Please contact me for PPT version w/ media files:
hullin@cs.uni-bonn.de

Computational Transient Imaging

Matthias Hullin, University of Bonn

Image processing Computational photography

Demosaicing
CVPR 2008

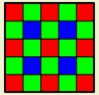

Simple lenses, deblurring
ACM ToG 2013, CVPR 2013

Computer vision

3D fluid capture
SIGGRAPH 2012

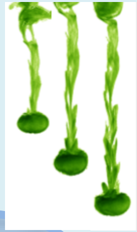

3D from diffuse indirect light
CVPR 2014

Automatic make-up suggestion
Eurographics 2011

3D scanning of glass
SIGGRAPH 2008

Light transport

Lens flare rendering
SIGGRAPH 2011

Transient imaging
SIGGRAPH 2013

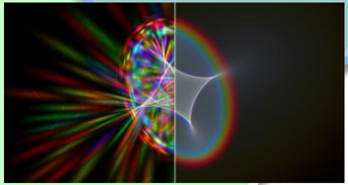

Polynomial optics
EGSR 2012

Color and appearance

Fluorescent materials
SIGGRAPH 2010

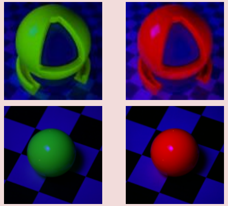

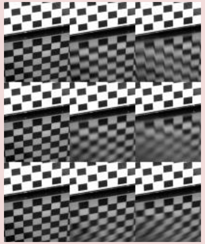

BRDF Display
Eurographics 2011

Physically-informed forward and inverse graphics

Why is computer vision hard?

- **Because images are ambiguous.**
- Every pixel value stands for a high-dimensional integral

$$I(x,y) = \int \cdots \int L\underbrace{(x', y', \theta_x, \theta_y, \lambda, \ldots, t)}_{\text{Most of this gets lost}}\, dA\, d\boldsymbol{\omega}\, d\lambda \cdots dt$$

- For any given image, there are infinitely many possible explanations
- Option 1: Imitate how our brain does it
- Option 2: Provide more information

Vision on generalized images

- λ: Spectral imaging [Cao11]

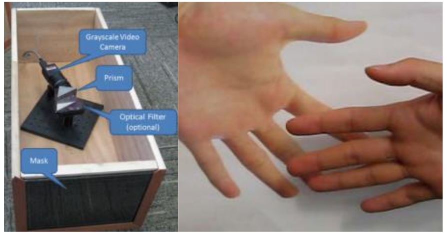

- θ_x, θ_y: 4D light fields [Wanner12]

- Light transport probing [O'Toole12]

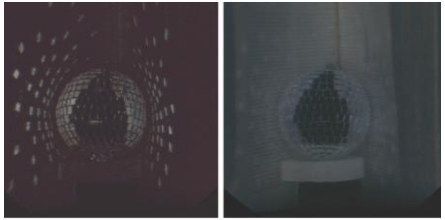

- …

Filming light in flight (= transient imaging)

- [Abramson78]

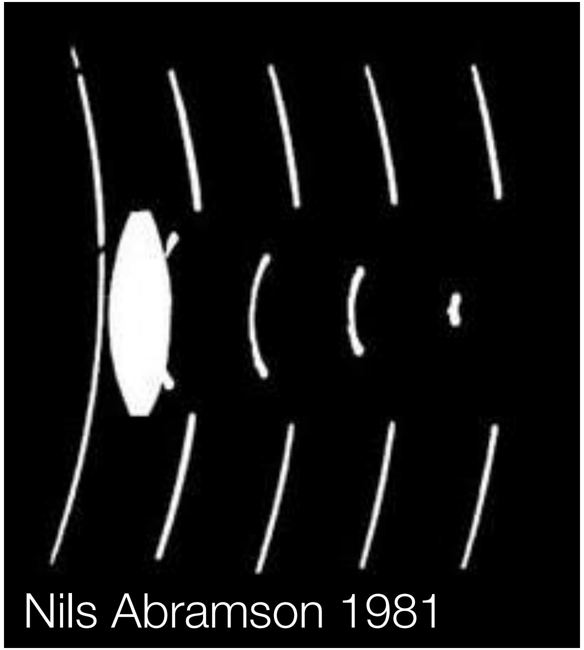
Nils Abramson 1981

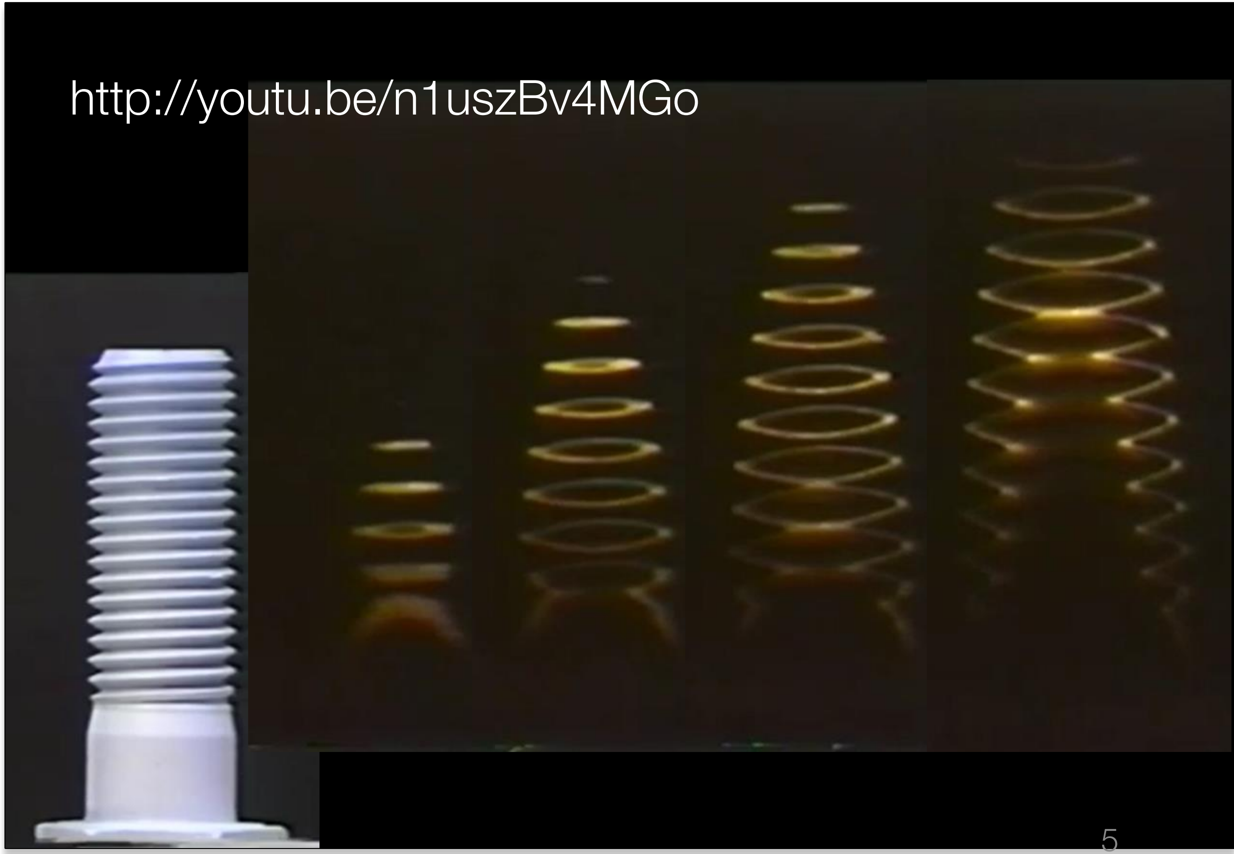

Transient imaging is important because...

$$I_{\text{out}}(x, y, t) = I_{\text{in}}(t) \underbrace{*}_{t} I_{\text{response}}(x, y, \tau)$$

- Impulse responses fully characterize linear time-invariant systems.

- Optical path length $\leftrightarrow$ scene geometry and reflectance

Research goals

1. **Instrumentation**
   Make it affordable and practical

2. **Image reconstruction and capture strategies**
   Fast and accurate acquisition of transient data

3. **Multi-path image formation**
   Understanding and inversion

(More on that from Ayush, Refael, Yebin)

Transient image as temporal impulse response

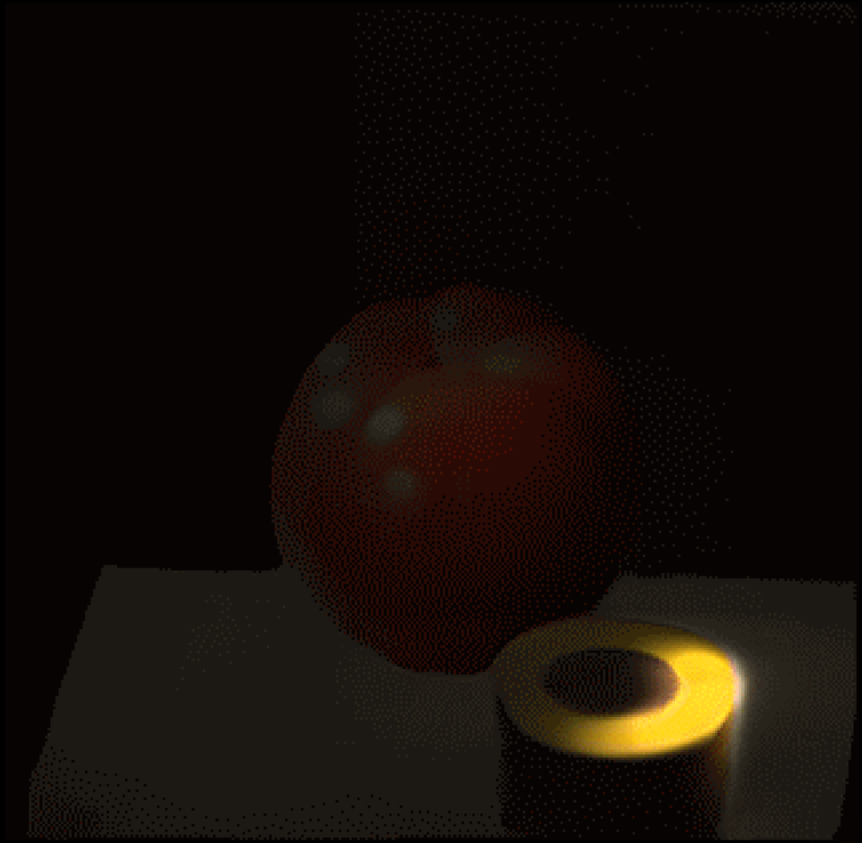

$$I(x,y,\tau)$$

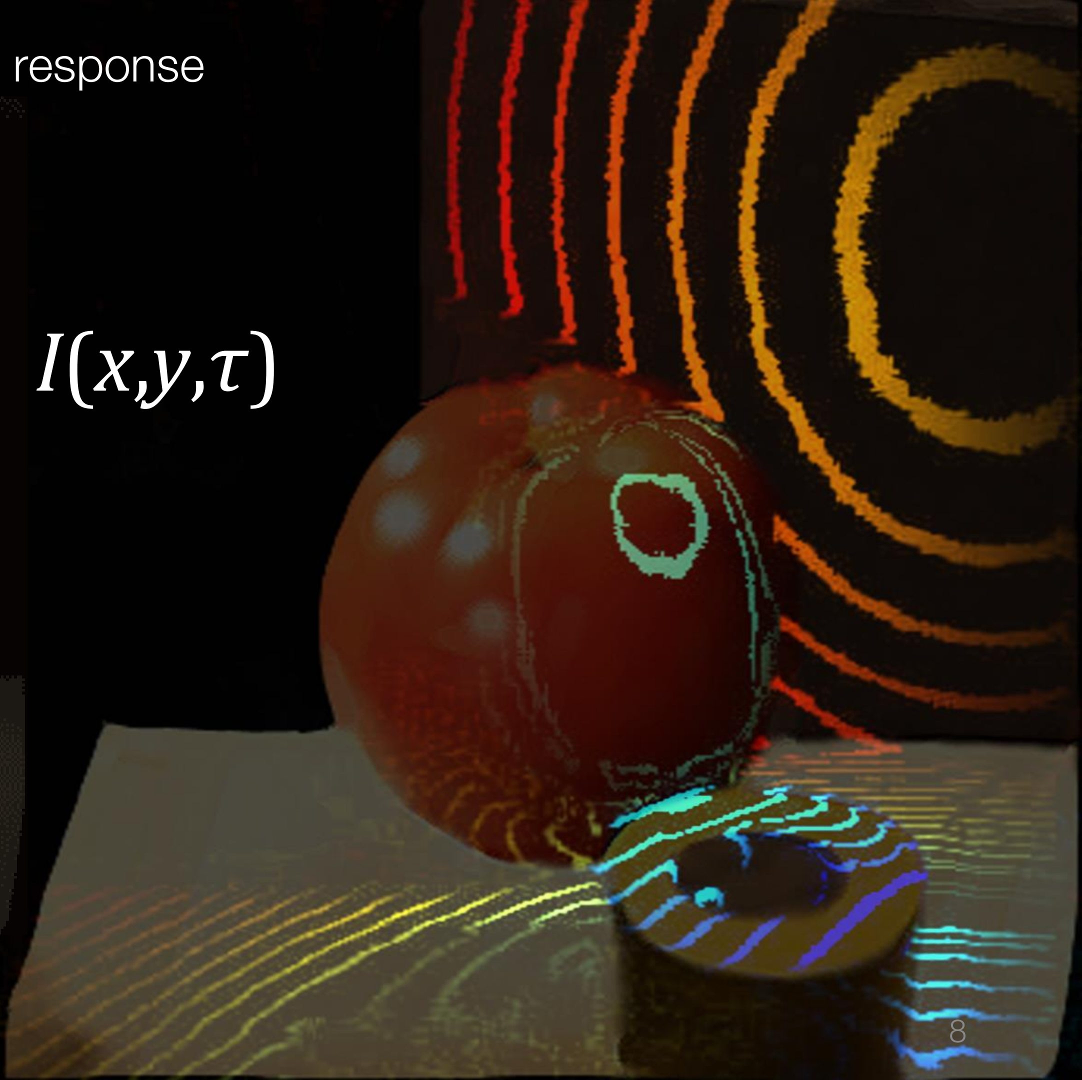

[Velten et al. 2011, 2012]

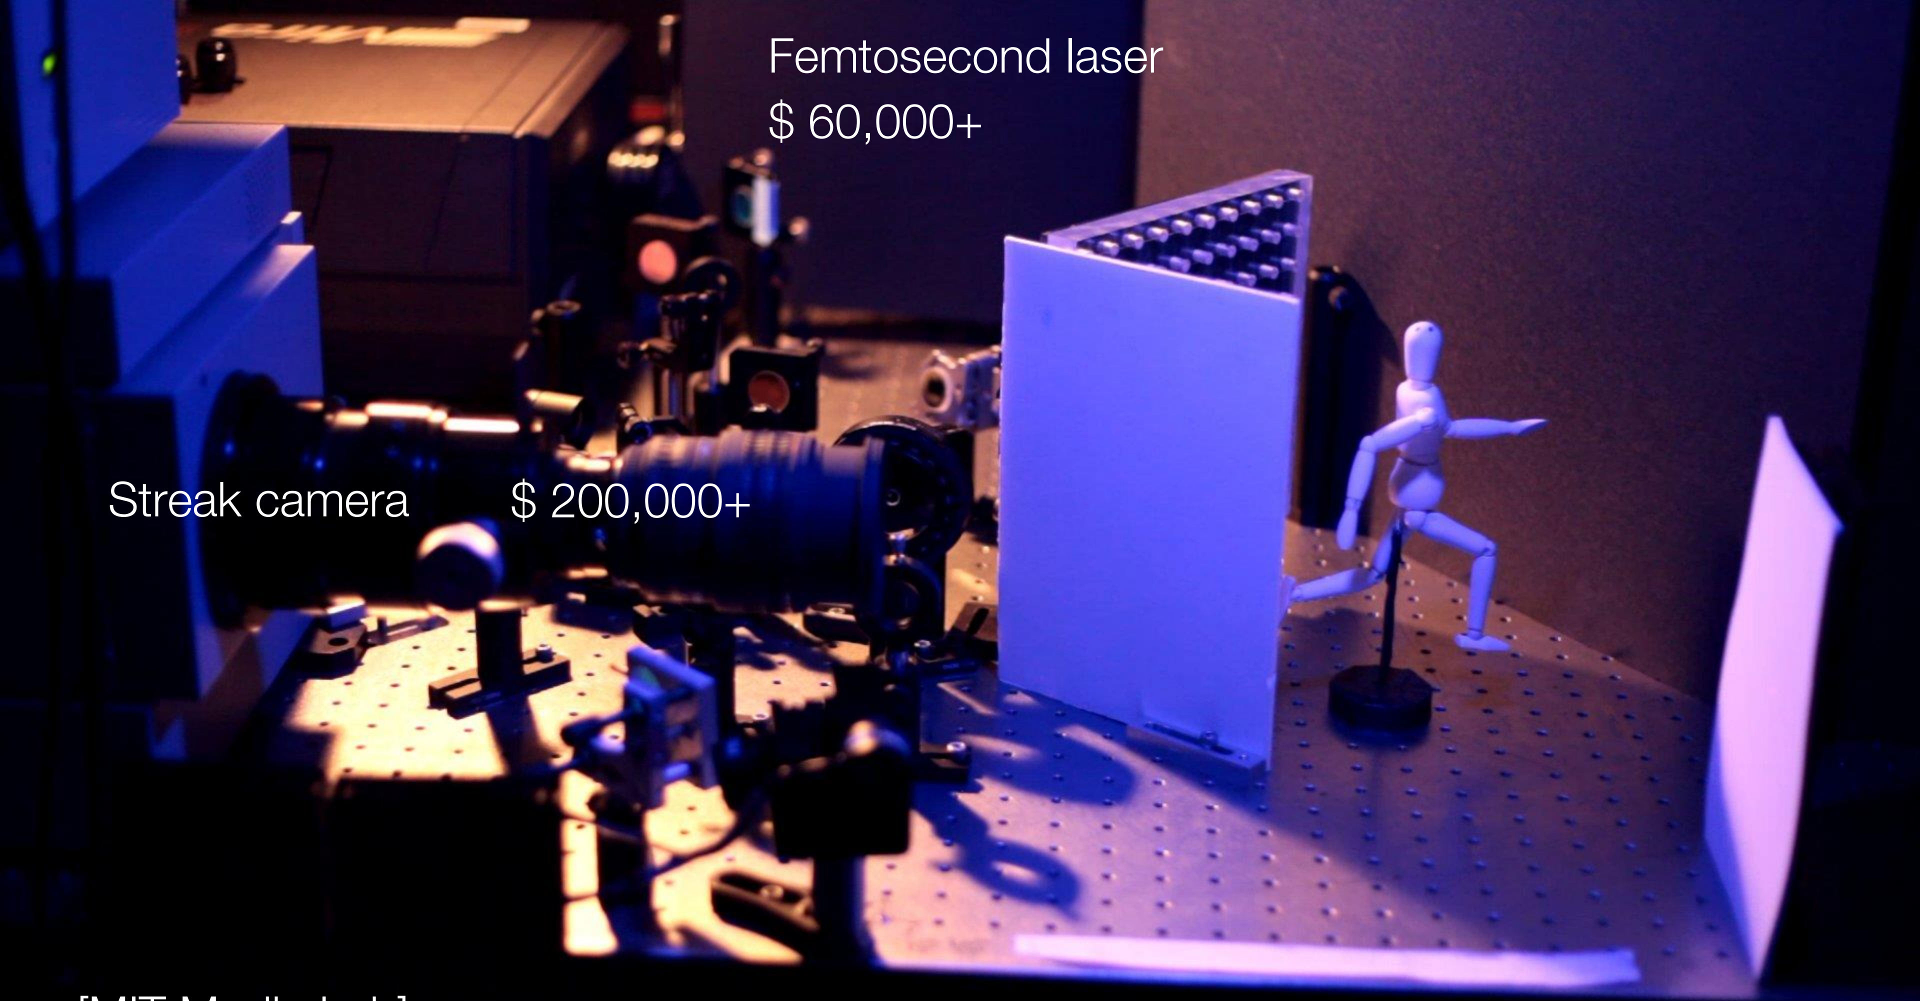
Femtosecond laser
$ 60,000+
Streak camera
$ 200,000+
[MIT Media Lab]

Low-Budget Transient Imaging using Photonic Mixer Devices

with Felix Heide, James Gregson, Wolfgang Heidrich

SIGGRAPH 2013 Technical Paper

Capturing modes

- Direct temporal sampling:
  streak sensors, fast photodetectors

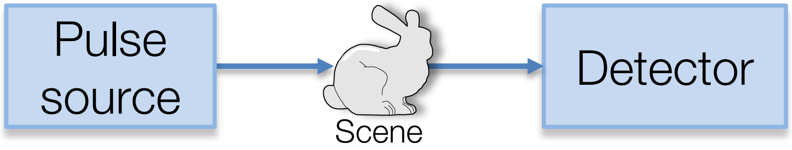

[Kirmani09]
[Velten12]

- Temporal coding

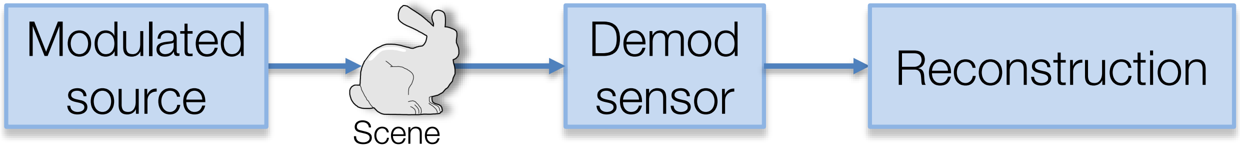

[Heide13]
[Kadambi13]
[Bhandari14]

Photonic Mixer Devices

PMDTechnologies PhotonICs® 19k-S3

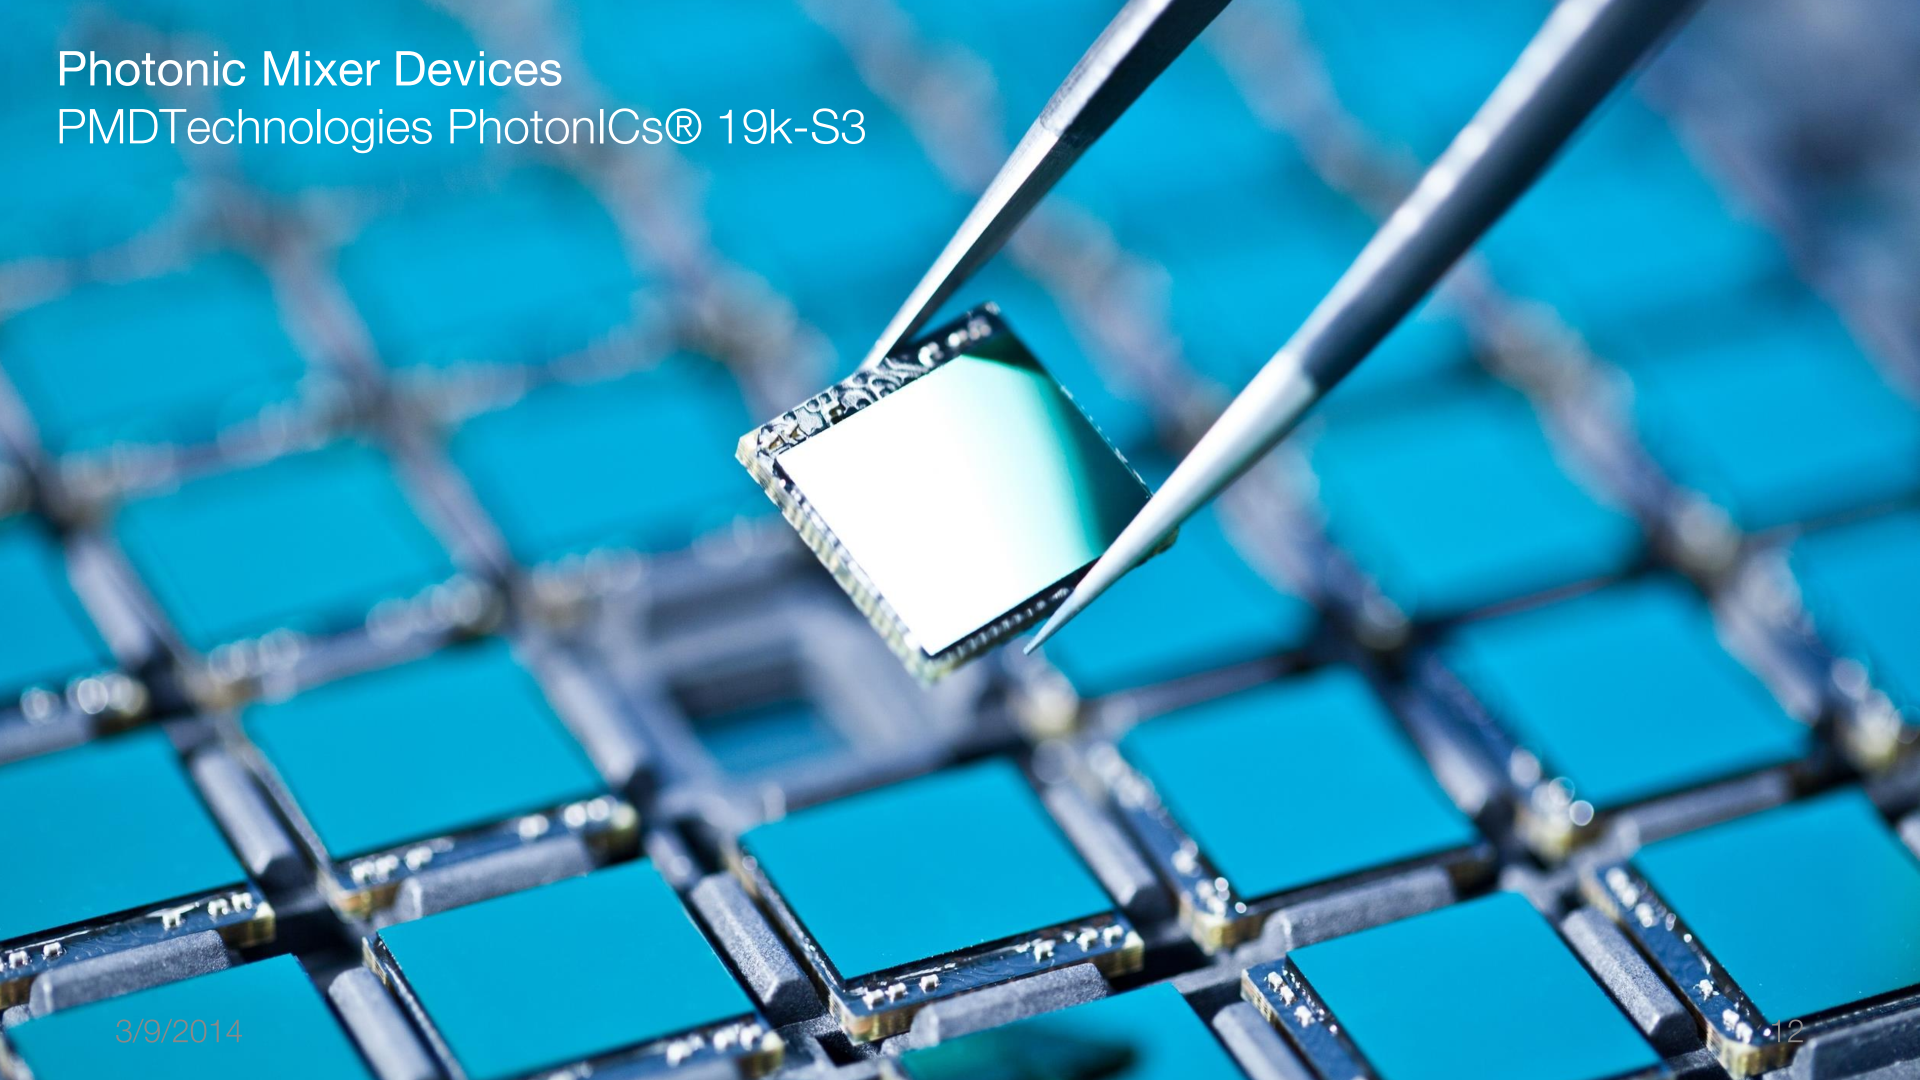

Hardware: PMDTechnologies CamBoard nano

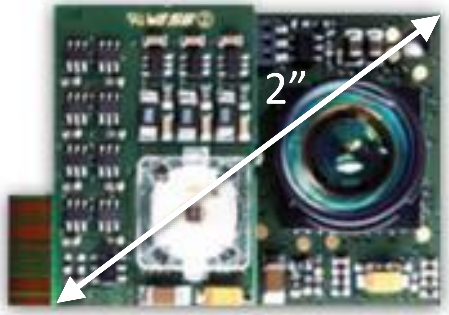

Too slow

Fixed frequency

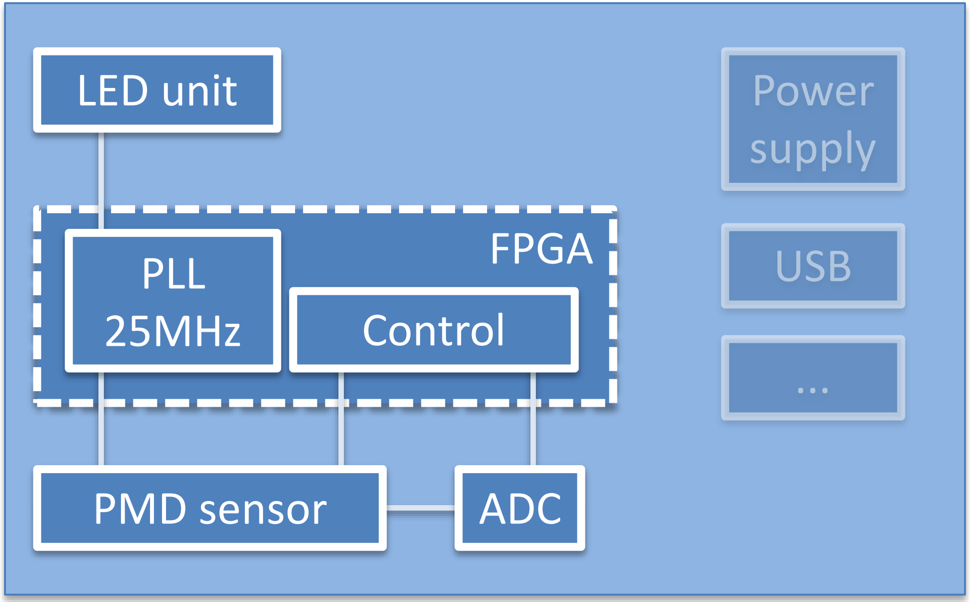

Modifications to the hardware

Late-2013 version: Max mod. frequency 180 MHz, stable up to 150 MHz

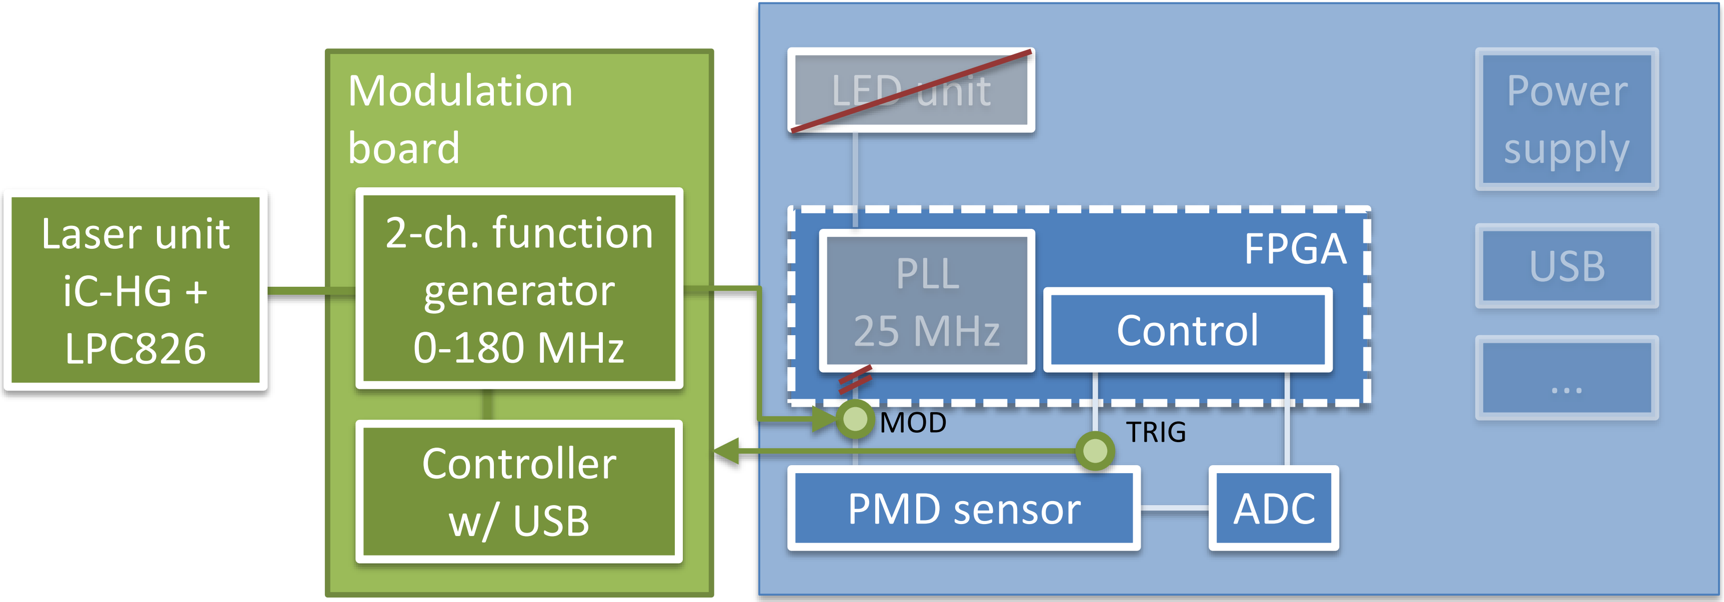

Illumination $g_\omega(t)$

Reference $f_{\omega,\phi}(t)$

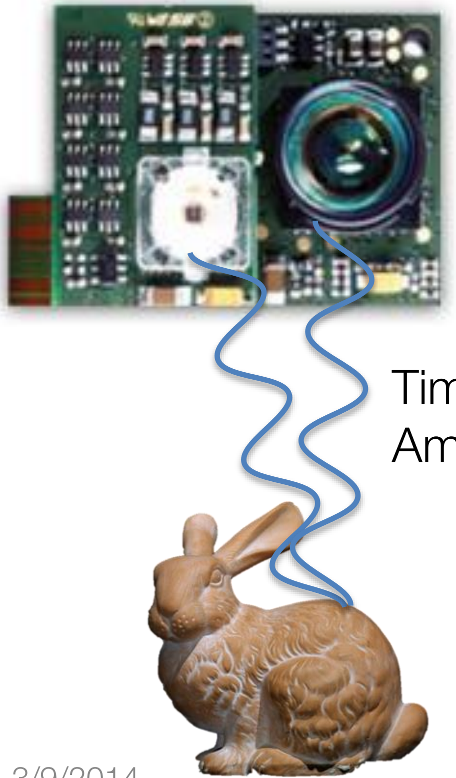

$$\tilde{g}_\omega(t) = \alpha g_\omega(t-\tau)$$

Time of flight τ

Amplitude α

Observed pixel value:

$$H_{\omega,\phi} = \int_0^T \tilde{g}_\omega(t) f_{\omega,\phi}(t)\,dt$$

$$= \alpha \int_0^T g_\omega(t-\tau) f_{\omega,\phi}(t)\,dt$$

$$\boxed{\begin{aligned} g_\omega(t) &= \sin(\omega t) \\ f_{\omega,\phi}(t) &= \sin(\omega t + \phi) \end{aligned}}$$

$$= \alpha \frac{T}{2} \cos(\phi - \omega\tau)$$

Traditional time-of-flight imaging:

Solve for τ, α for each pixel

Illumination $g_\omega(t)$

Reference $f_{\omega,\phi}(t)$

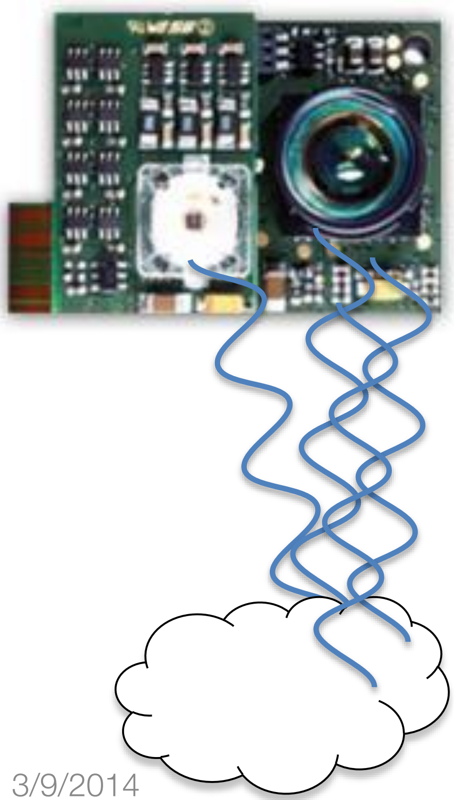

$$\tilde{g}_\omega(t) = \int_{\tau_1}^{\tau_2} \alpha(\tau) g_\omega(t-\tau)\, d\tau$$

Transient pixel $\alpha(\tau)$:
Superposition of delayed and attenuated reflections

("backscatter" in
Daniel Freedman's nomenclature)

Discretize τ:

$$H_{\omega,\phi} = \int_0^T \sum_k \alpha_k g_\omega(t-\tau_k) f_{\omega,\phi}(t)dt$$

$$= \sum_k \alpha_k \underbrace{\int_0^T g_\omega(t-\tau_k) f_{\omega,\phi}(t)dt}_{c_{\omega,\phi}(\tau_k)=\frac{T}{2}\cos(\phi-\omega\tau_k)}$$

Correlation coefficient

Vary $\omega = \omega_j$:
Linear system

$$H_j = \sum_k \alpha_k c_{j,k}$$

$$\mathbf{h} = \mathbf{Ci}$$

Capturing modes

- Direct temporal sampling:
  streak sensors, fast photodetectors

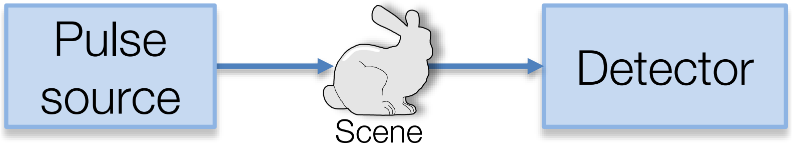

$\mathbf{i}$

[Kirmani09]
[Velten12]

- Temporal coding

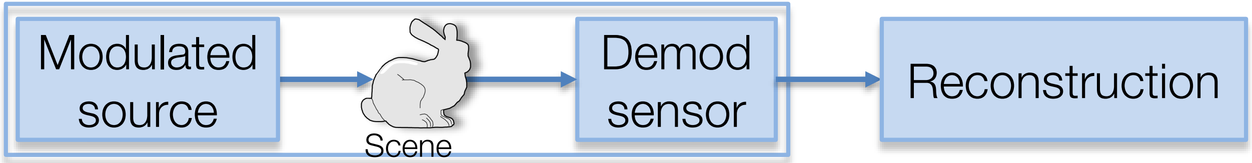

$\mathbf{h} = \mathbf{C}\mathbf{i}$ $\quad$ $\mathbf{i} = \mathbf{C}^{-1}\mathbf{h}$

[Heide13]
[Kadambi13]
[Bhandari14]

Correlation matrix $\mathbf{C}$

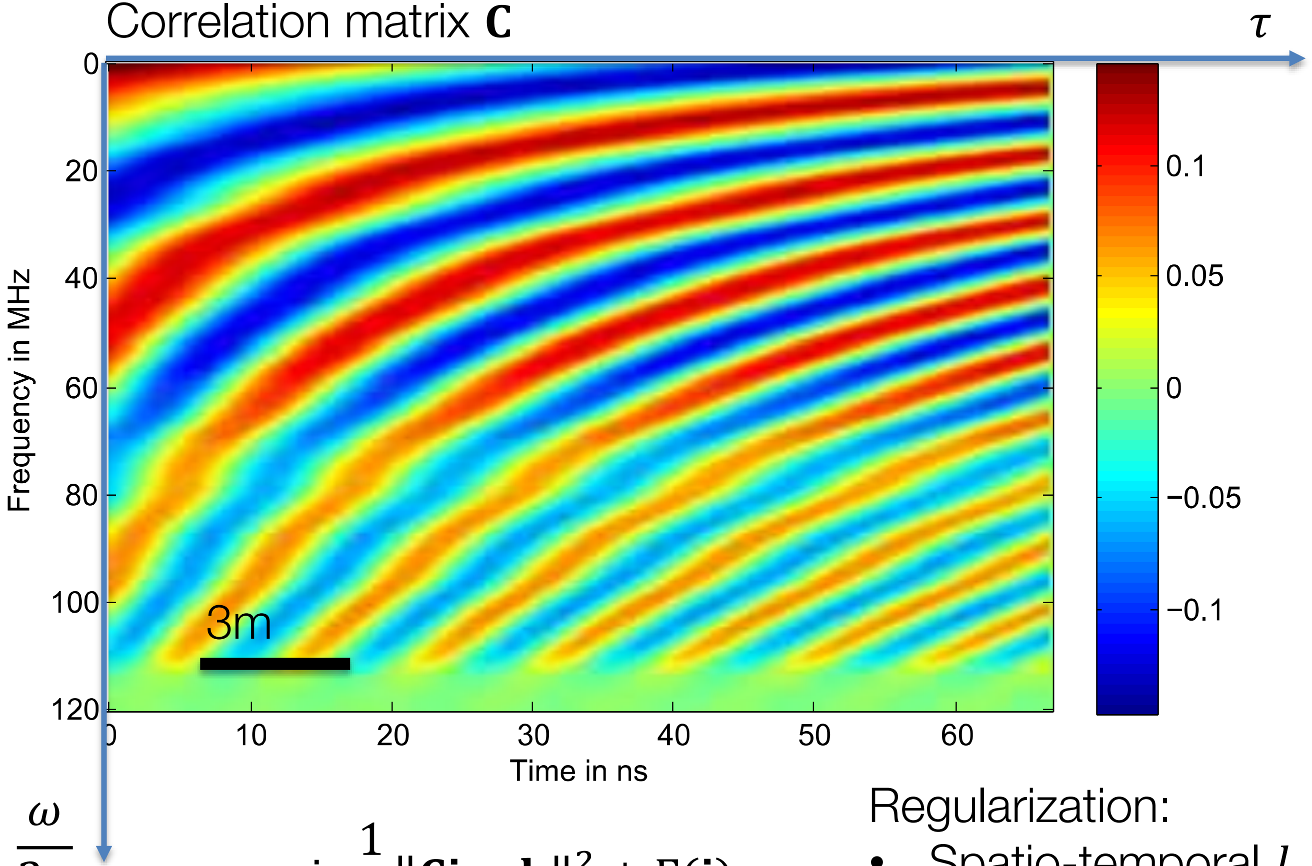

$$\frac{\omega}{2\pi}$$

$$\underset{\mathbf{i}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{C}\mathbf{i} - \mathbf{h}\|_2^2 + \Gamma(\mathbf{i})$$

Regularization:

- Spatio-temporal l_1/l_2 hybrid prior [Heide13]
- Parametric model [Wu12]
- **Better: application-specific [Heide14]**

Results ‘Bottles’:

Different time-steps:

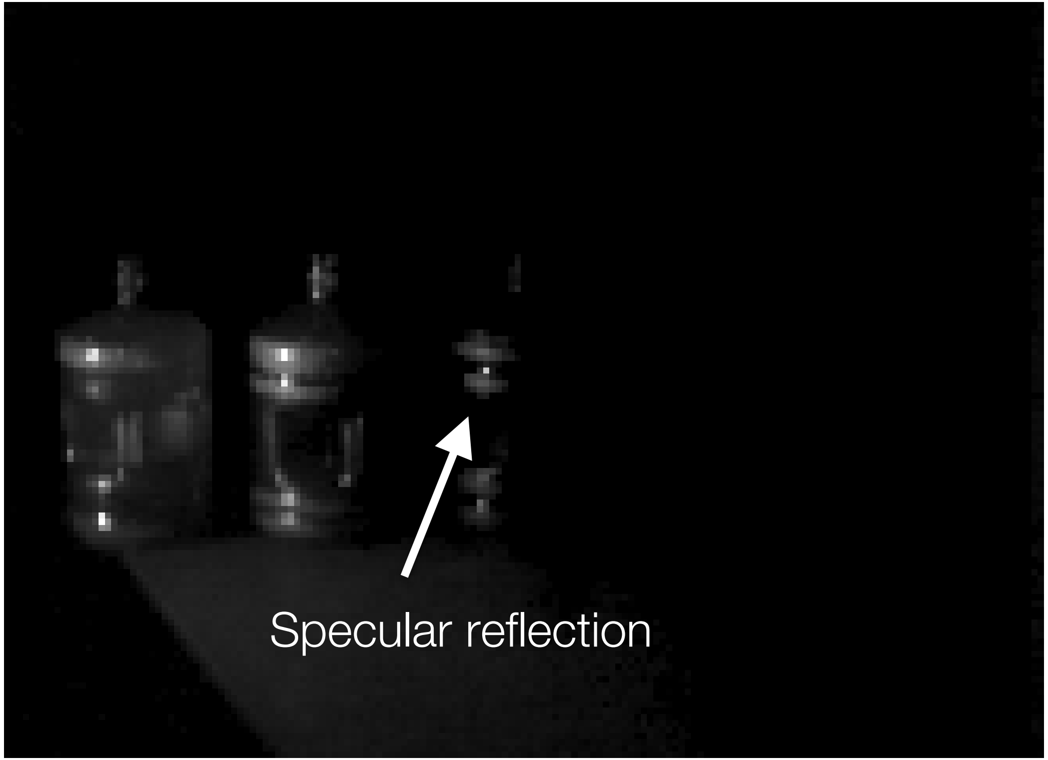

Results ‘Bottles’:

Different time-steps:

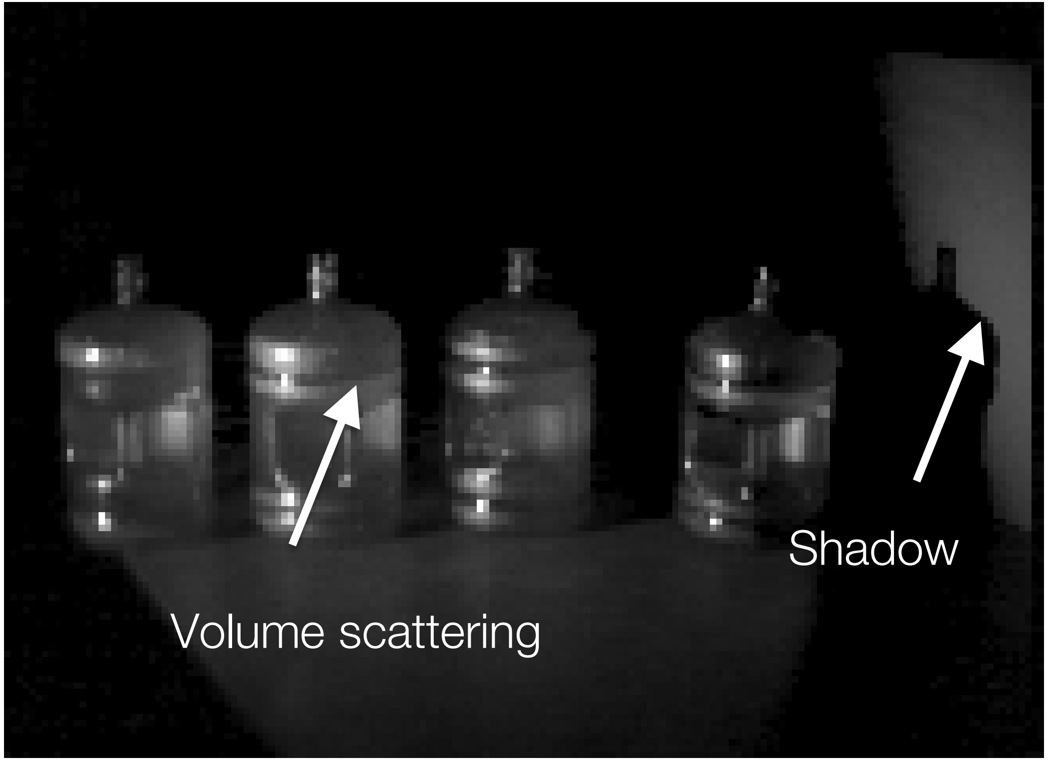

Results ‘Bottles’:

Different time-steps:

Results ‘Bottles’:

Different time-steps:

Demo at SIGGRAPH 2013

SIGGRAPH 2013
E-Tech implementation:
3 frequencies only;
fully automatic calibration
in 2 min, TI capture 30 sec

Comparison

	Streak camera [Velten12]	Correlation-based [Heide13]
Slicing	Scanline	Frequency/Phase
Cost	$300,000	$500
Ambient illum.	Very sensitive	Insensitive
Capture time	Hours	Seconds
Resolution	> 2ps	Scene-dependent

Camera we actually used in SIGGRAPH 2013 submission

- (bonus slide for ITоF2014 handouts)

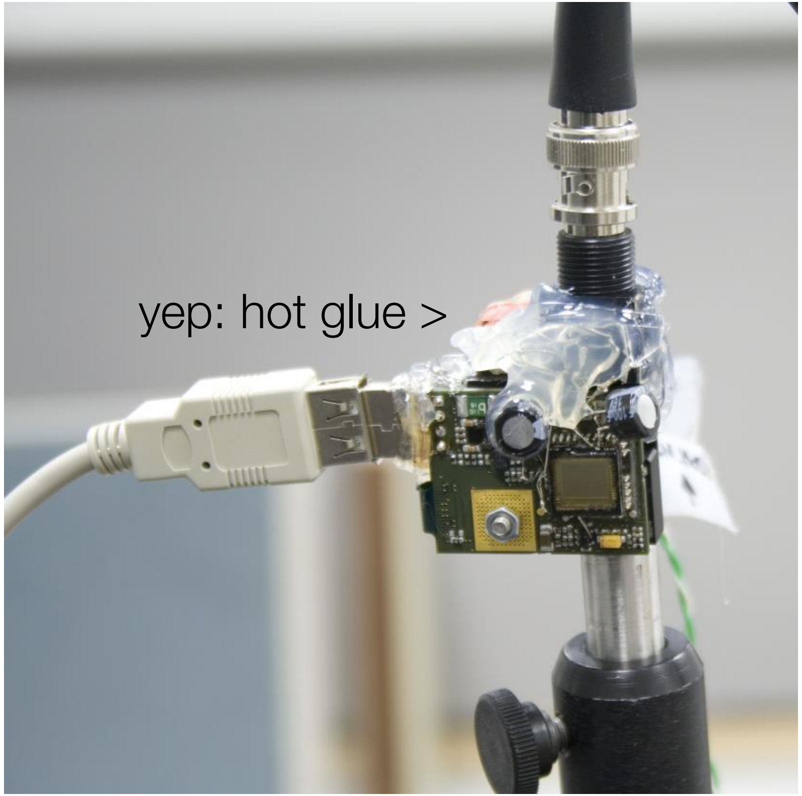

Research goals

- **Instrumentation**
  Make it affordable and practical

- **Image reconstruction and capture strategies**
  Fast and accurate acquisition of transient data

- **Multi-path image formation**
  Understanding and inversion

Indirect Diffuse Time-of-Flight Imaging

2010, with Johannes Hanika, Hans-Peter Seidel, Hendrik Lensch

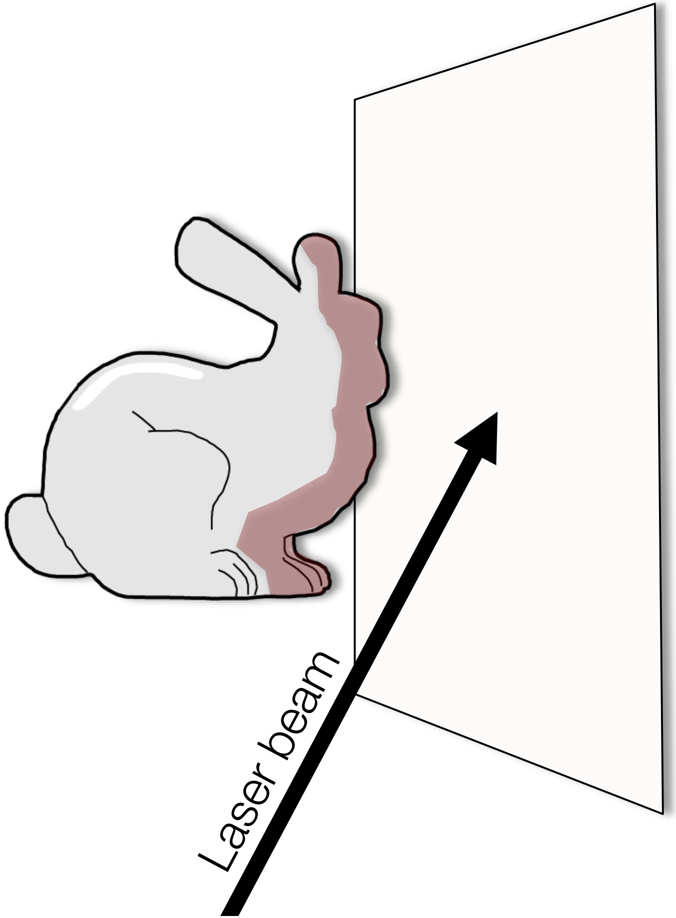

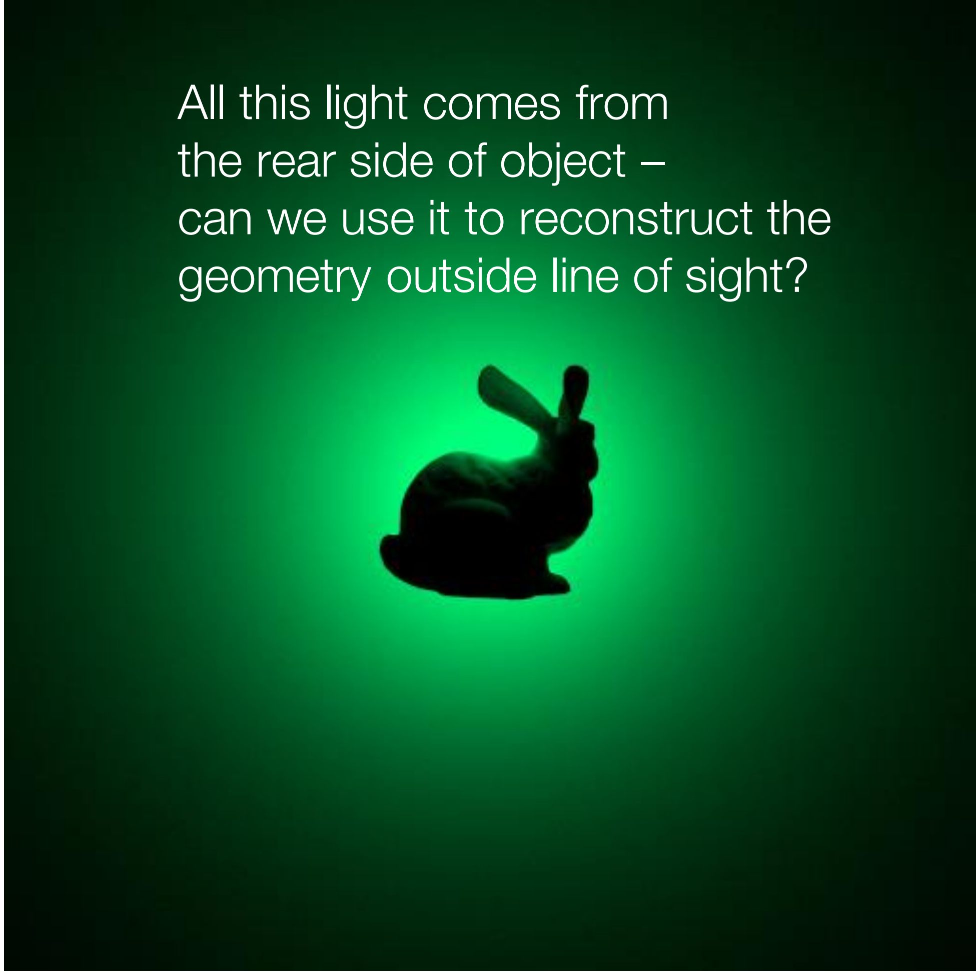

What can **indirect diffuse** light tell us about an unknown scene?

Transient response

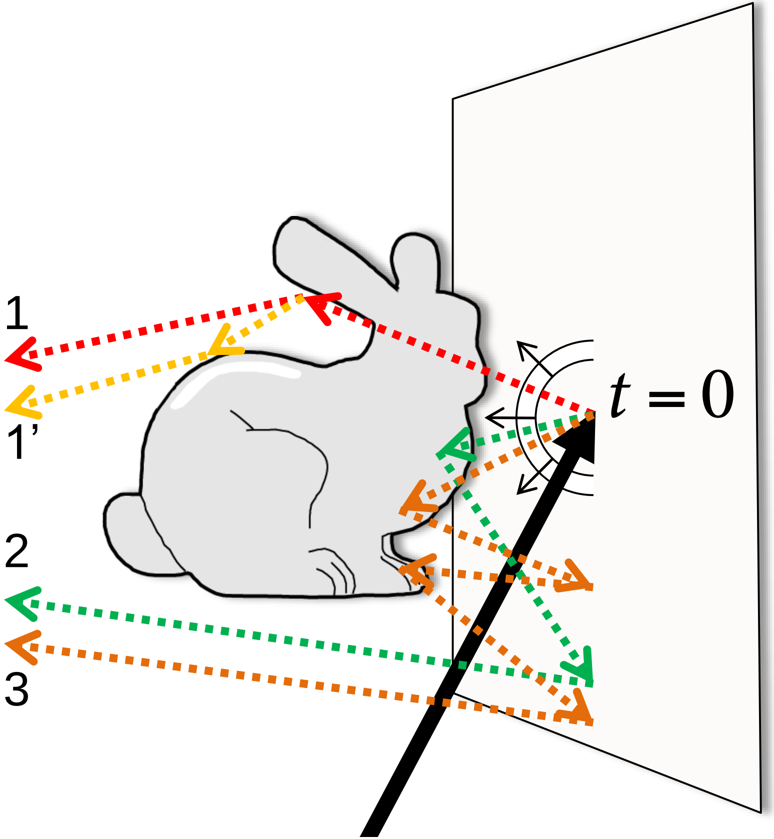

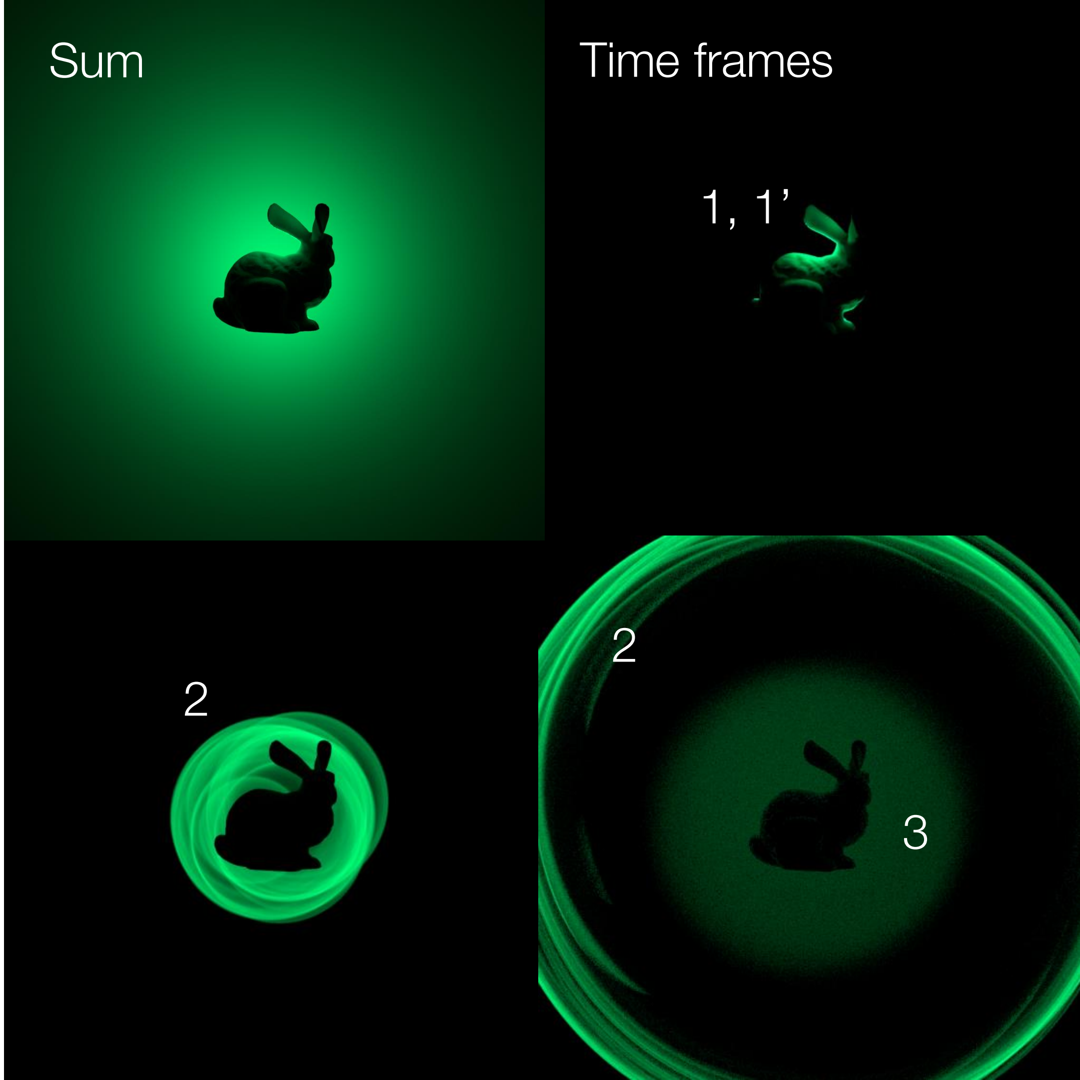

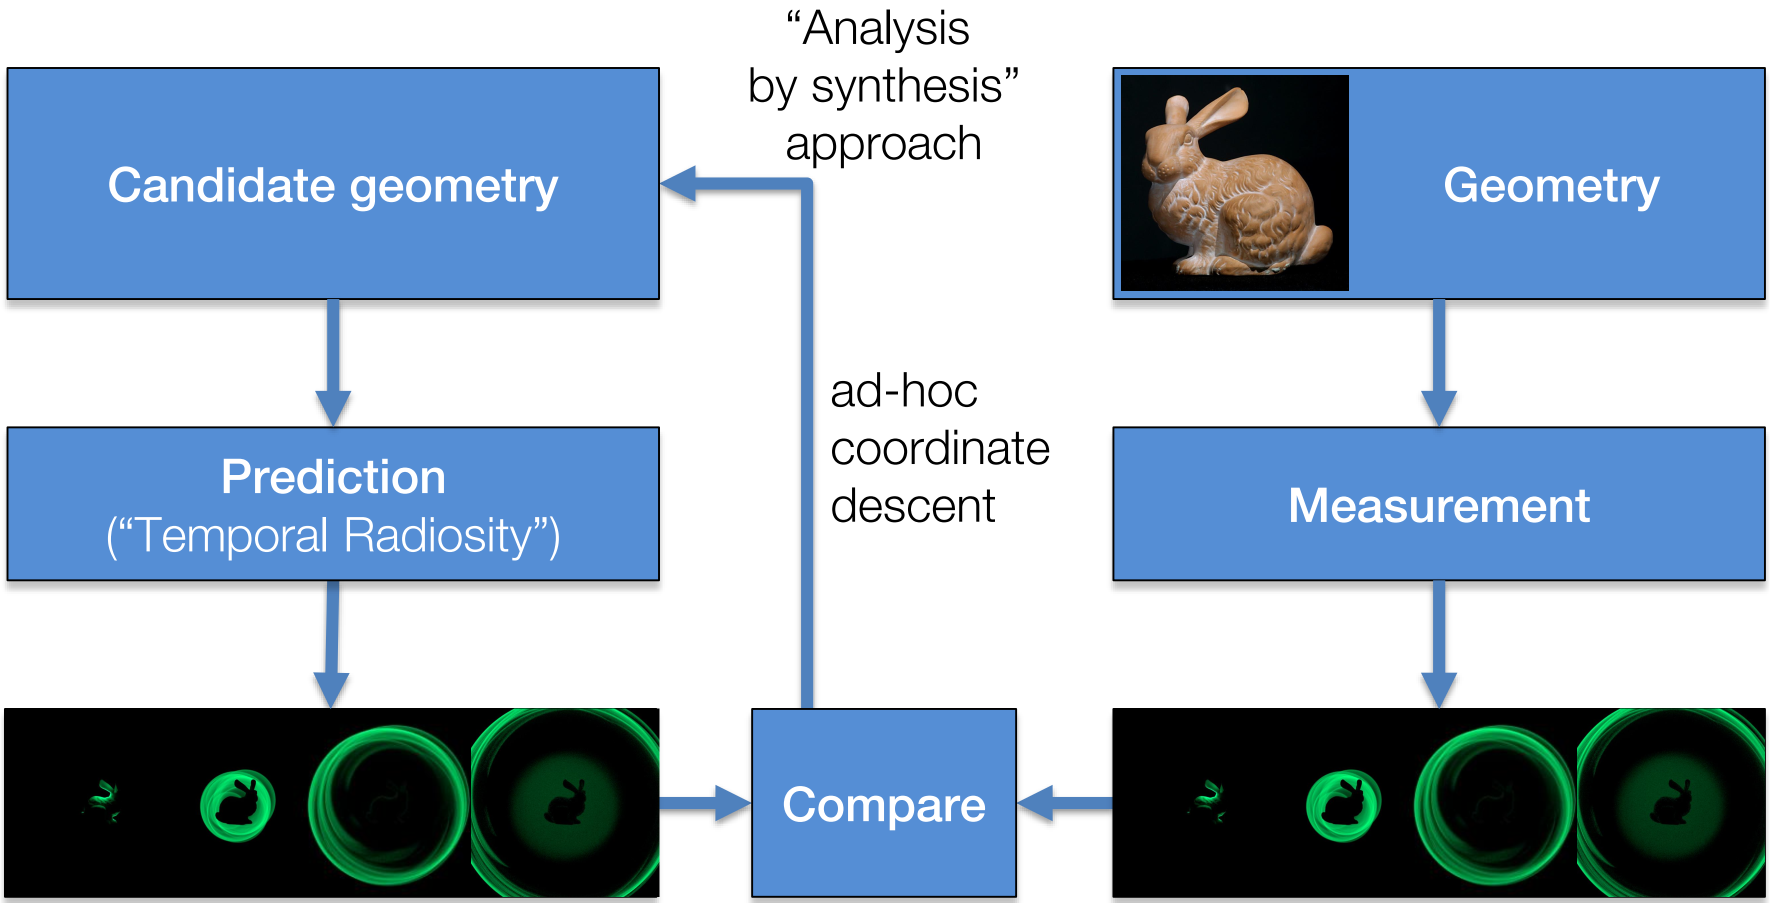
"Analysis
by synthesis"
approach
Candidate geometry
Geometry
ad-hoc
coordinate
descent
Prediction
("Temporal Radiosity")
Measurement
Compare

“Looking around corners”

- How can it work?

- Diffuse reflection destroys angles, but leaves time intact

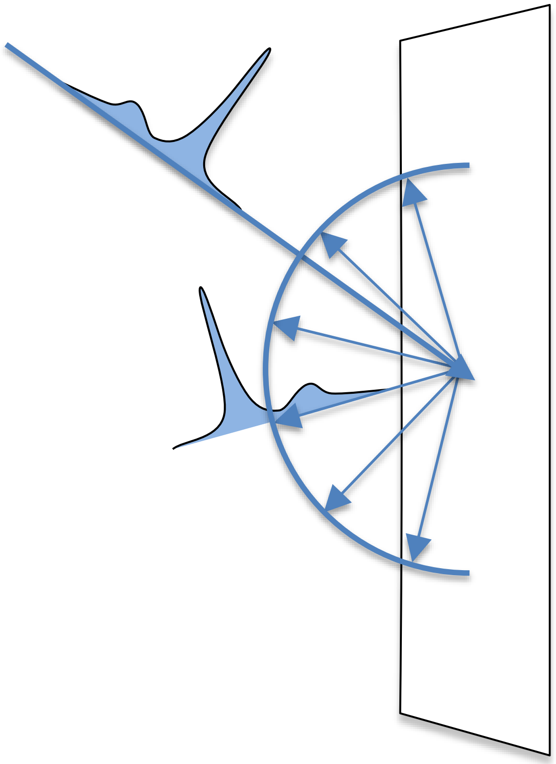

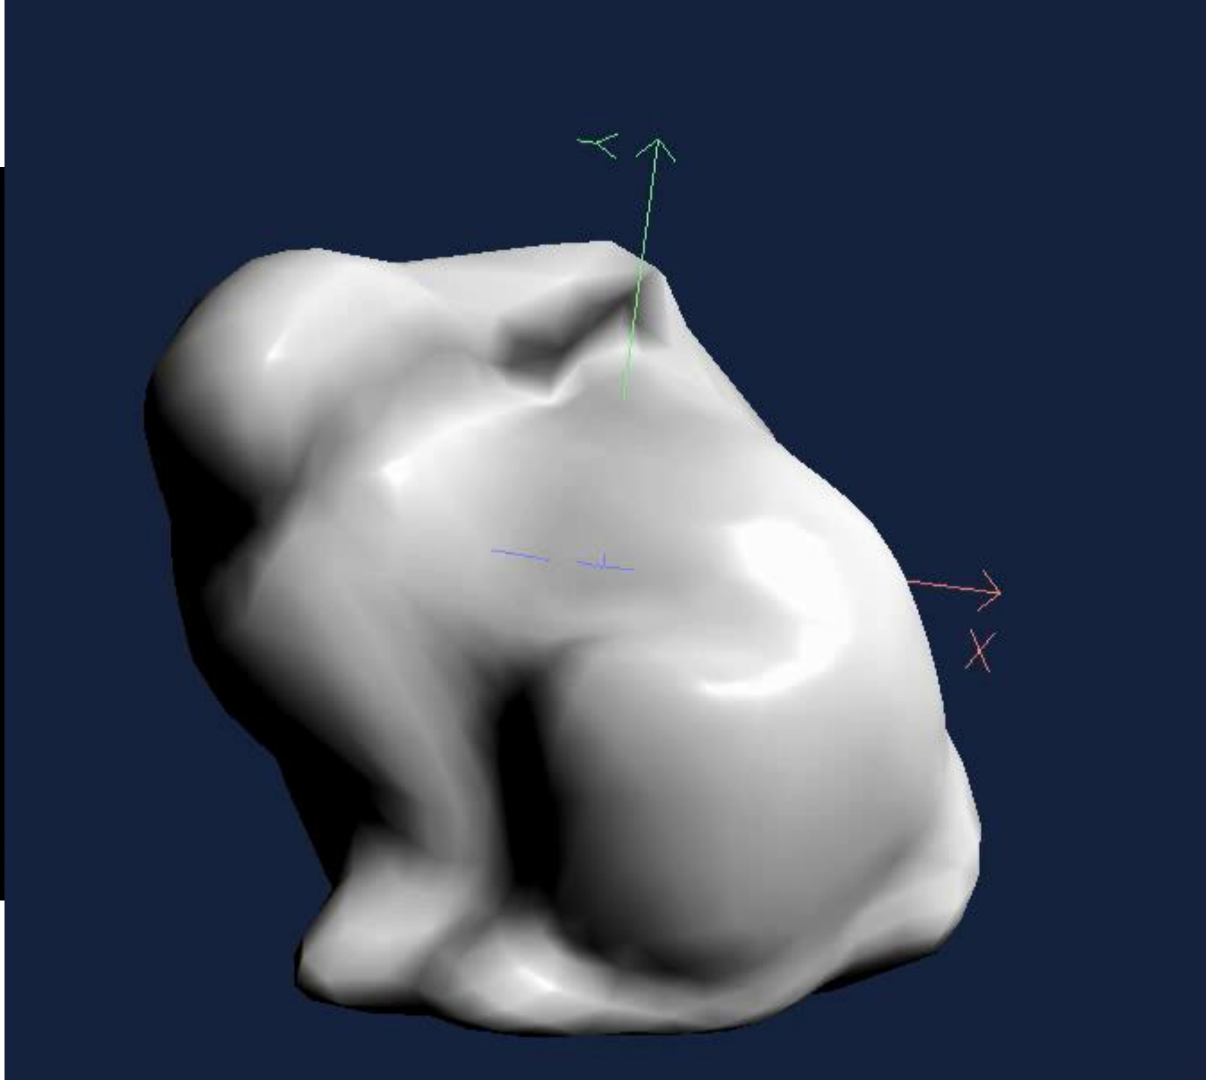

Pilot experiment (2010)

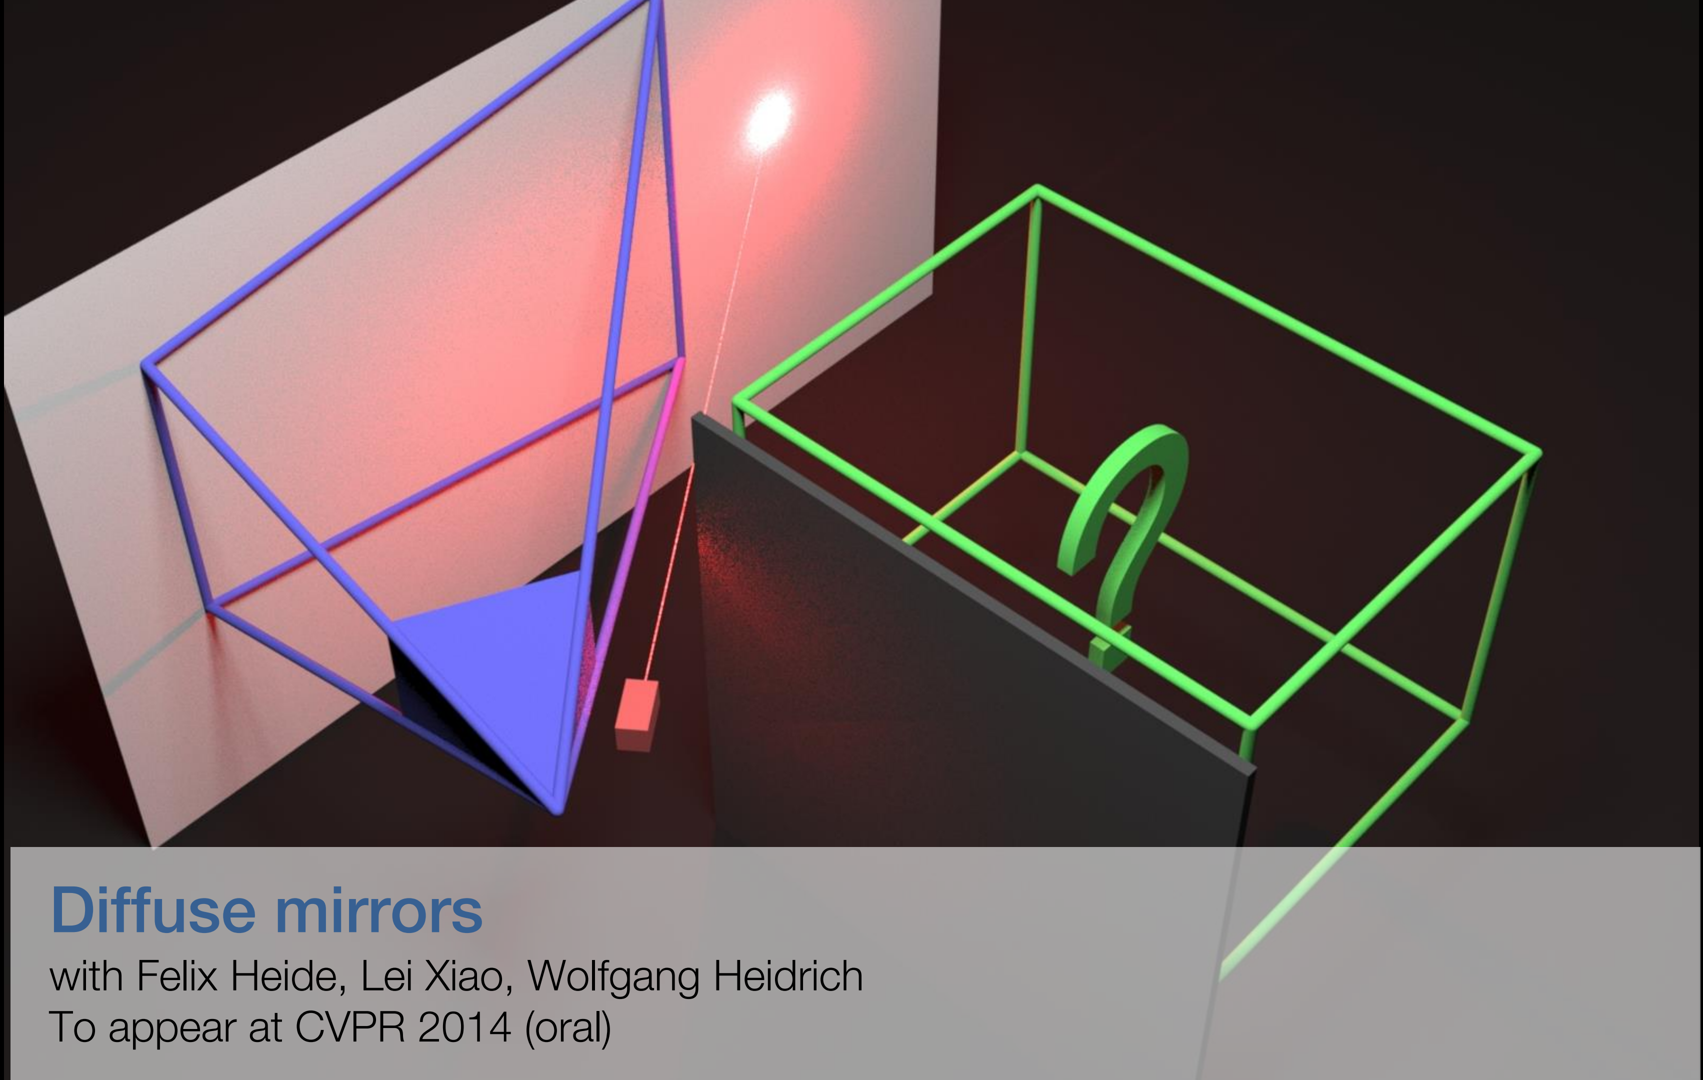

Diffuse mirrors

with Felix Heide, Lei Xiao, Wolfgang Heidrich

To appear at CVPR 2014 (oral)

“Looking around corners”

[Velten12]

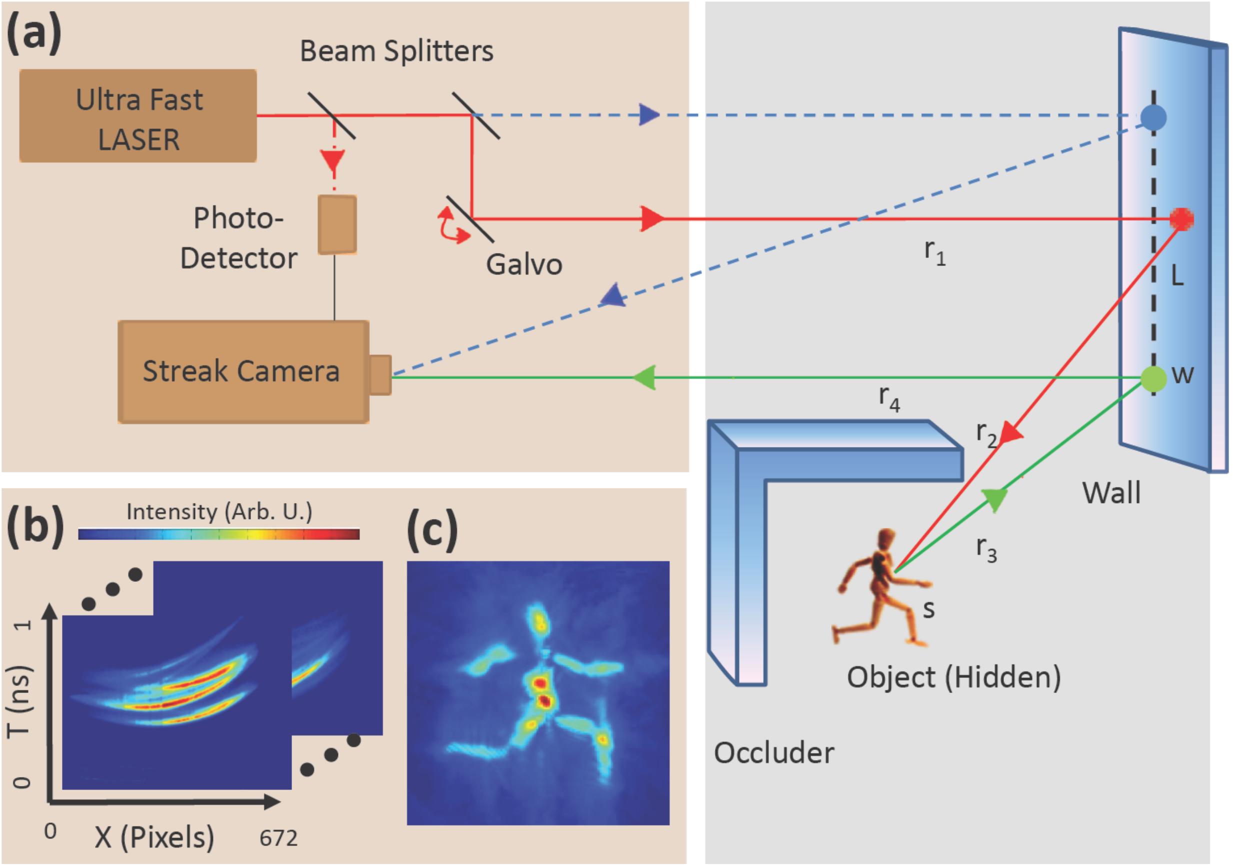

Scene

- Lambertian scene and wall
- Single-bounce, occlusion-free scattering in the scene

[Heide14]:

- Introduce generalized albedo $\mathbf{v}$
- Linear image formation model:

$$\mathbf{i} = \mathbf{P}\mathbf{v}$$

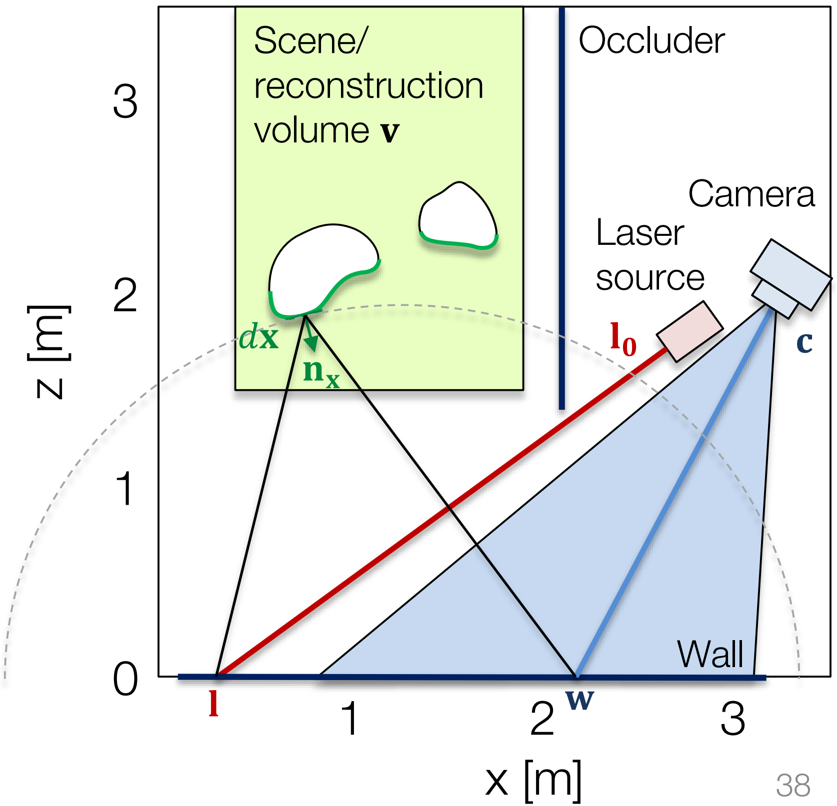

Joint solution

- Direct temporal sampling [Velten12]:

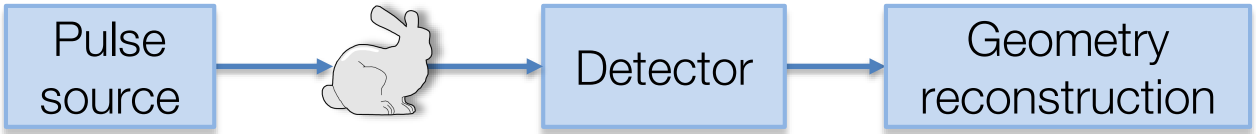

$$\mathbf{i} = \mathbf{P}\mathbf{v} \qquad \mathbf{v} \cong \mathbf{P}^T(\mathbf{i})$$

- **Temporal coding:**

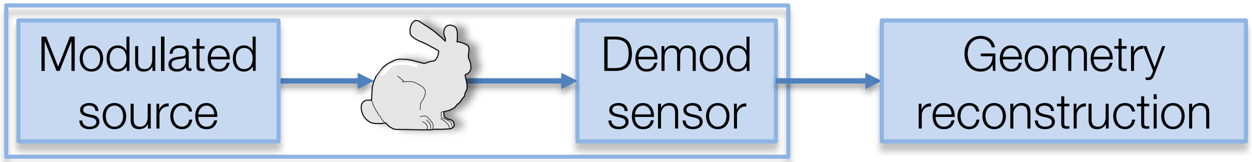

$$\mathbf{h} = \mathbf{C}\mathbf{i} = \mathbf{C}\mathbf{p}(\mathbf{v}) \qquad \mathbf{v} = \mathbf{p}^{-1}(\mathbf{C}^{-1}\mathbf{h})$$

Linearized model: $\mathbf{h} = \mathbf{C}\mathbf{i} \cong \mathbf{C}\mathbf{P}\mathbf{v}$ $\qquad \mathbf{v} \cong (\mathbf{C}\mathbf{P})^{-1}\mathbf{h}$

Inverse problem

$$\mathbf{v}_{\text{opt}} = \underset{\mathbf{v}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{CPv} - \mathbf{h}\|_2^2 + \Gamma(\mathbf{v})$$

$$\Gamma(\mathbf{v}) = \underbrace{\lambda \sum_z \left\| \nabla_{x,y} \mathbf{v}_z \right\|_1}_{\text{Sparse xy-gradient}} + \underbrace{\theta \|\mathbf{Wv}\|_1}_{\text{Volume sparsity 1}} + \underbrace{\omega \sum_{x,y} \operatorname{ind}_{\mathcal{C}}\left(\mathbf{v}_{\mathbf{x,y}}\right)}_{\text{Volume sparsity 2}}$$

Learn how to solve this at CVPR! :-)

Results

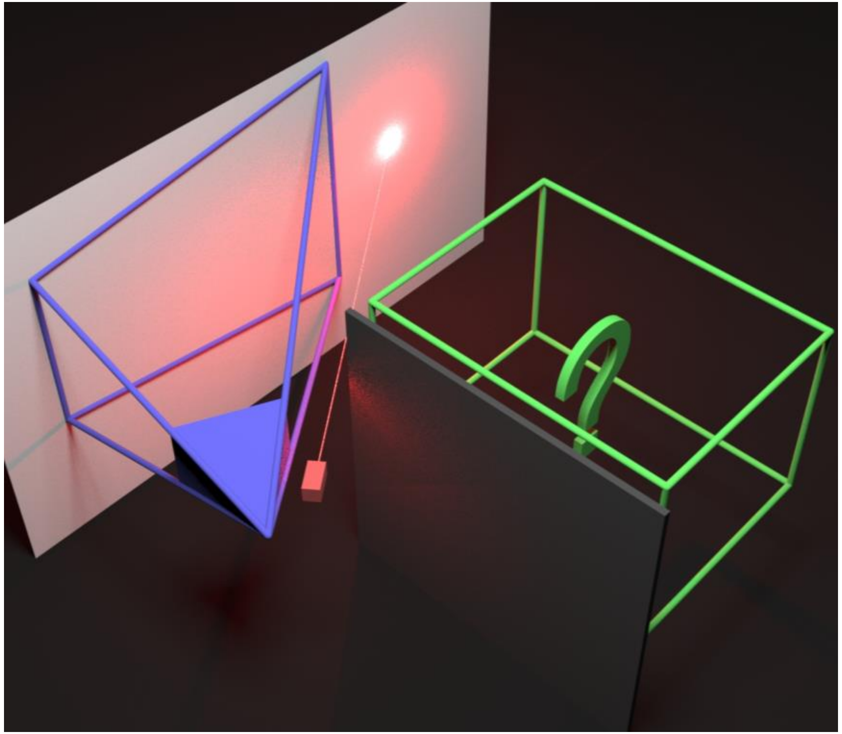

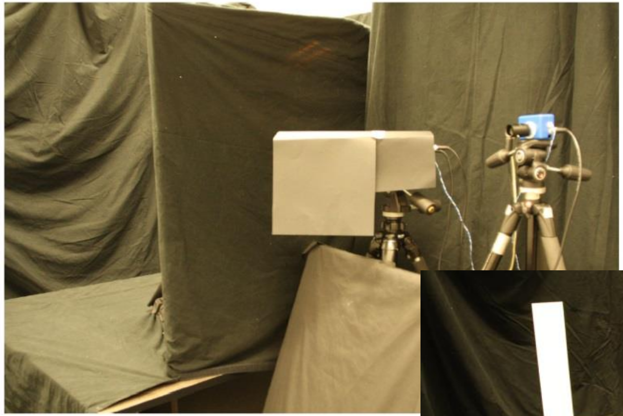

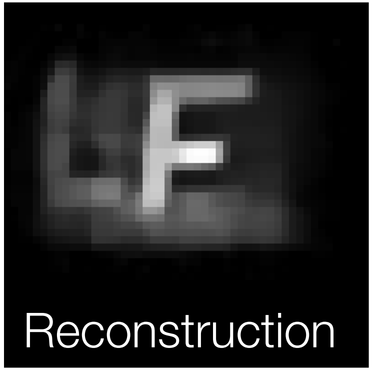

Effect of ambient light on geometry:

Room lights off:

All room lights on:

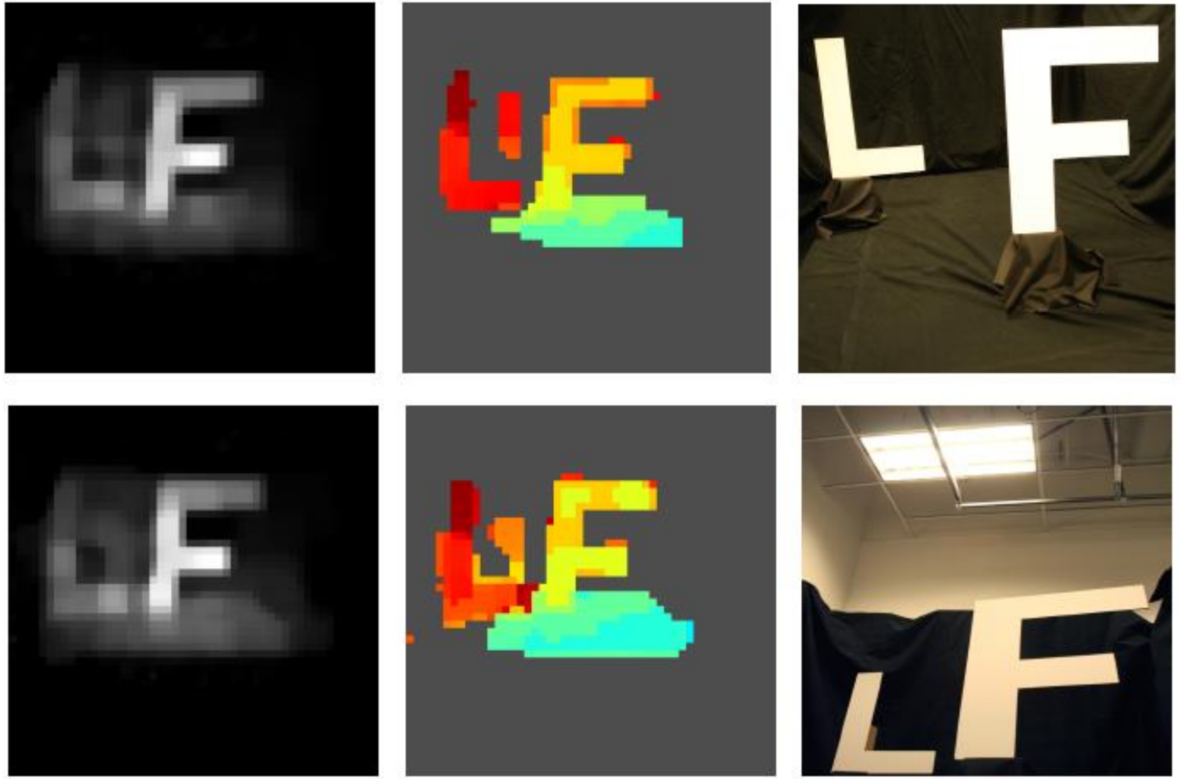

Albedo | Reconstructed depth | Unknown scene

Different materials:

Diffuse cardboard:

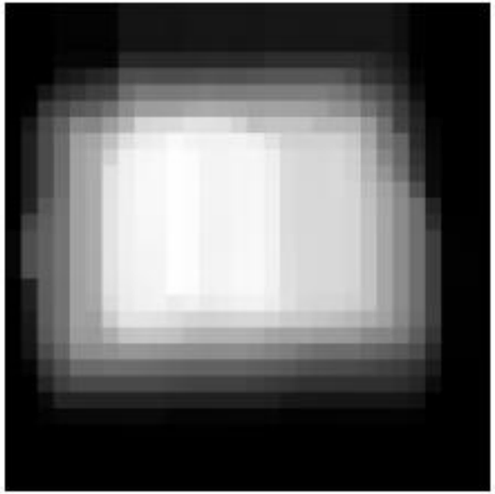

Albedo

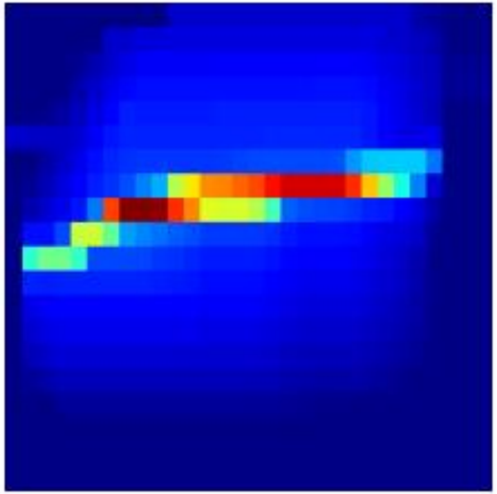

Slice through volume

Unknown scene

Different materials:

Polished whiteboard:

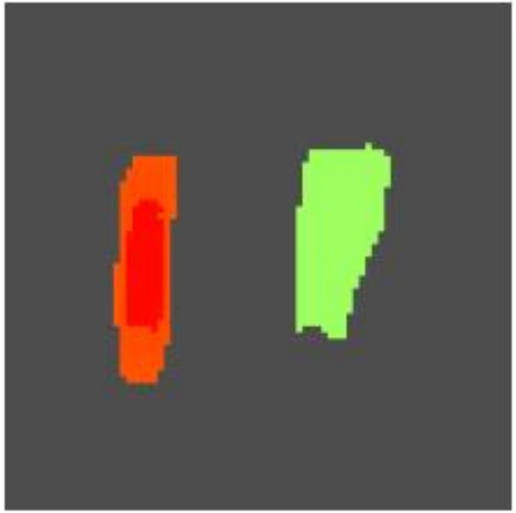

Depthmap

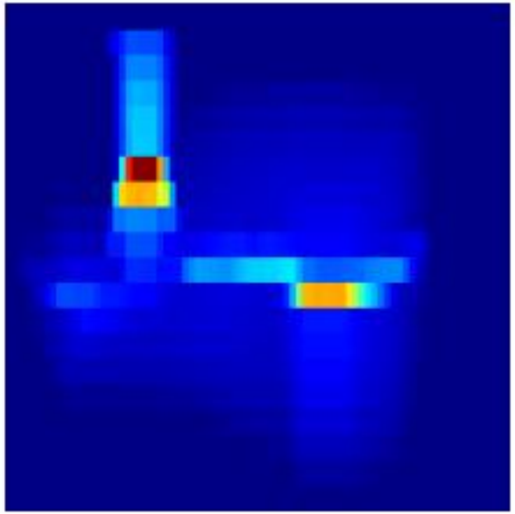

Slice through volume

Unknown scene

Different materials:

Mirrors:

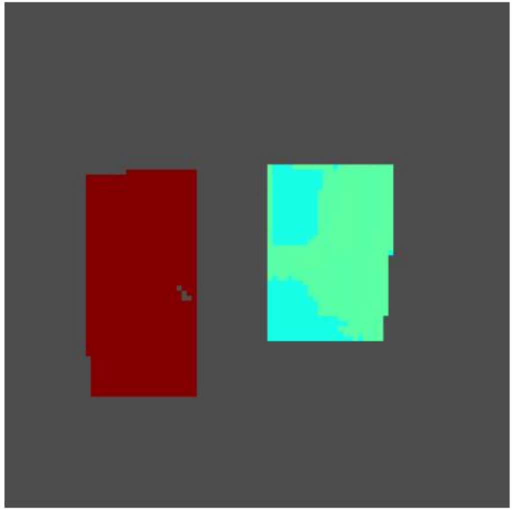

Depthmap

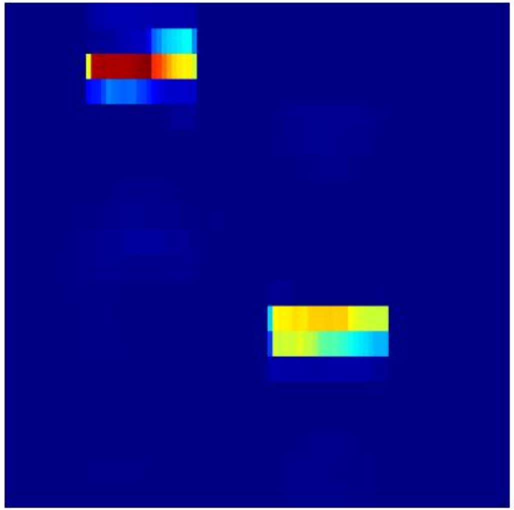

Slice through volume

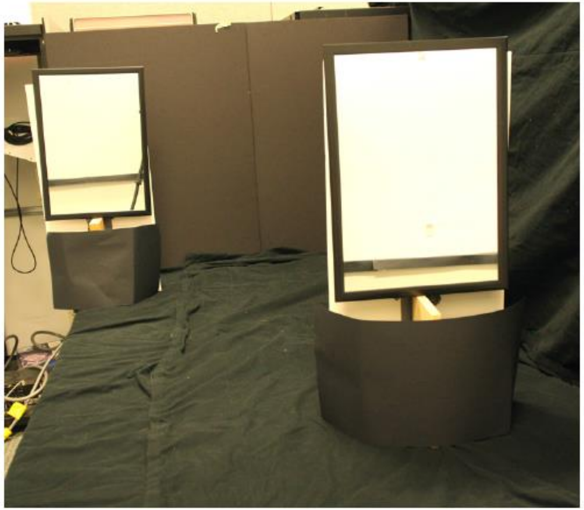

Unknown scene

Different materials:

Brushed metal and diffuse whiteboard:

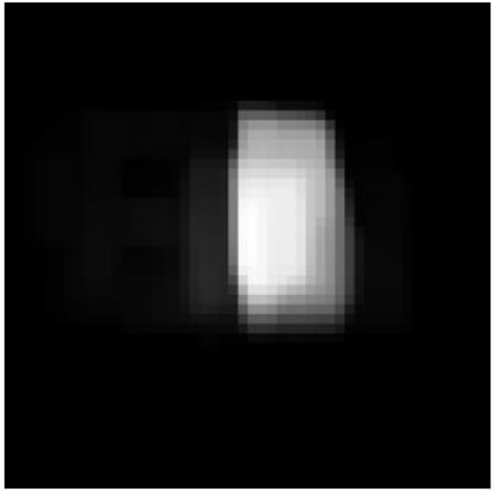

Depthmap

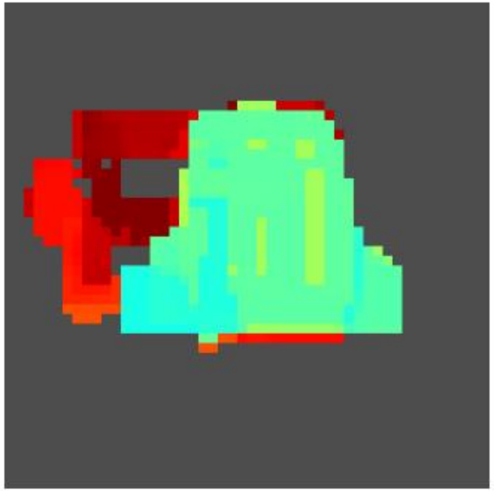

Slice through volume

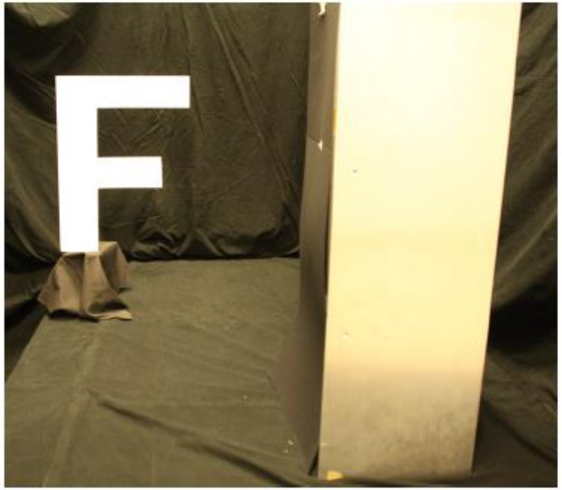

Unknown scene

Conclusion

- Lowered entry barrier for transient imaging
  - No more laser lab required

- ToF-sensor based setups can do things
  that used to require much more expensive gear
  - Looking around corners

Future work

- **Lots!** From theoretical foundations to new applications
- Alternative sensors – Microsoft, SoftKinetic?
- Your project ideas here – 7 imager systems available

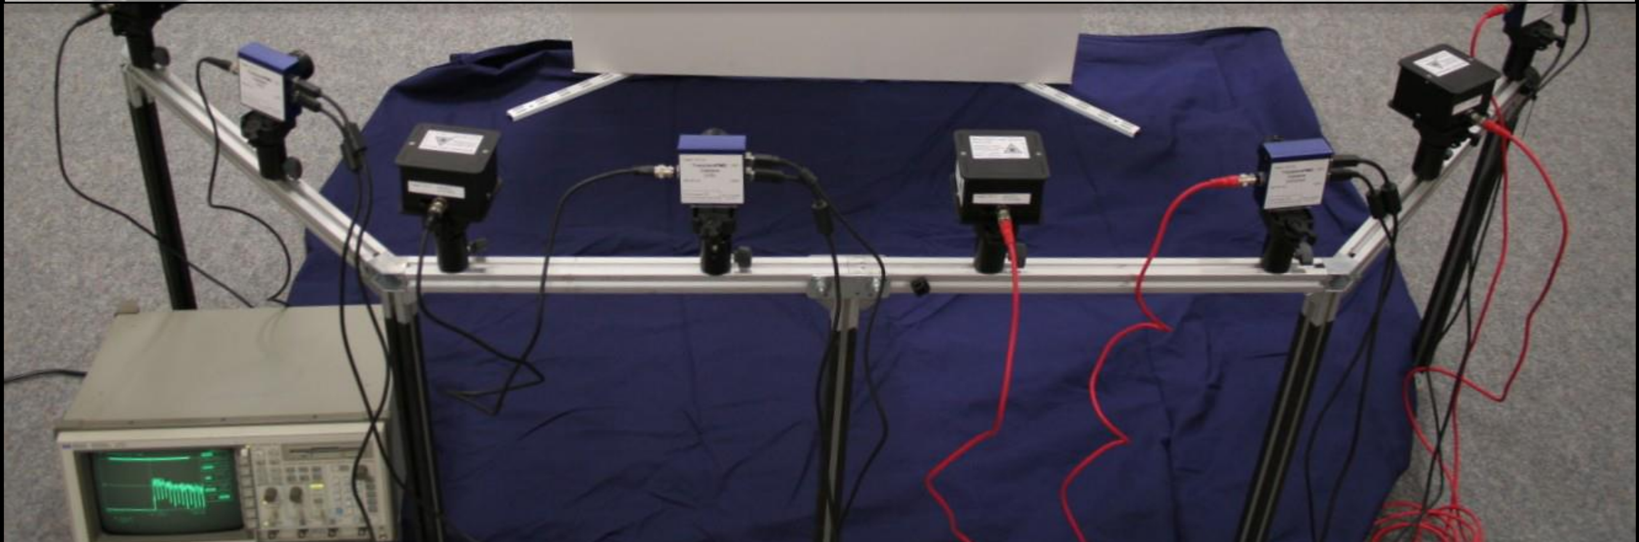

Thanks

- People:
  Felix Heide, Lei Xiao, James Gregson, Wolfgang Heidrich, Hendrik Lensch, Johannes Hanika, Hans-Peter Seidel

- [Heide14] supported through Max Planck Center for Visual Computing and Communication