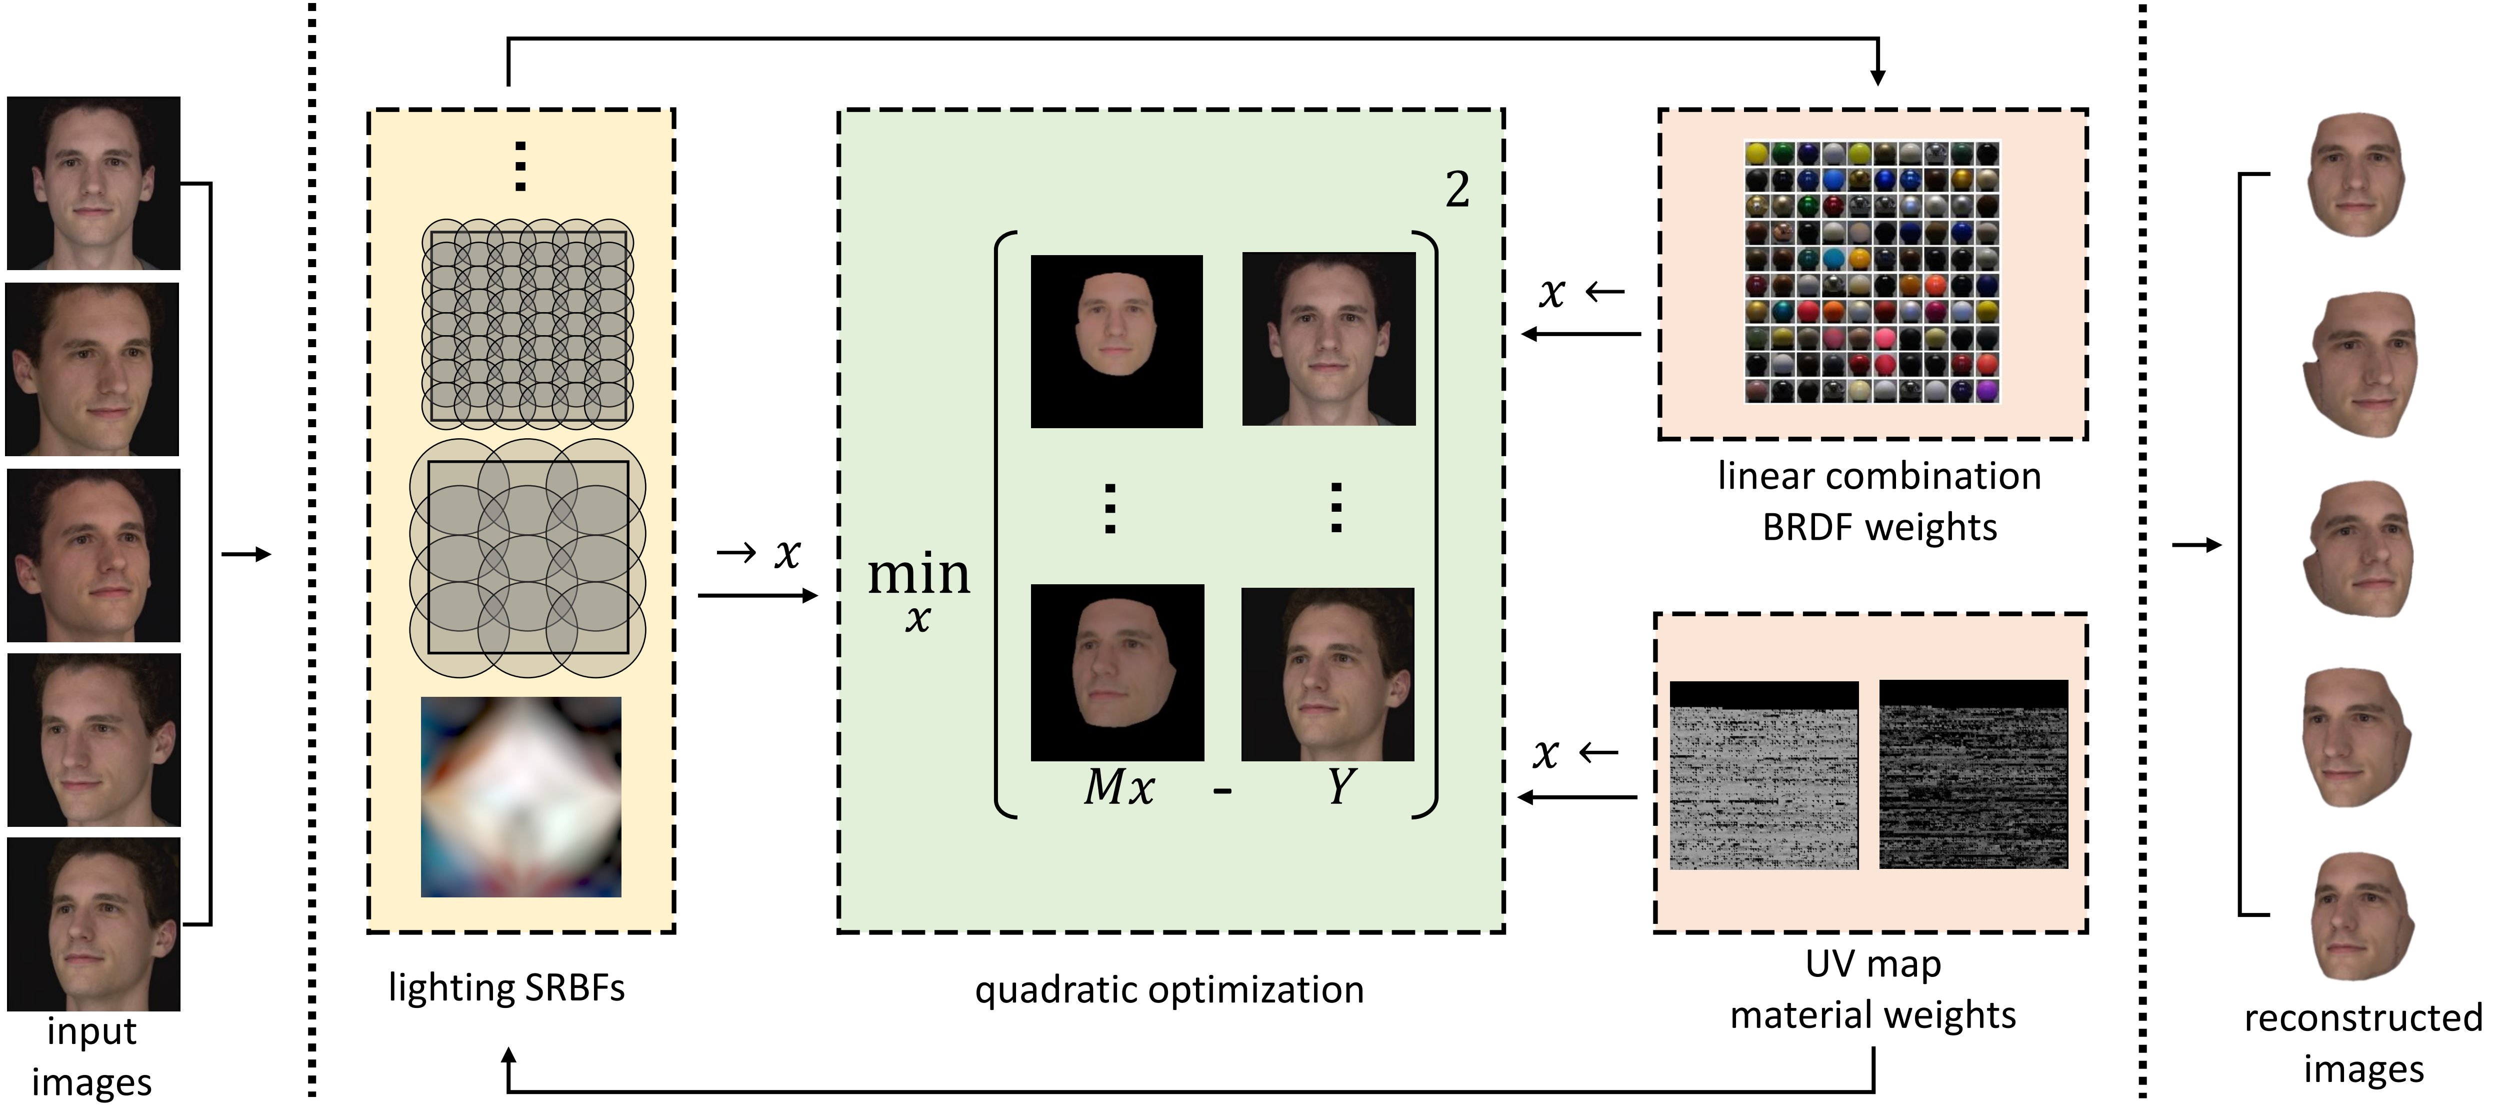

universiteit hasselt | EDM | iMinds

Relighting of Real Objects - **Inverse Rendering using SRBFs**

	Haar wavelets [Haber et al., 2009]	SRBF [ours]
# lighting coefficients	1024	**1020**
Δt optimize lighting	1089 s	**10 s**
# BRDF weights	2689 (per vertex)	**906240 (per texel)**
Δt optimize BRDF weights	3 s	**180 s**
Δt optimize one BRDF weights	9.4×10^{-4} s	$\mathbf{1.9 \times 10^{-4}}$ **s**
Δt full optimization	184 min	**29 min**
Δt rendering reconstruction	13.403 s	$\mathbf{< 0.1}$ **s**

Conclusions

- General triple product theorem for high-order wavelets
  - \+ Triple product rendering for a mixture of different wavelets bases
  - \+ Wavelets tailored to the signal
  - \+ Few coefficients to estimate
  - \- Rather slow

- Dynamic triple product rendering using SRBFs
  - \+ Interactive and real-time triple product rendering
  - \+ All three factors are dynamic
  - \+ Preview rendering of estimates
  - \- Slightly more coefficients: quality/speed tradeoff

Take Home Messages

- Underlying representation does have an impact on relighting applications
  - Representation tailored to the signal
  - Representation tailored to the application
  - Compression/time tradeoff

Thank you for your attention!

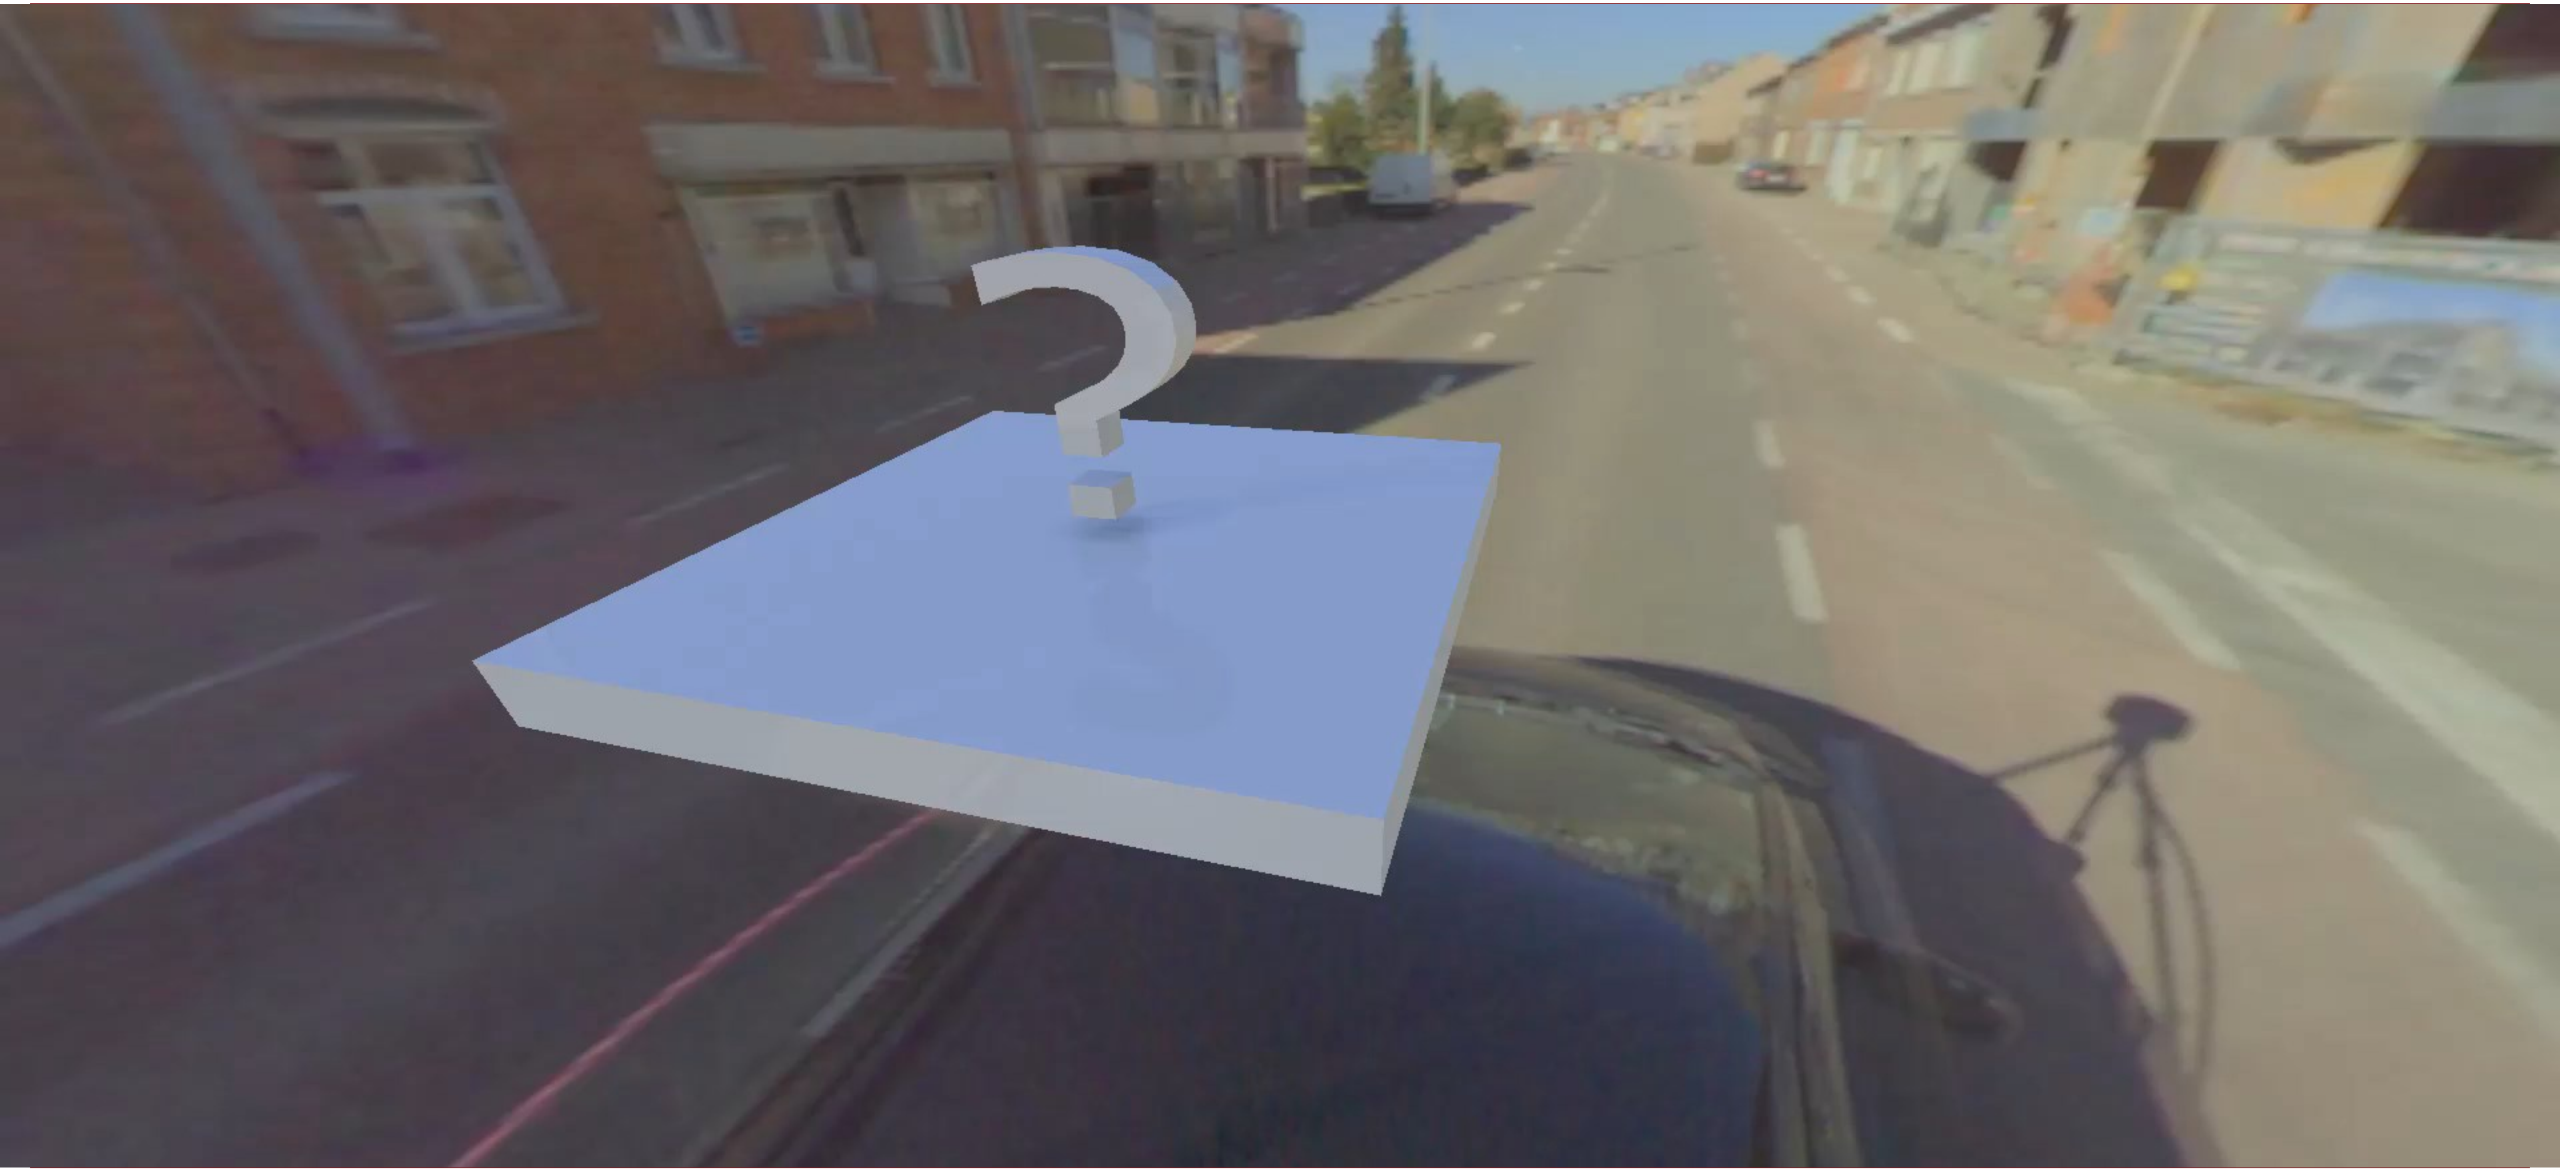

Thank you for your attention

High-Order Wavelets in Render Application

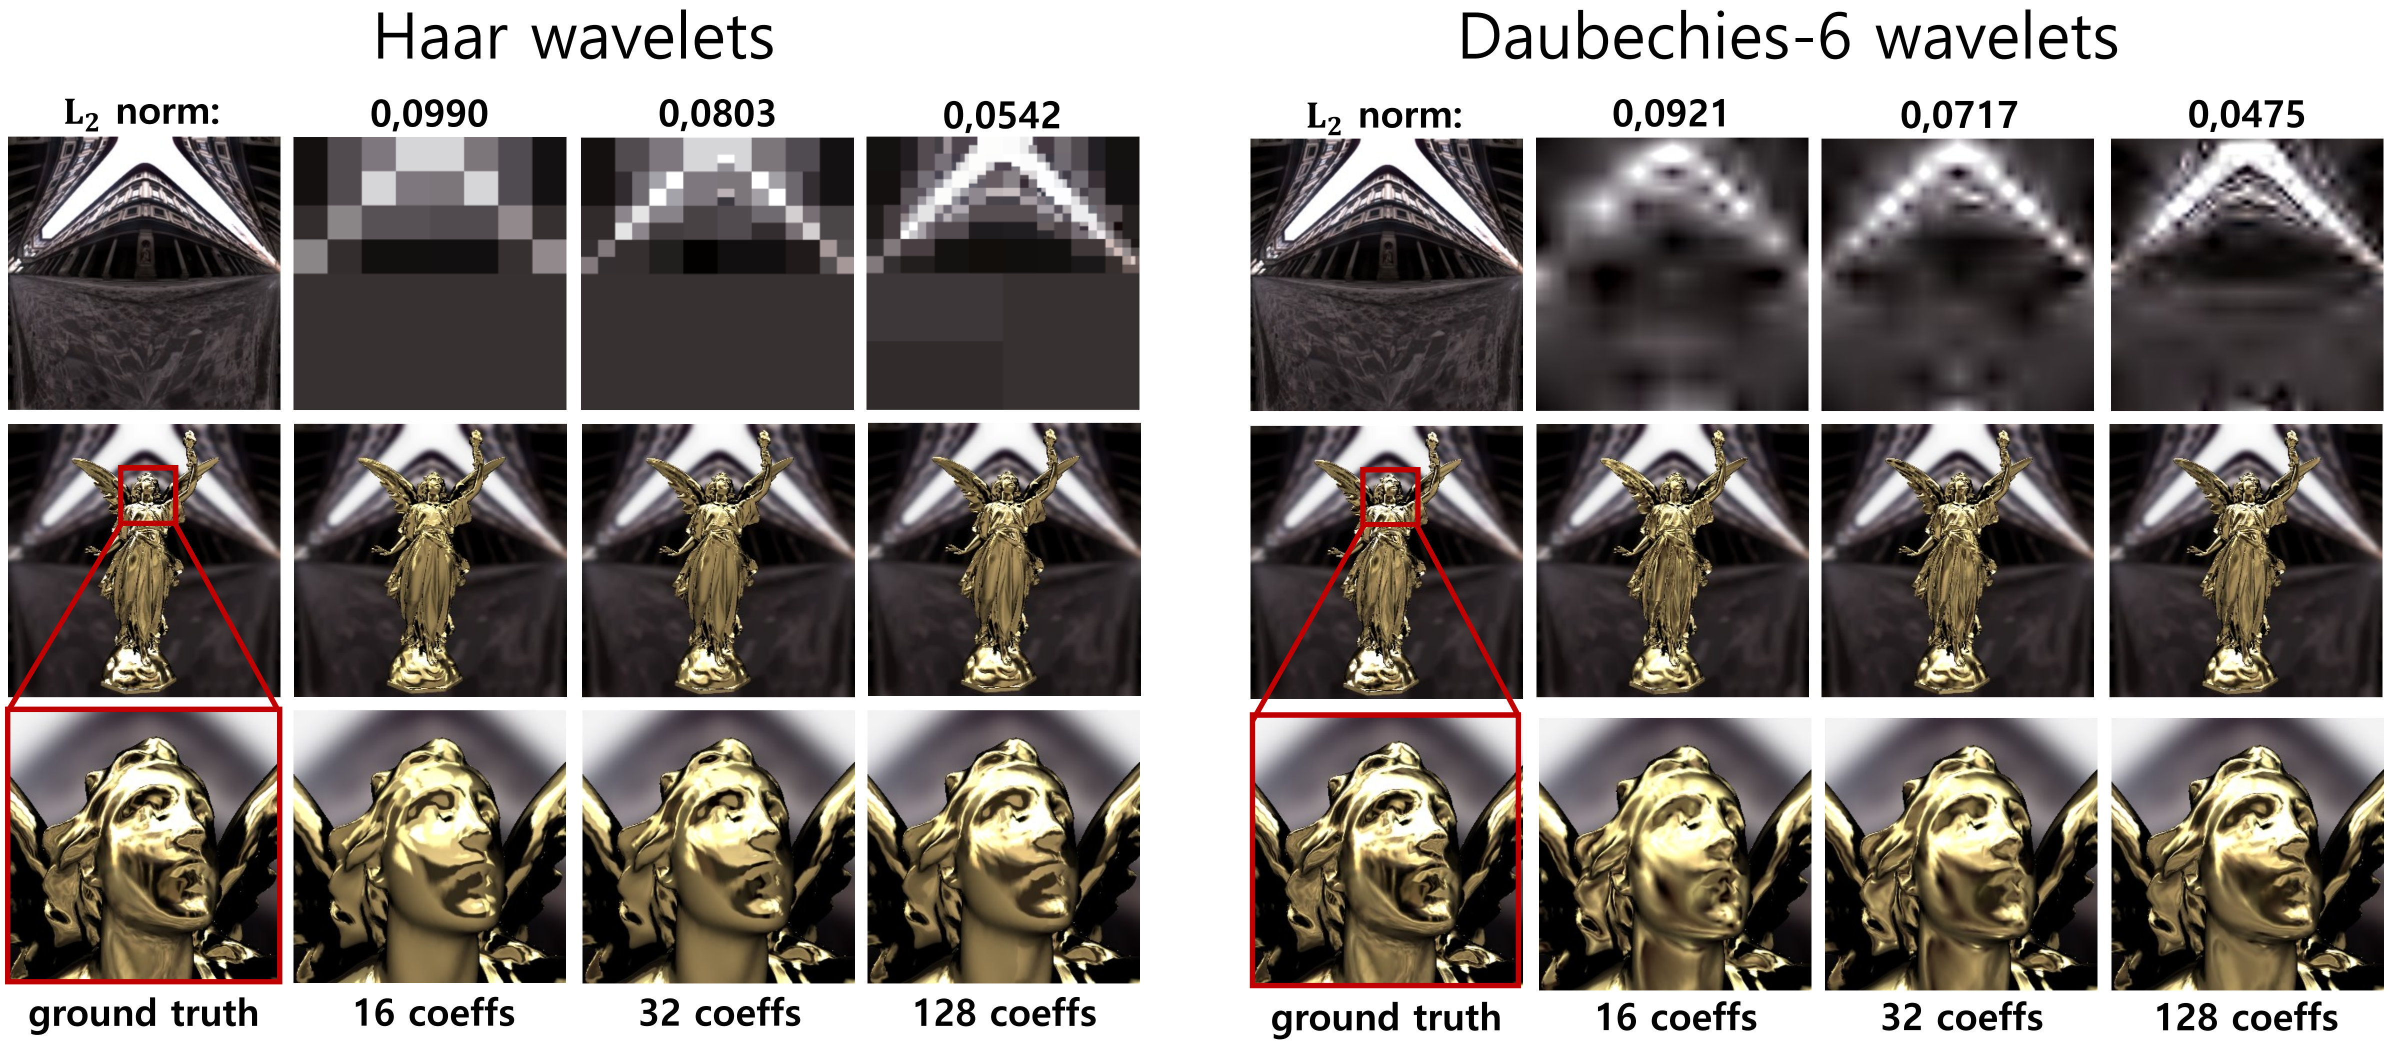

High-Order Wavelets in Render Application

universiteit hasselt | EDM | iMinds

High-Order Wavelets in Render Application

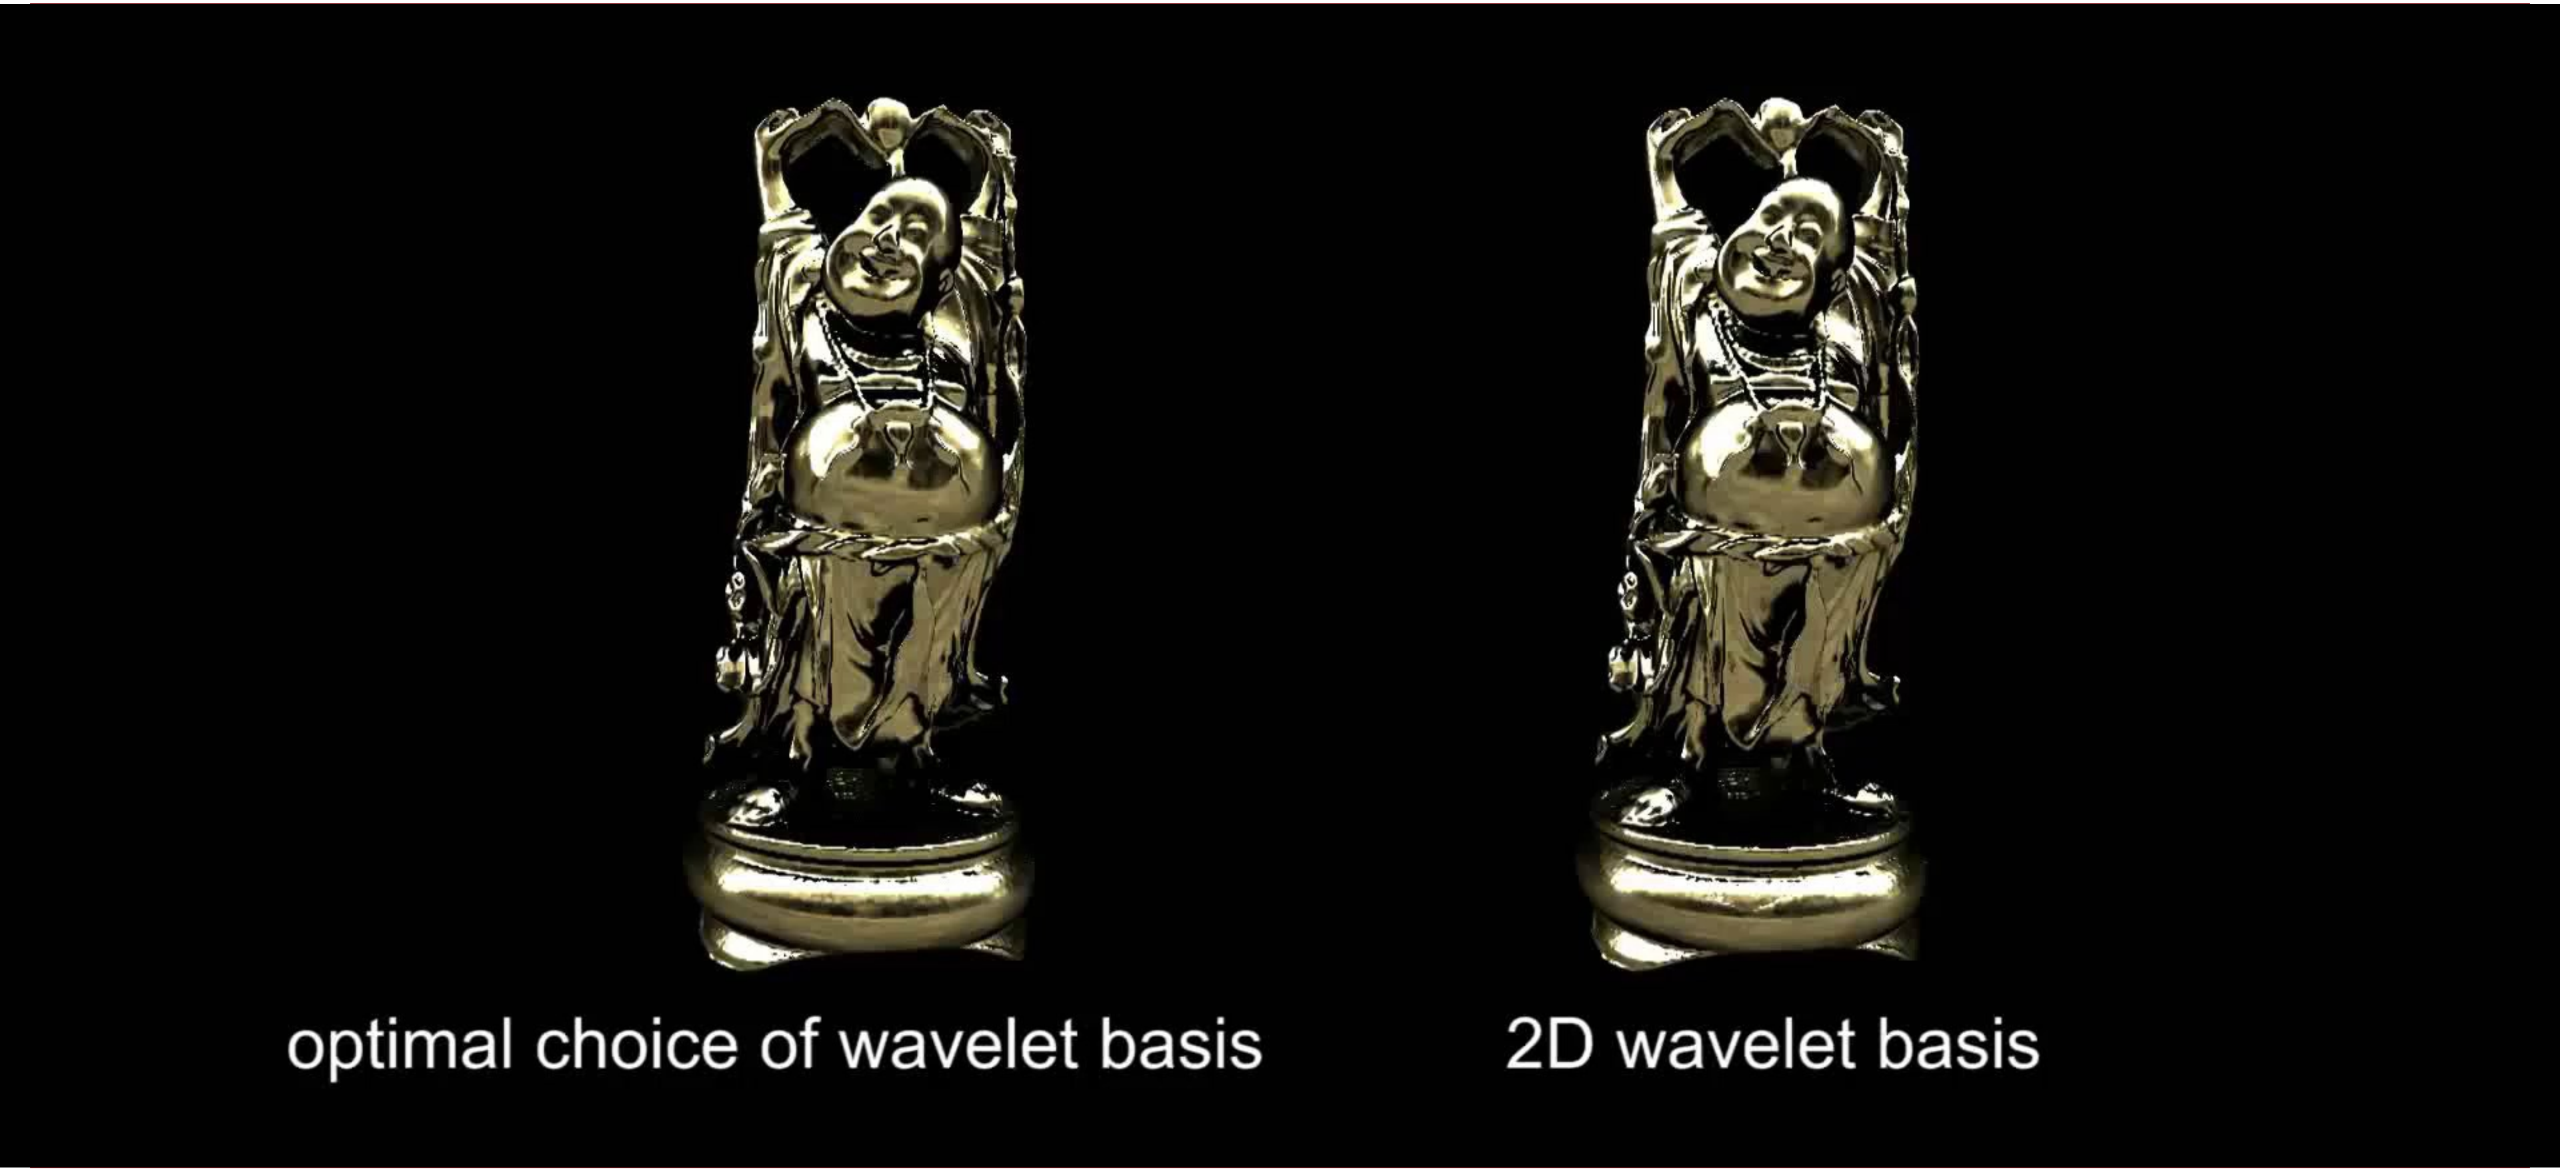

Triple product rendering with SRBFs

+ Analytic Evaluation of the binding coefficients

$$C_{ijk} = \int_{\Omega} \Psi_i(\omega)\Psi_j(\omega)\Psi_k(\omega)d\omega$$

$$\boldsymbol{\Psi_{i,j,k} = G(\omega, c, \lambda, \mu) = \mu e^{\lambda(\omega \cdot c - 1)}}$$

$$= \int_{\Omega} G(\omega, \boldsymbol{c}_V, \lambda_V, \mu_V) G(\omega, \boldsymbol{c}_L, \lambda_L, \mu_L) G(\omega, \boldsymbol{c}_\rho, \lambda_\rho, \mu_\rho) d\omega$$

$$= \mu_V \mu_L \mu_\rho e^{-(\lambda_V + \lambda_L + \lambda_\rho)} \int_{\Omega} e^{\omega \cdot (\lambda_V \boldsymbol{c}_V + \lambda_L \boldsymbol{c}_L + \lambda_\rho \boldsymbol{c}_\rho)} d\omega$$

$$= \mu_V \mu_L \mu_\rho e^{-(\lambda_V + \lambda_L + \lambda_\rho)} 4\pi \frac{\sinh \left\| \lambda_V \boldsymbol{c}_V + \lambda_L \boldsymbol{c}_L + \lambda_\rho \boldsymbol{c}_\rho \right\|}{\left\| \lambda_V \boldsymbol{c}_V + \lambda_L \boldsymbol{c}_L + \lambda_\rho \boldsymbol{c}_\rho \right\|}$$

[Tsai and Shih, 2006]

Rotation Operator

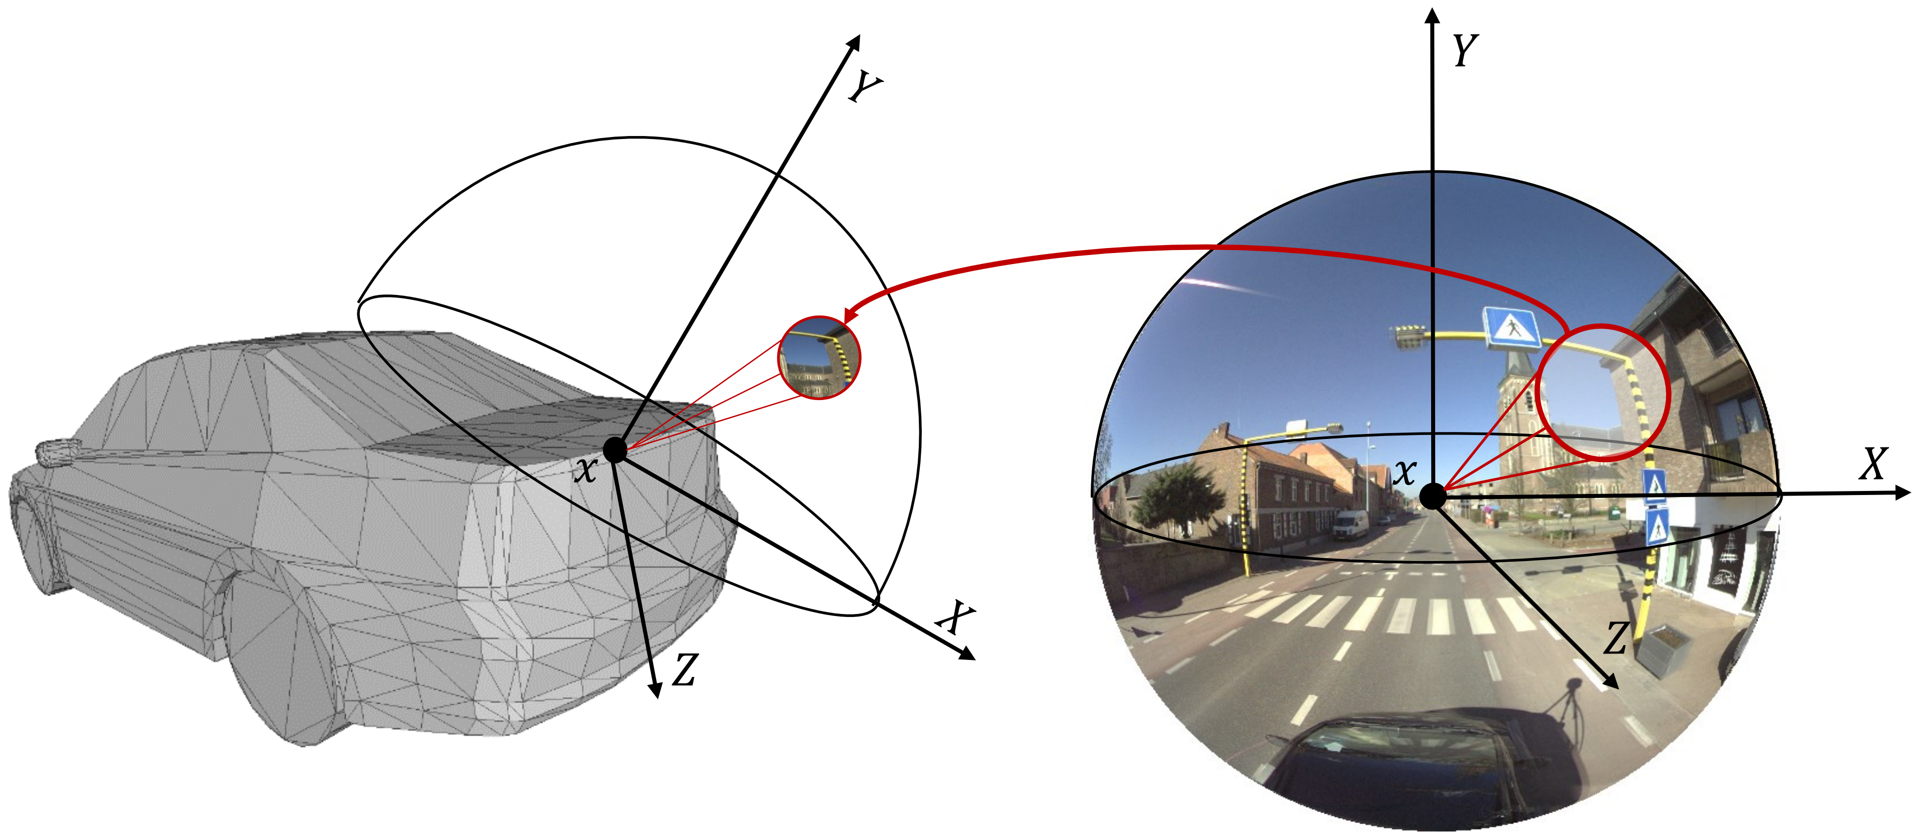

Overlapping Lighting SRBFs

bin level	BRDF λ	# bins	mean # overlapping SRBFs	max level
1	2,0	45	60	1
2	4,0	153	252	2
3	16,0	561	312	3
4	32,0	2145	608	4
5	128,0	2145	182	5
6	256,0	8320	109	5
7	512,0	33153	70	5

Dynamic Visibility – Voxel Density

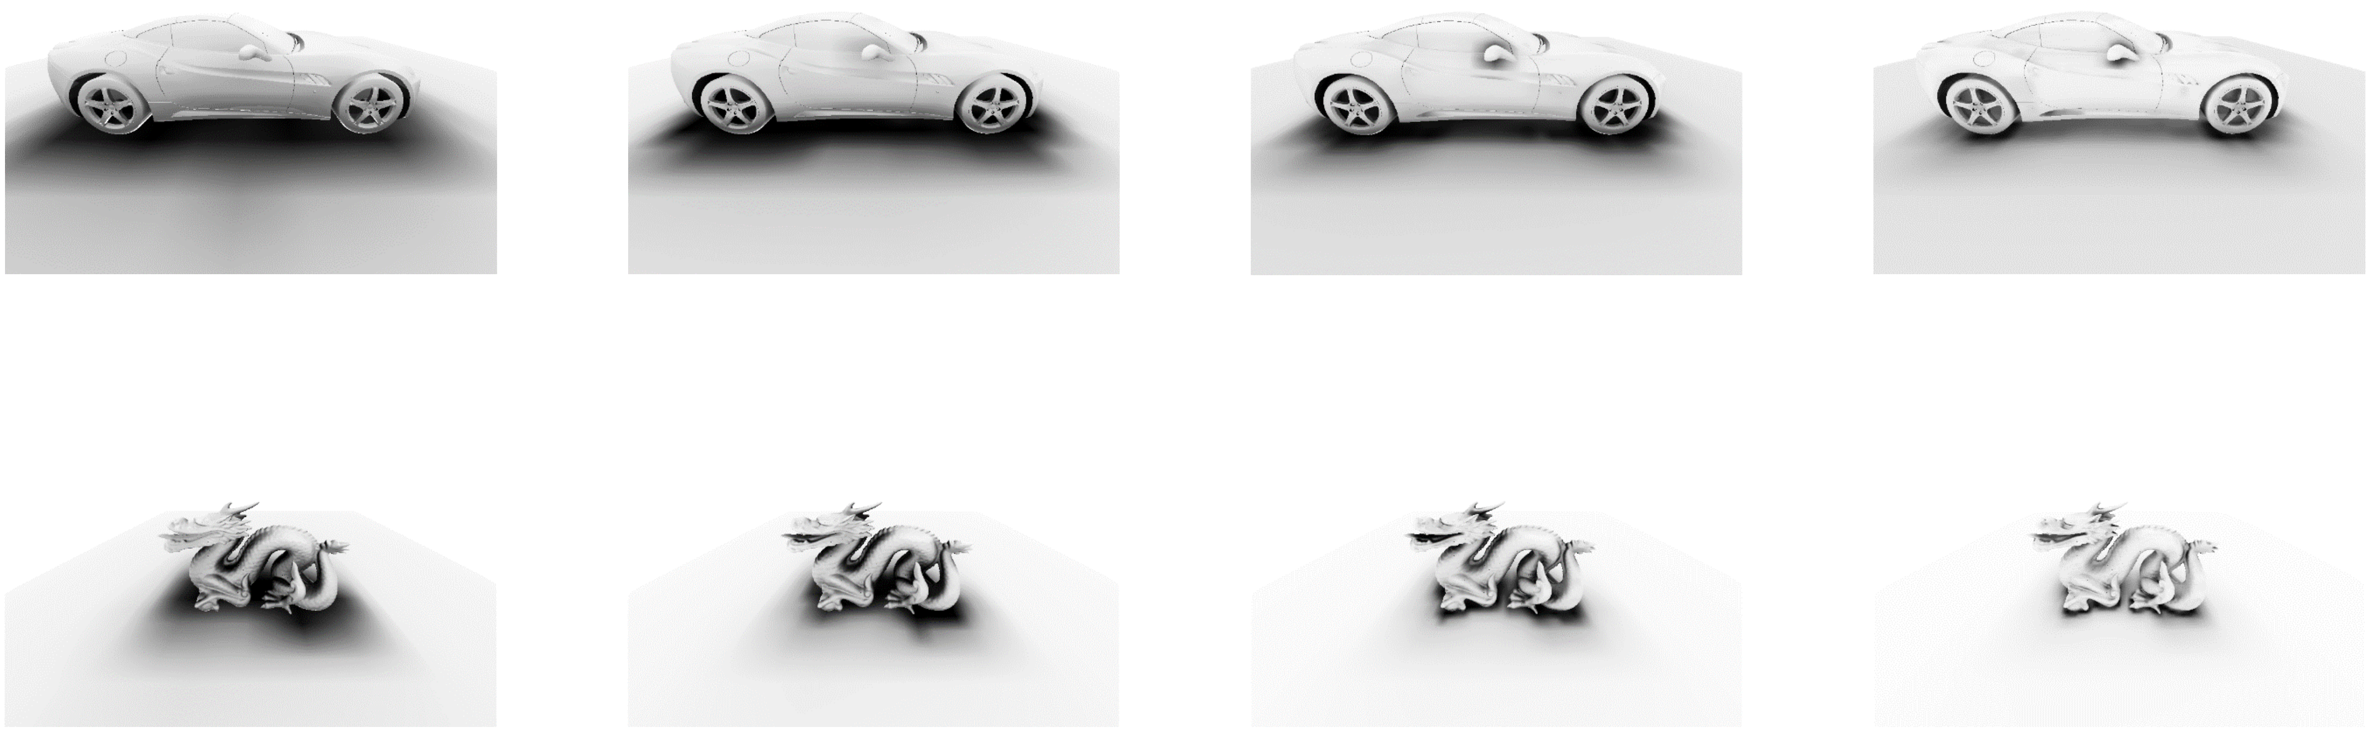

64 **increasing voxel volume** **512**

Dynamic Visibility – Voxel Density

64 increasing voxel volume 512

Dynamic Lighting

Peak Detection

without peak detection

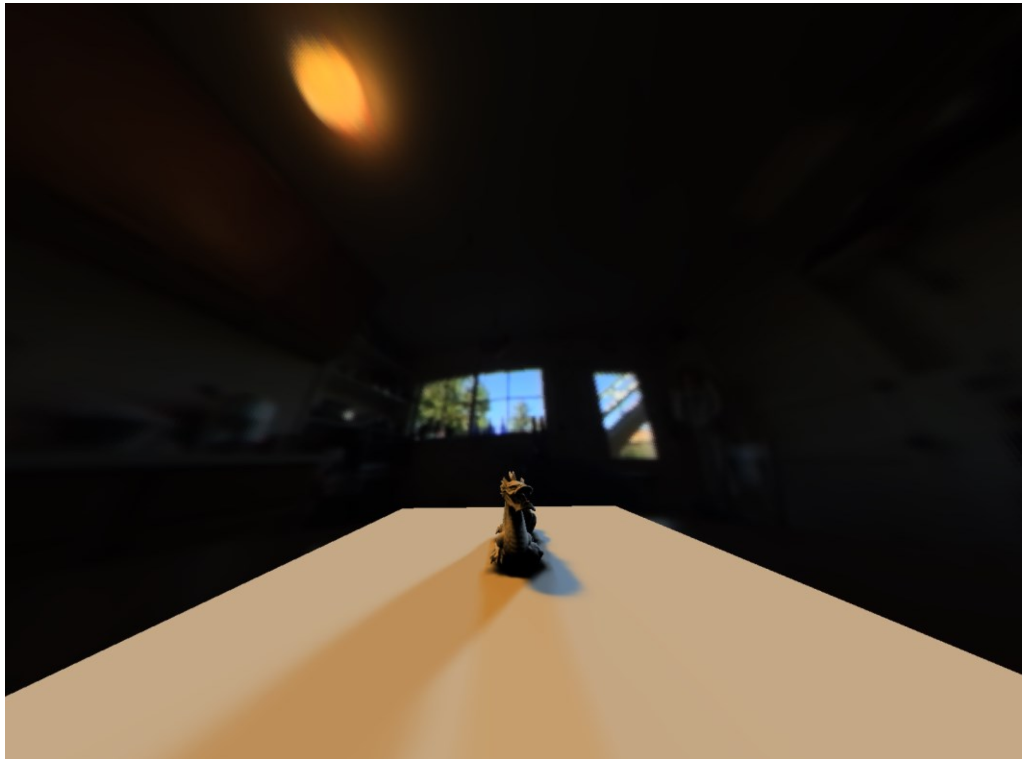

with peak detection

Relighting of Real Objects - **Inverse Rendering using SRBFs**

	Haar wavelets [Haber et al., 2009]	SRBF (ours)
input resolution	960×540	960×540
# images	5	5
# vertices	2689	2689
# lighting coefficients	1024	1020
# materials	3	3
# BRDF weights	**2689 (per vertex)**	**906240 (per texel)**
Δt optimize lighting	$1089.340\ s$	$10.189\ s$
Δt optimize materials	$2.516\ s$	$6.973\ s$
Δt optimize BRDF weights	$2.518\ s$	$180.837\ s$
Δt optimize one BRDF weights	$\mathbf{1.36 \times 10^{-4}\ s}$	$\mathbf{1.99 \times 10^{-4}\ s}$
Δt one iteration	$1099.680\ s$	$201.240\ s$
Δt full optimization	184 min	29 min
Δt rendering reconstruction	$\mathbf{13.403\ s}$	$\mathbf{0.022\ s}$

Relighting of Real Objects – **Intrinsic Image Decomposition**

- Intrinsic images decomposition using SRBFs *[Barron and Malik, 2015]*

$$\underset{G,R,L}{\text{maximize}} \quad P(G)P(R)P(L)$$

$$\text{subject to} \quad I = R + S(G,L)$$

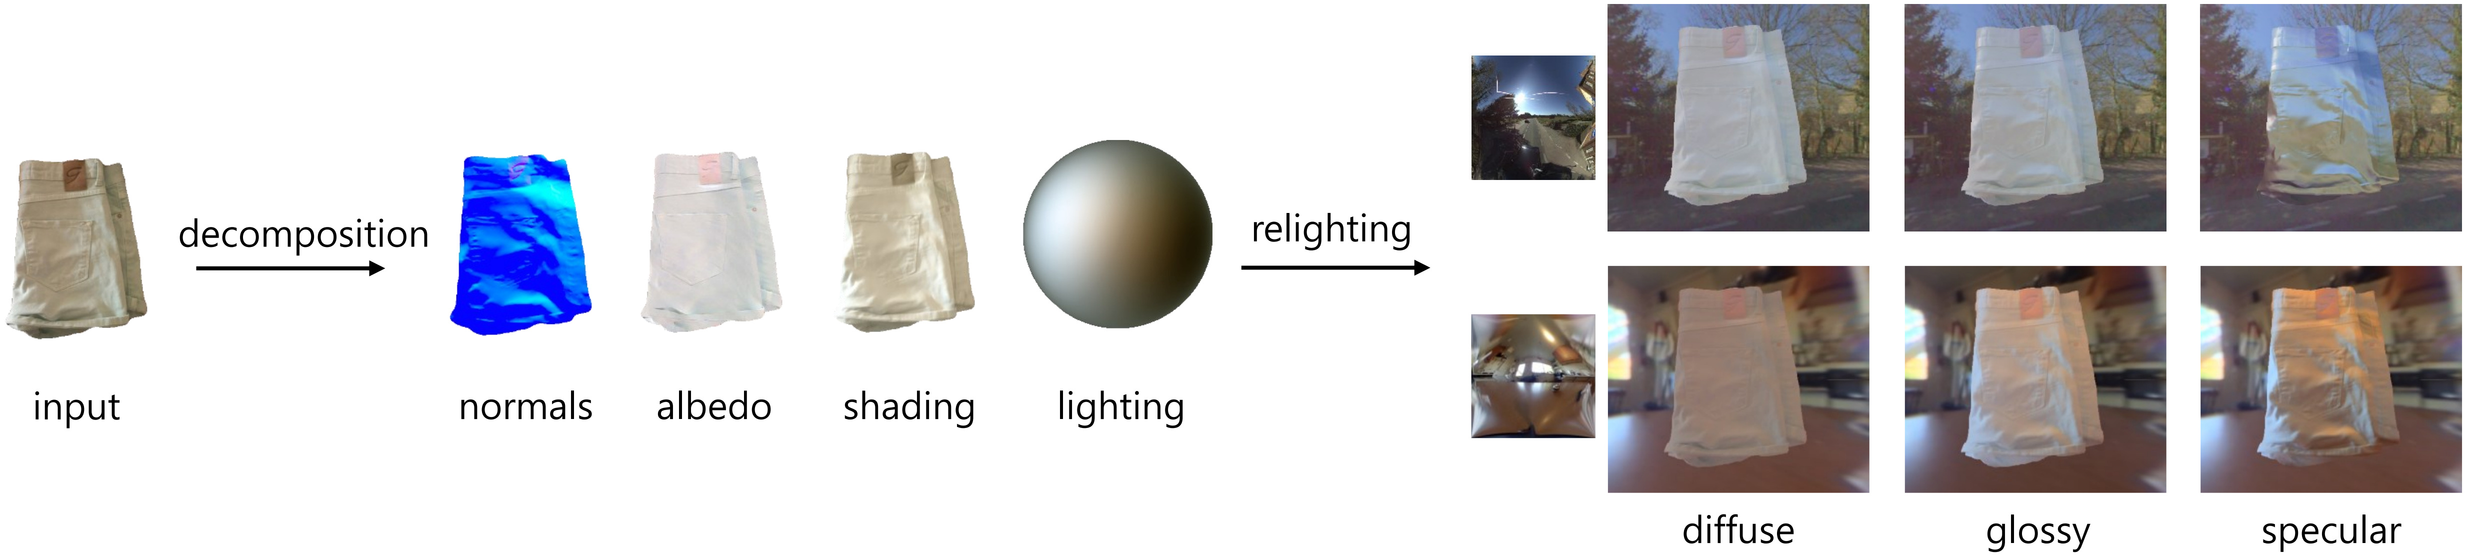

Relighting of **Virtual Objects** *[AIVIE, 2014]*

Relighting of **Real Objects** *[ICT-FP7 SCENE, 2014]*