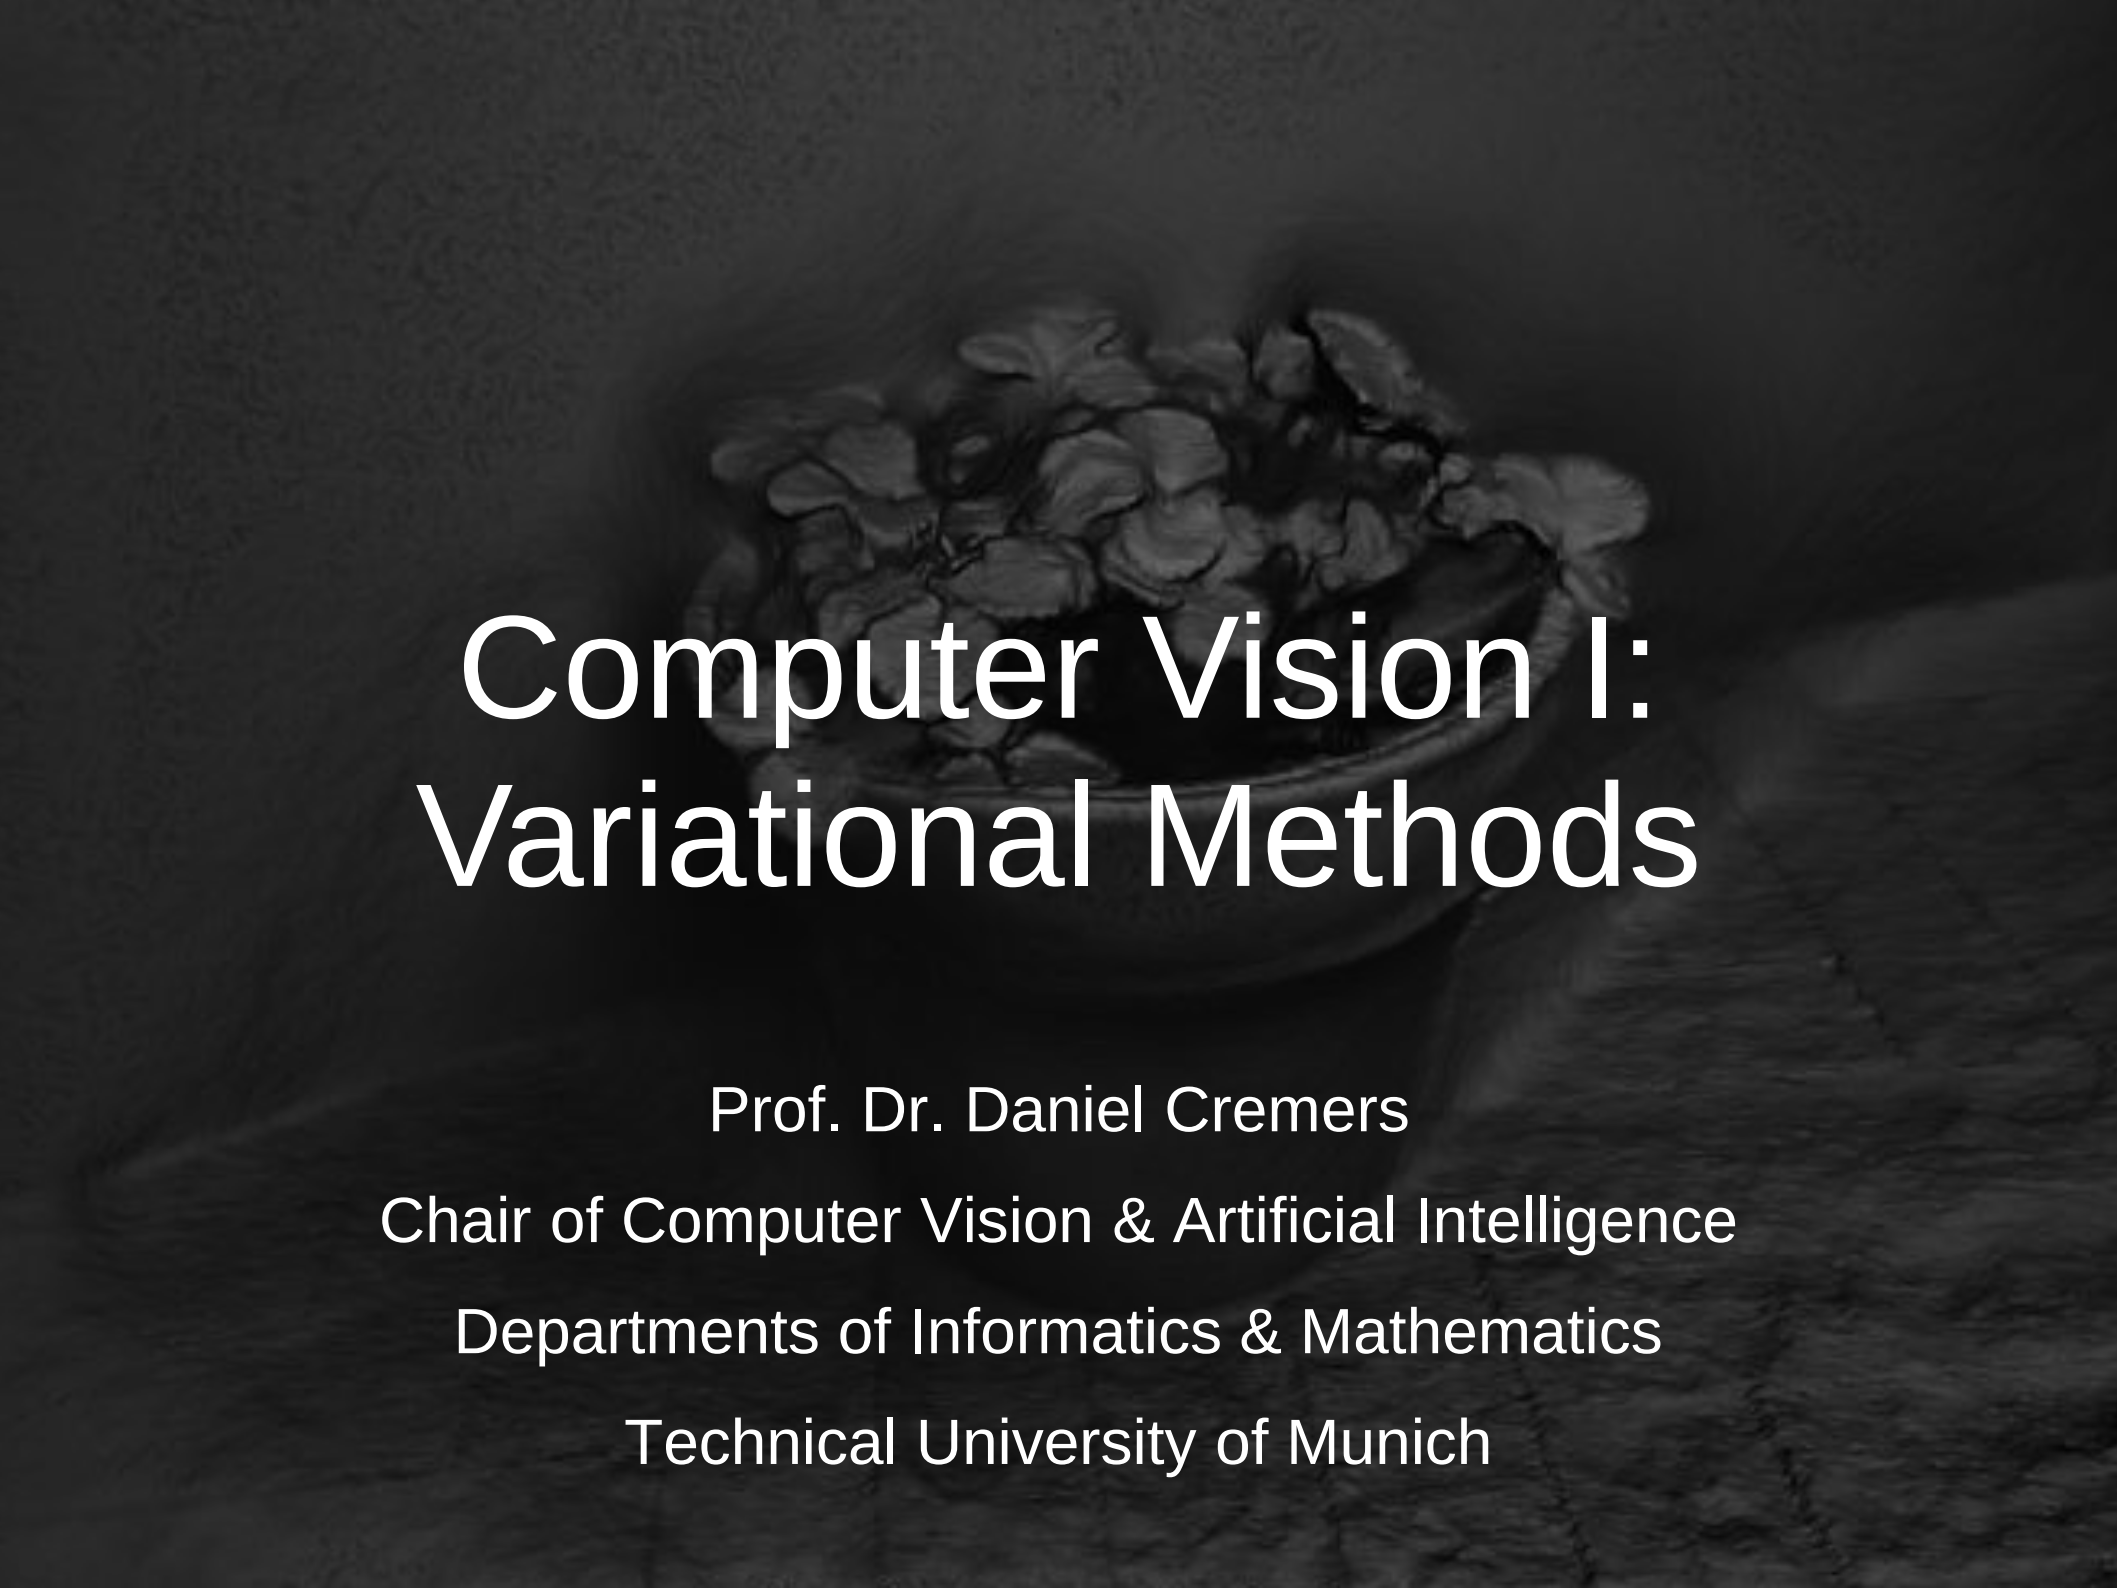

Computer Vision I: Variational Methods

Prof. Dr. Daniel Cremers

Chair of Computer Vision & Artificial Intelligence

Departments of Informatics & Mathematics

Technical University of Munich

Exercises

Marvin Eisenberger

Emanuel Laude


Original left image

Variational Scene Flow

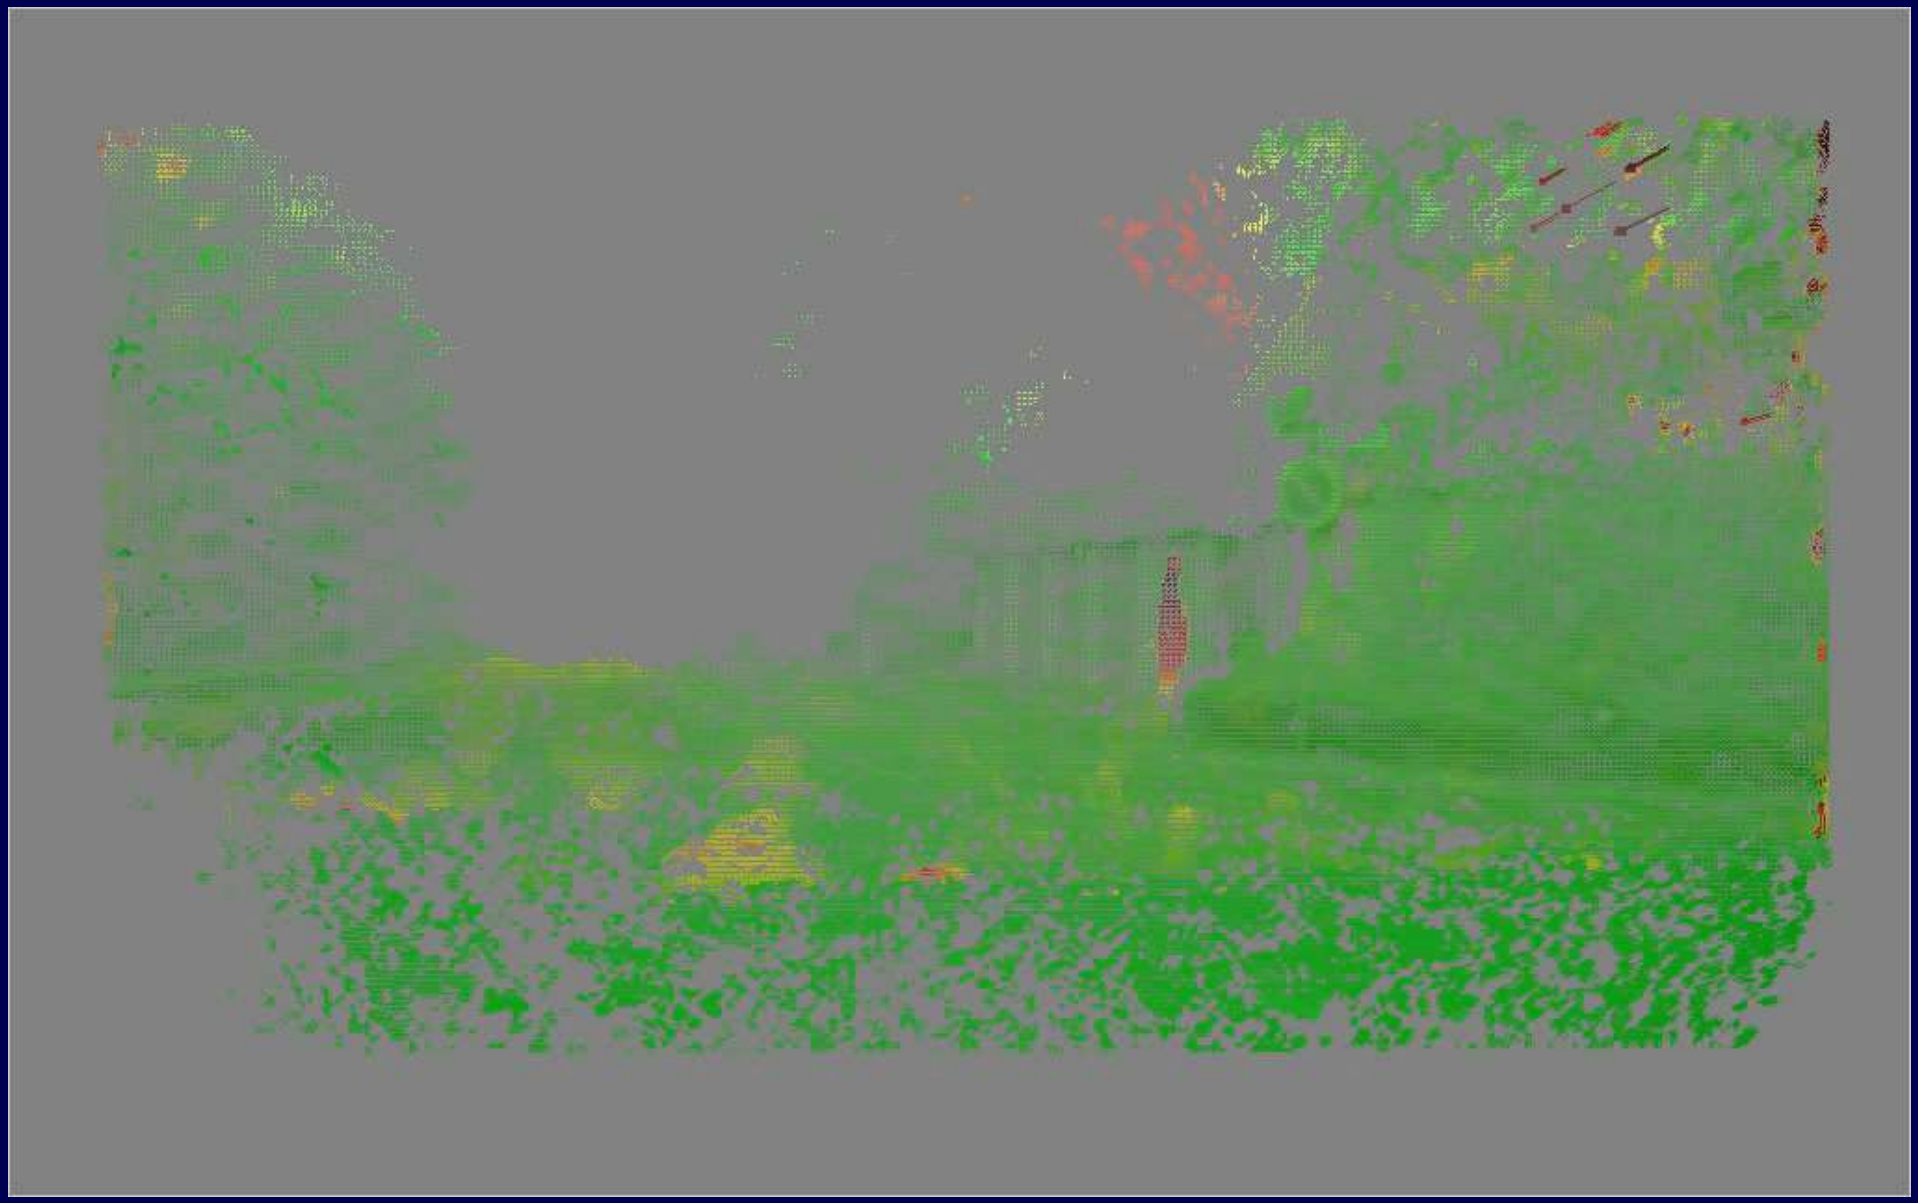

Wedel et al. IJCV ‘11, Wedel & Cremers, Springer 2011

ECCV 2018

European Conference on Computer Vision

8 – 14 September 2018 | Munich, Germany

TUM Chair of Computer Vision & AI

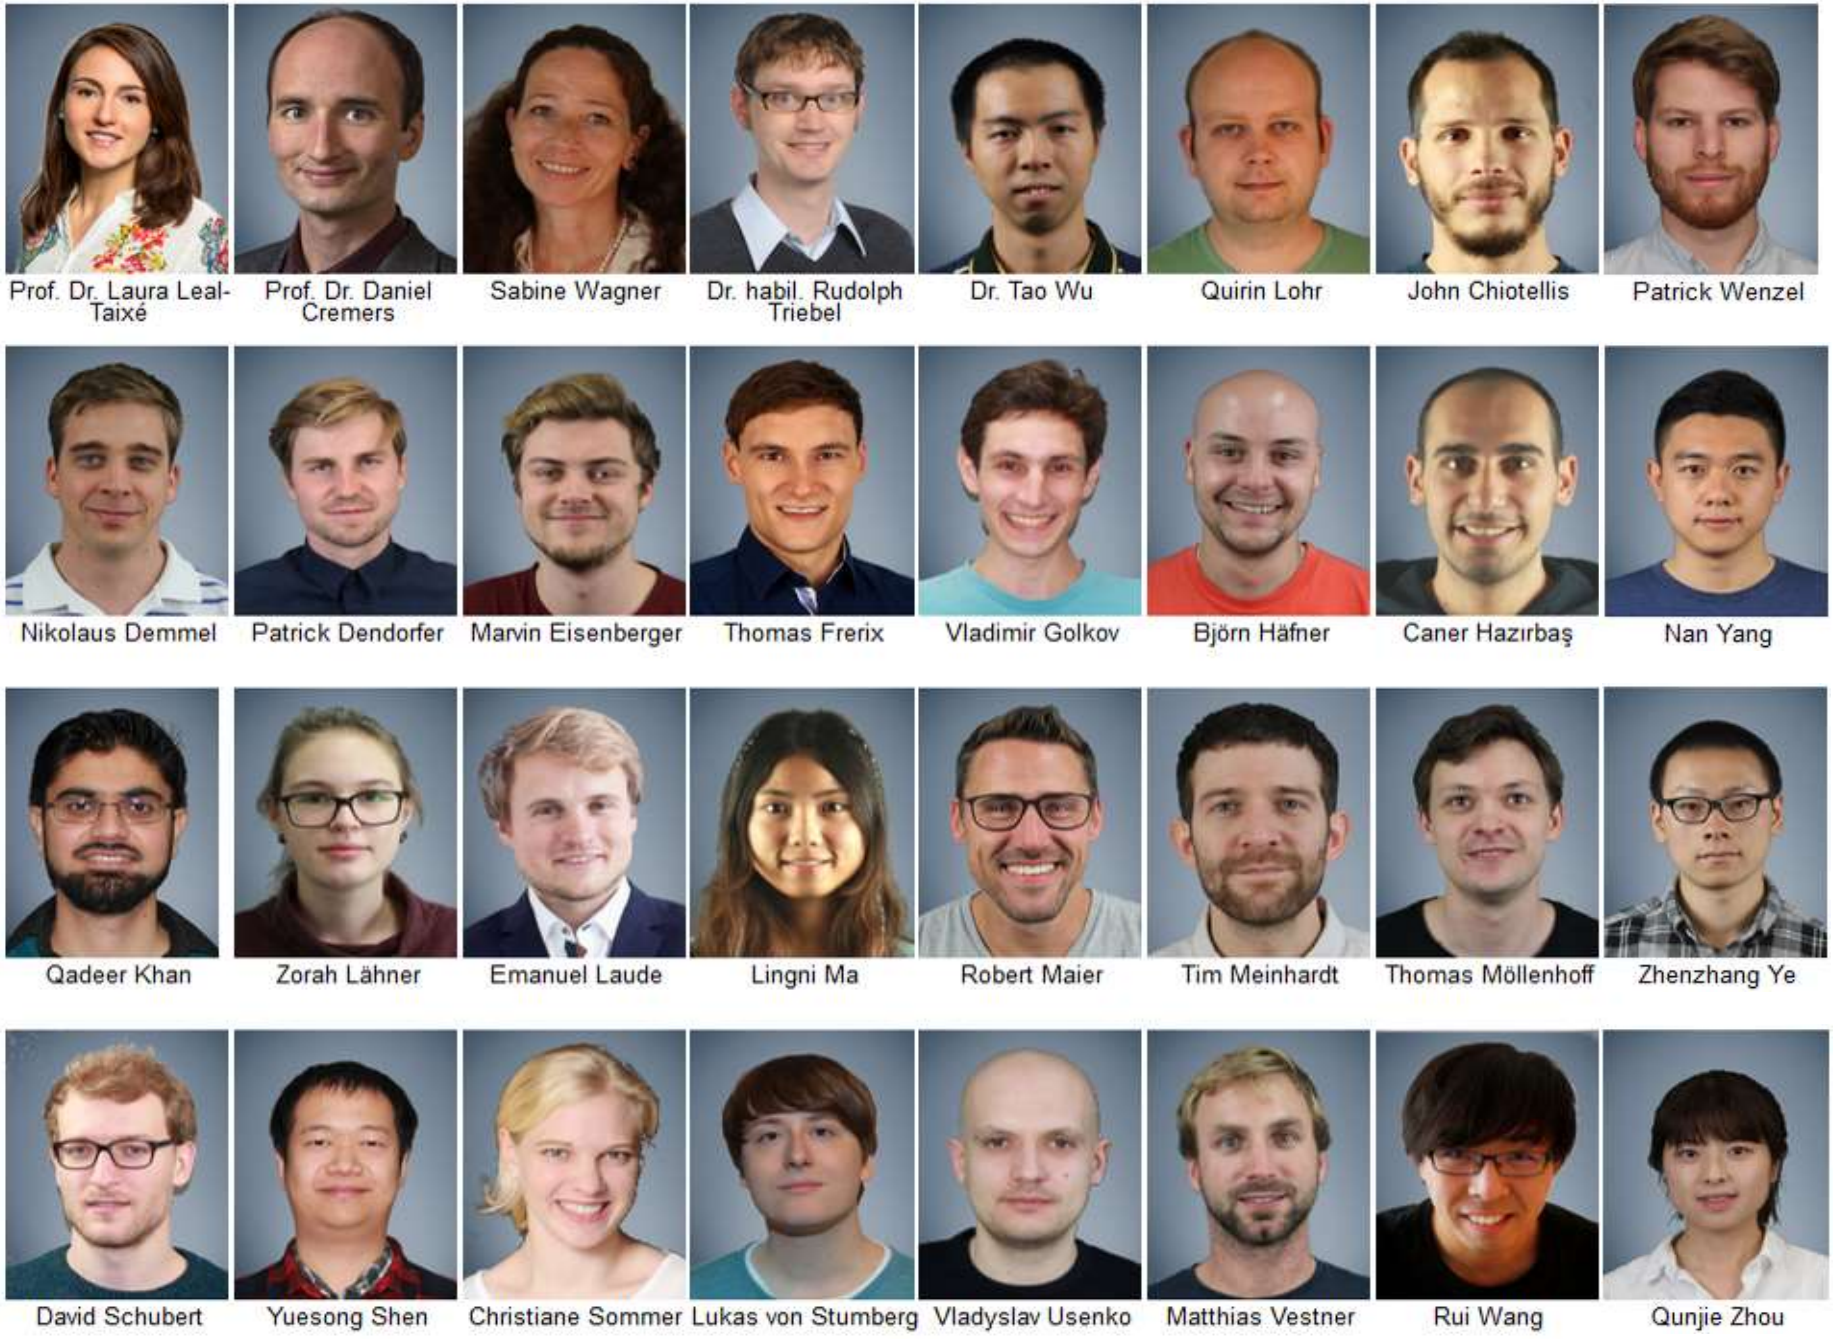

Spatially Dense 3D Reconstruction

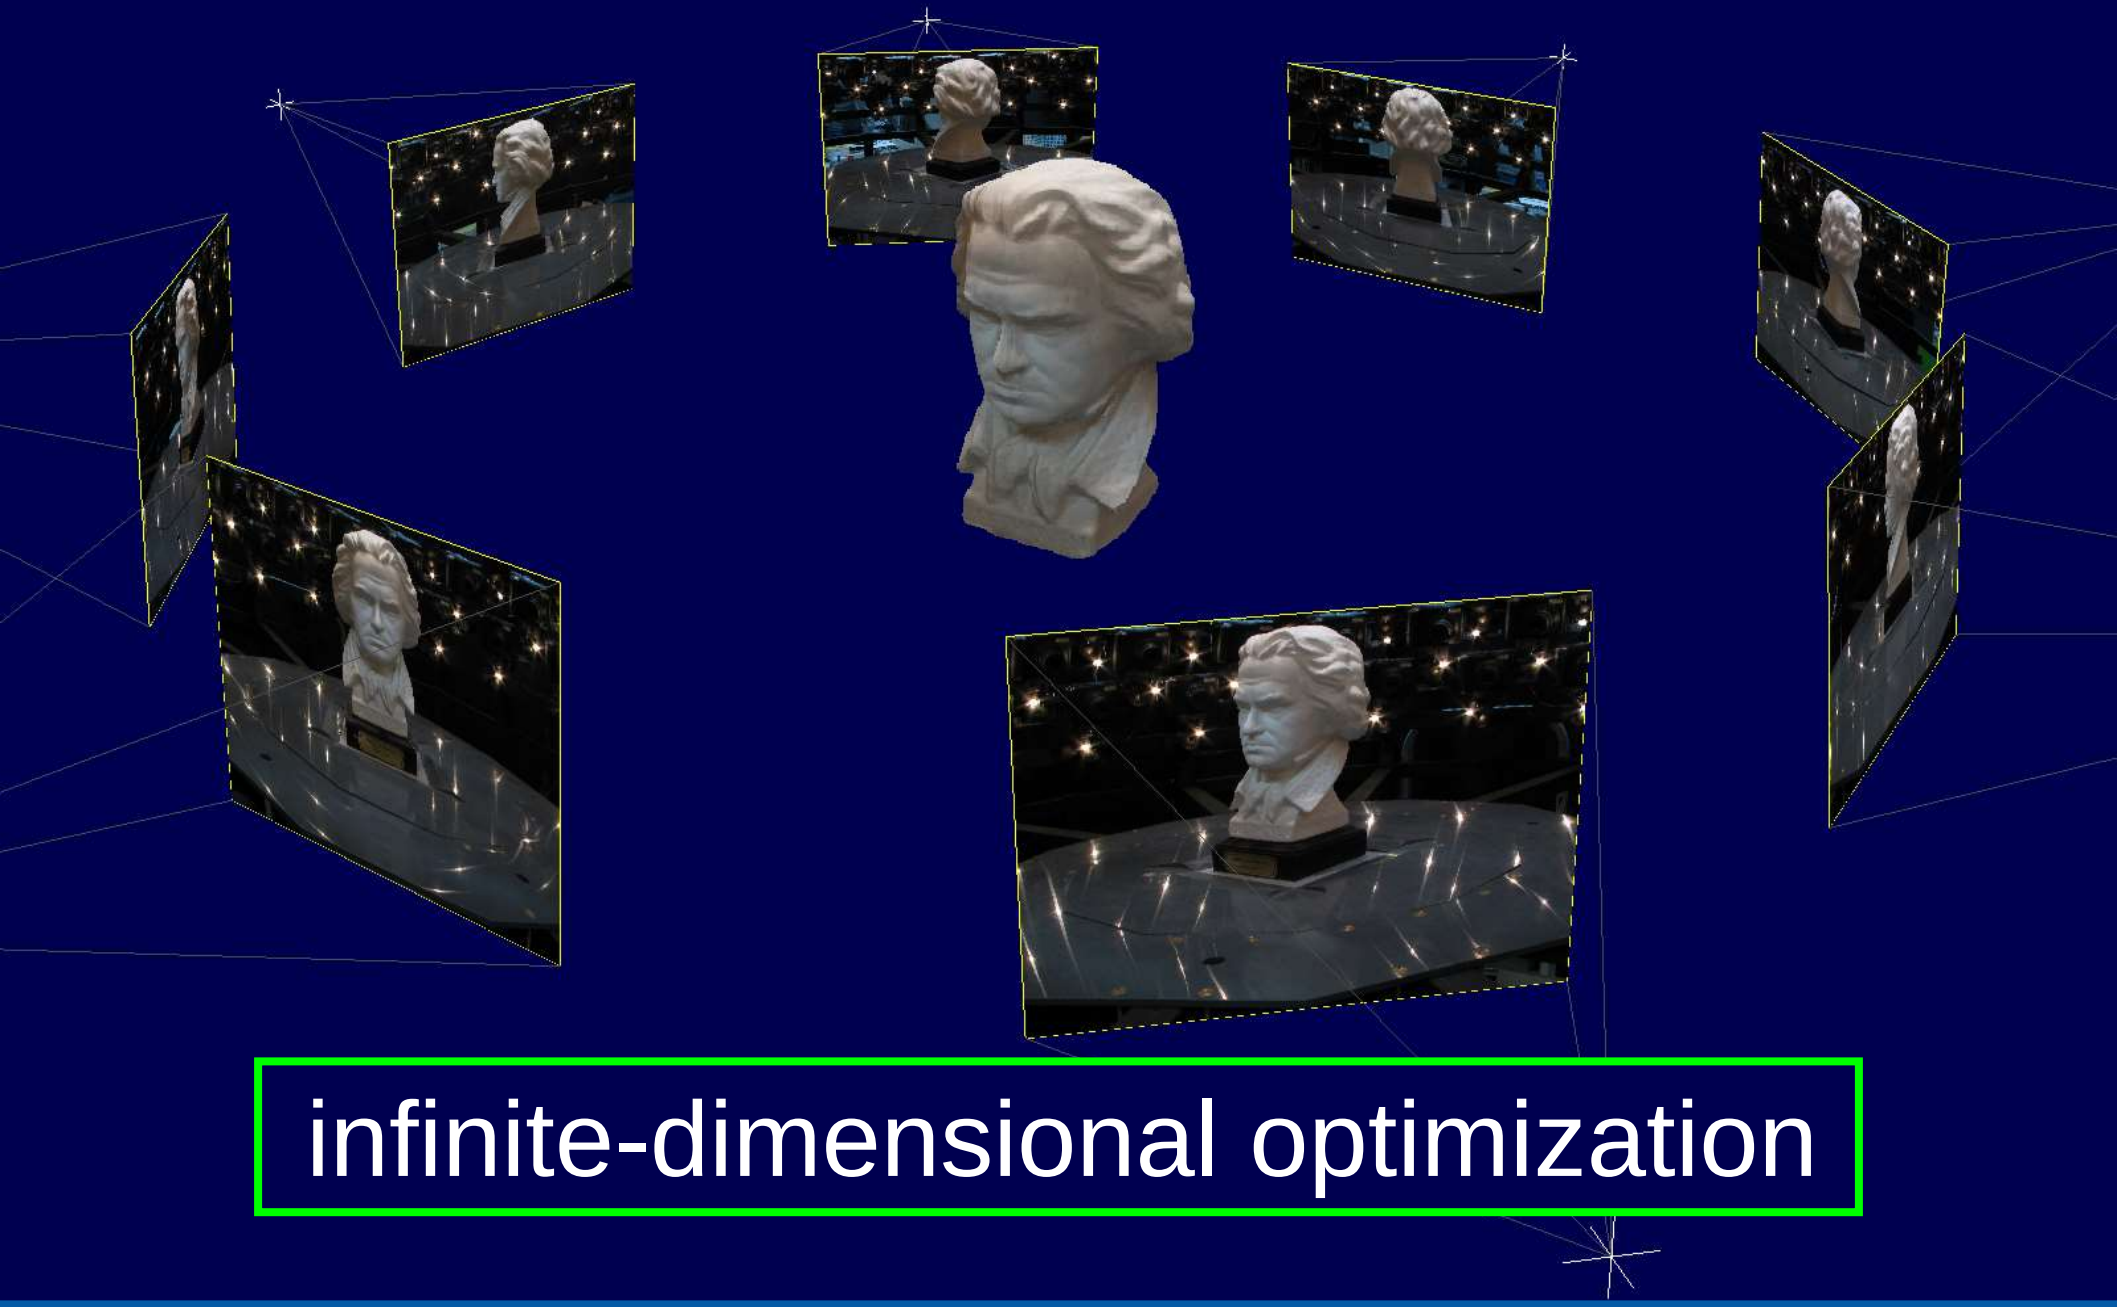

infinite-dimensional optimization

Which path is the fastest?

Bernoulli & The Brachistochrone

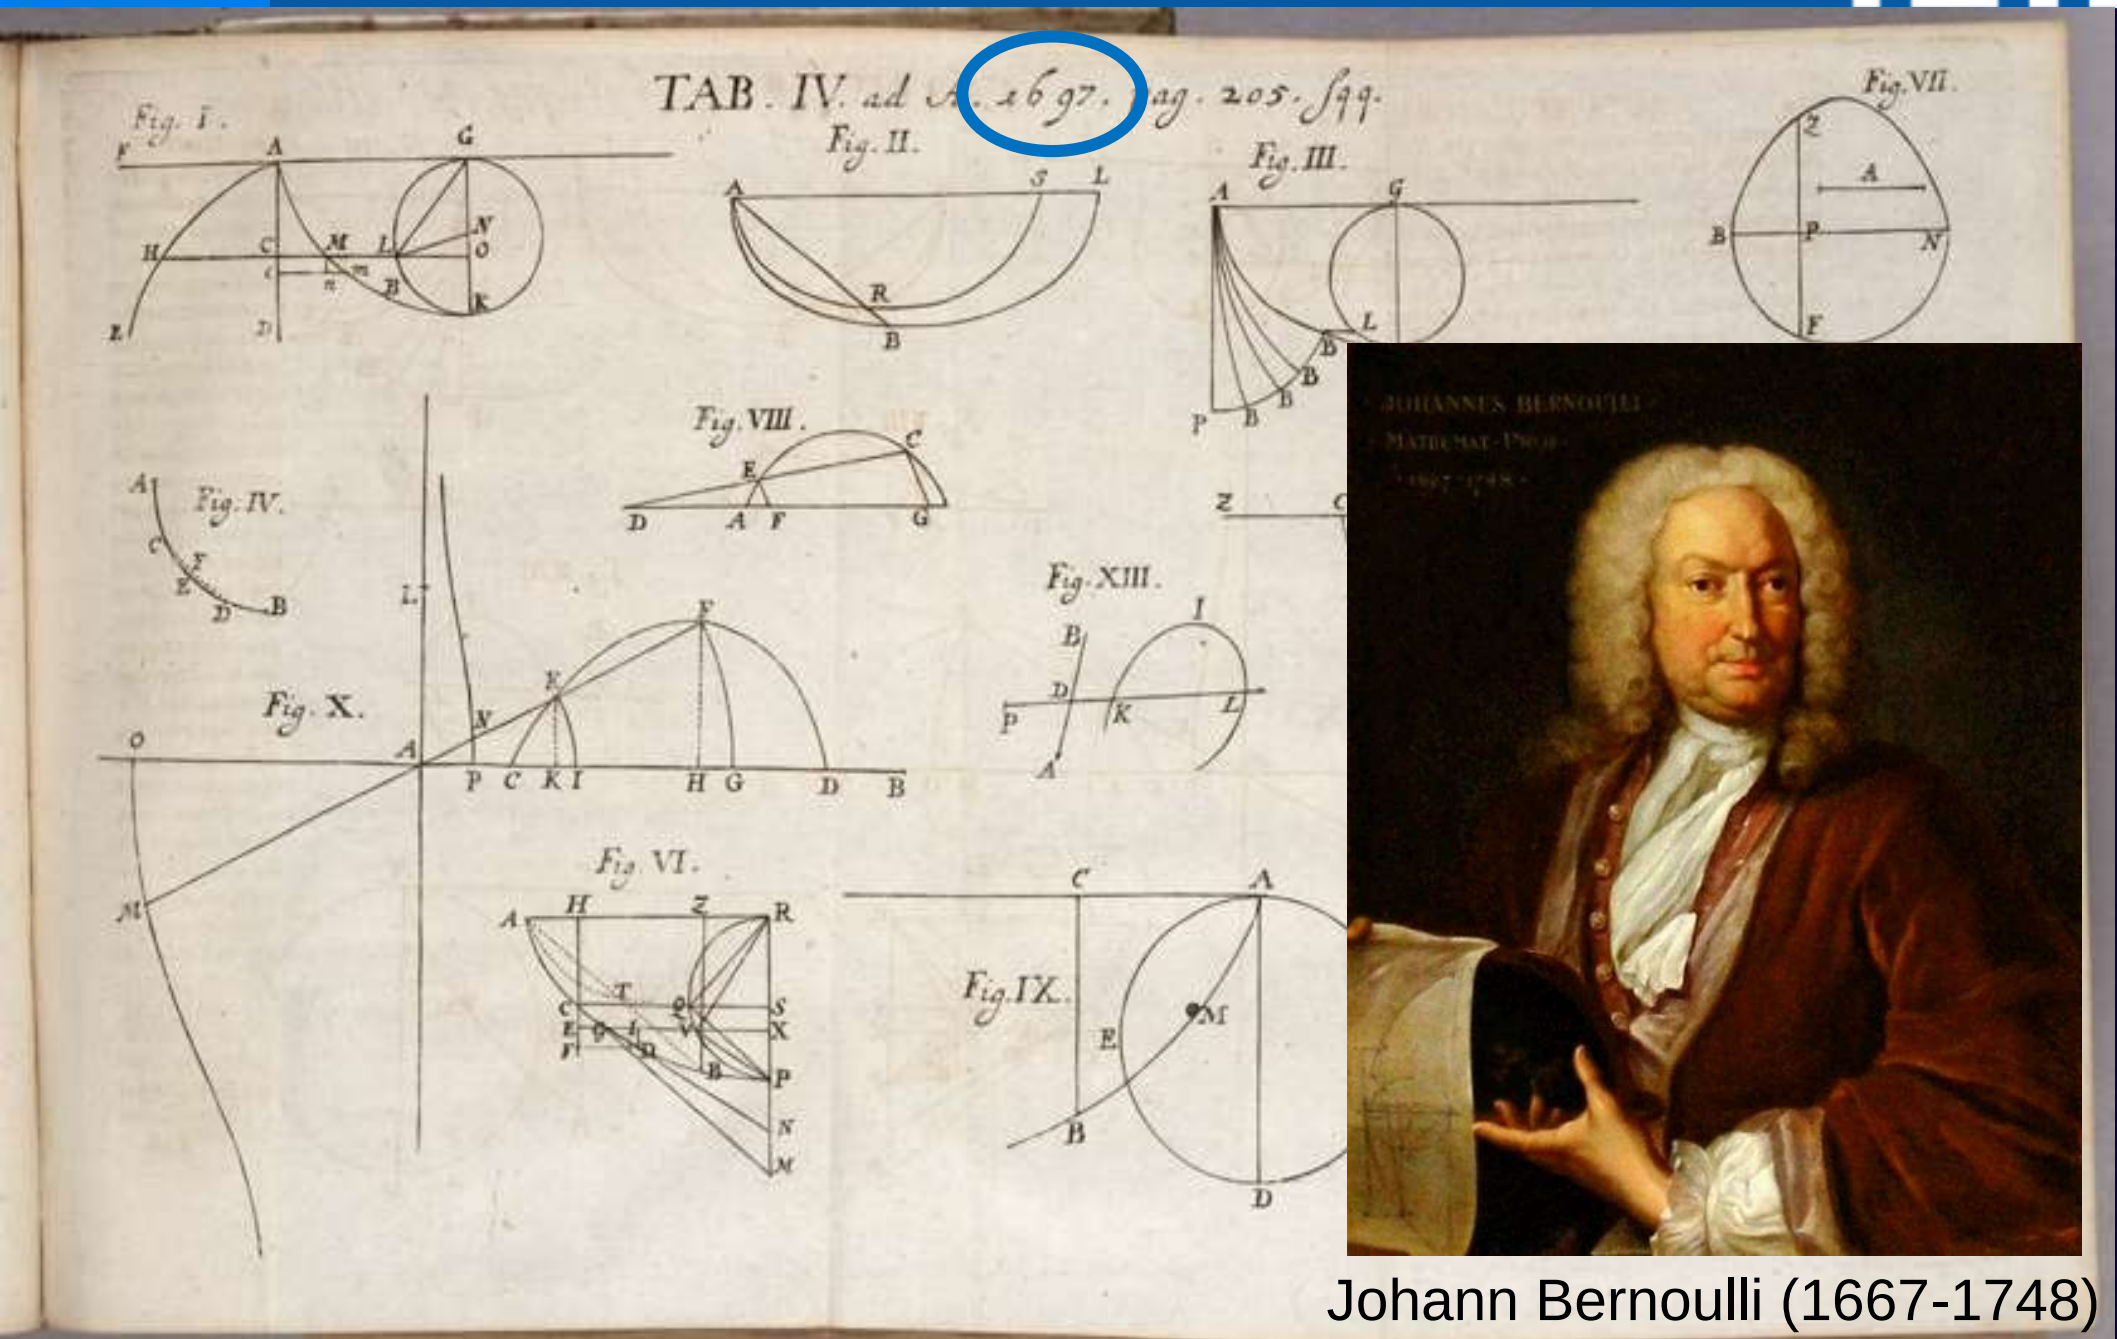

Johann Bernoulli (1667-1748)

Optimization in Computer Vision

Image segmentation:

Geman, Geman ’84, Blake, Zisserman ‘87, Kass et al. ’88, Mumford, Shah ’89, Caselles et al. ‘95, Kichenassamy et al. ‘95, Paragios, Deriche ’99, Chan, Vese ‘01, Tsai et al. ‘01, …

Multiview stereo reconstruction:

Faugeras, Keriven ‘98, Duan et al. ‘04, Yezzi, Soatto ‘03, Seitz et al. ‘06, Hernandez et al. ‘07, Labatut et al. ’07, …

Optical flow estimation:

Horn, Schunck ‘81, Nagel, Enkelmann ‘86, Black, Anandan ‘93, Alvarez et al. ‘99, Brox et al. ‘04, Baker et al. ‘07, Zach et al. ‘07, Sun et al. ‘08, Wedel et al. ’09, …

Non-convex energies

Optimization and Convexity

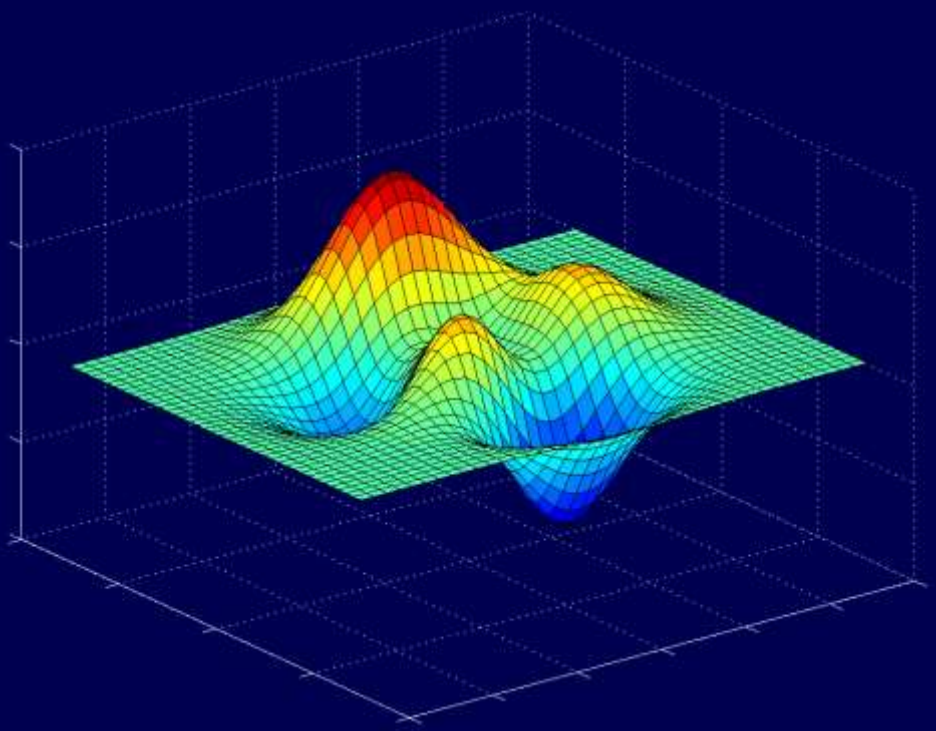

Non-convex energy

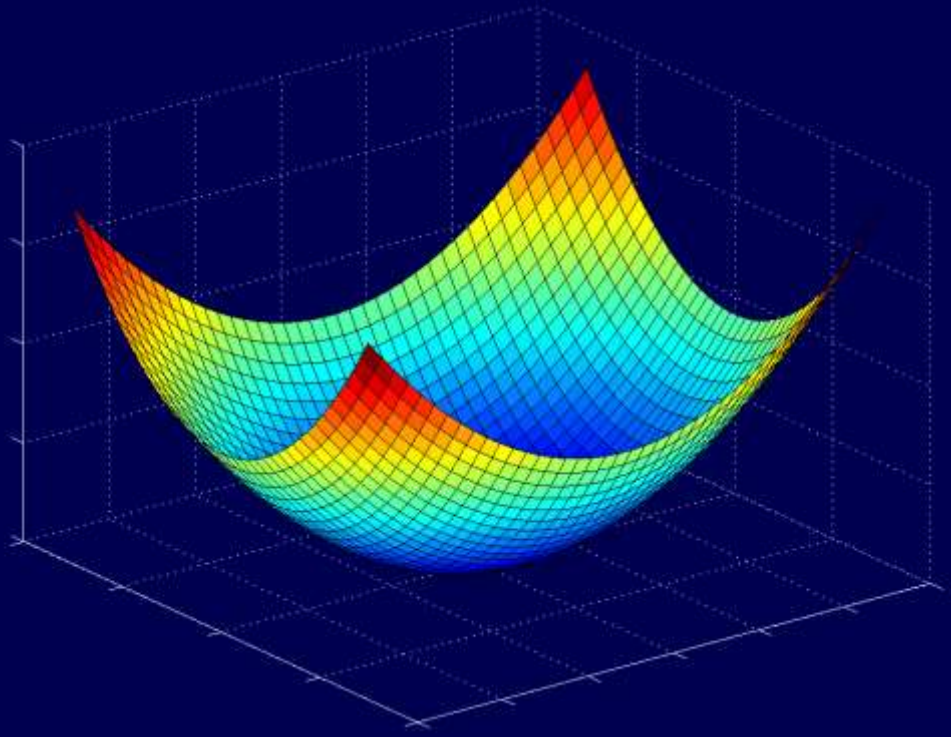

Convex energy

Inverse Problems: Denoising

image $f : \Omega \to \mathbb{R}^3, \quad \Omega \subset \mathbb{R}^2$

denoised image $u_{den} : \Omega \to \mathbb{R}^3$

$$u_{den} = \arg\min_u \int_\Omega (u - f)^2 \, dx + \lambda \int_\Omega |\nabla u| \, dx$$

Rudin, Osher, Fatemi 1992, Goldlücke, Strekalovskiy, Cremers 2012

Variational Methods

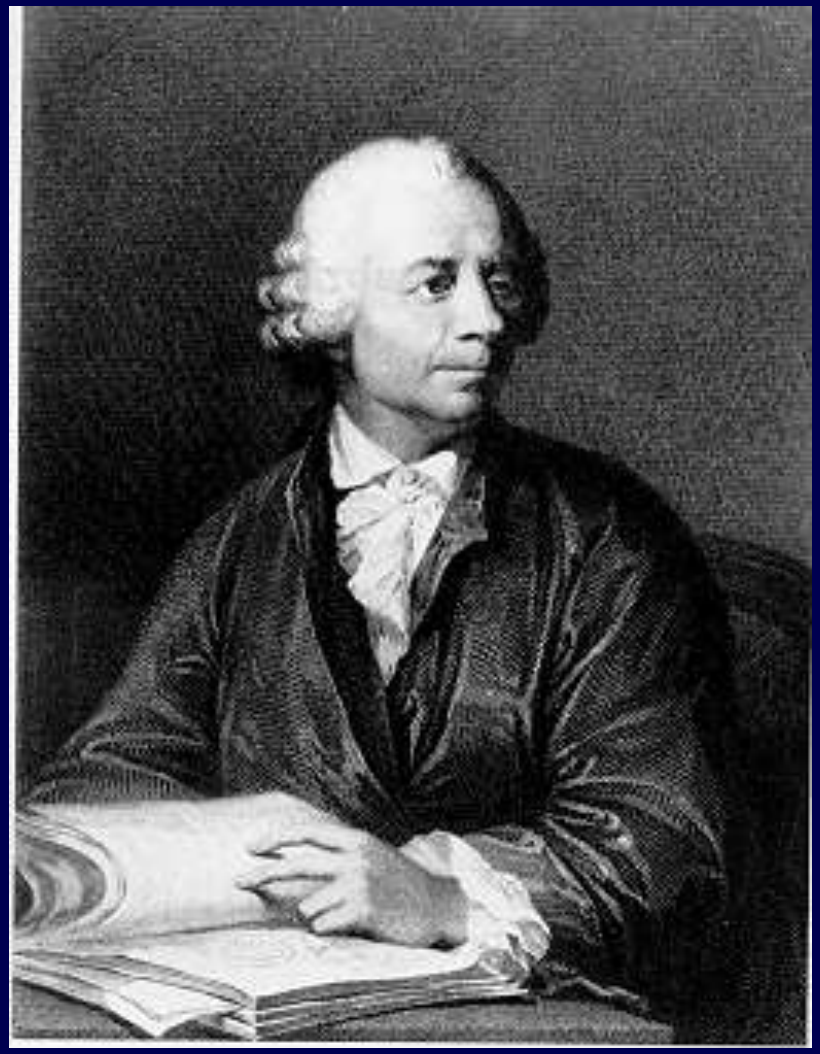

Leonhard Euler
(1703-1783)

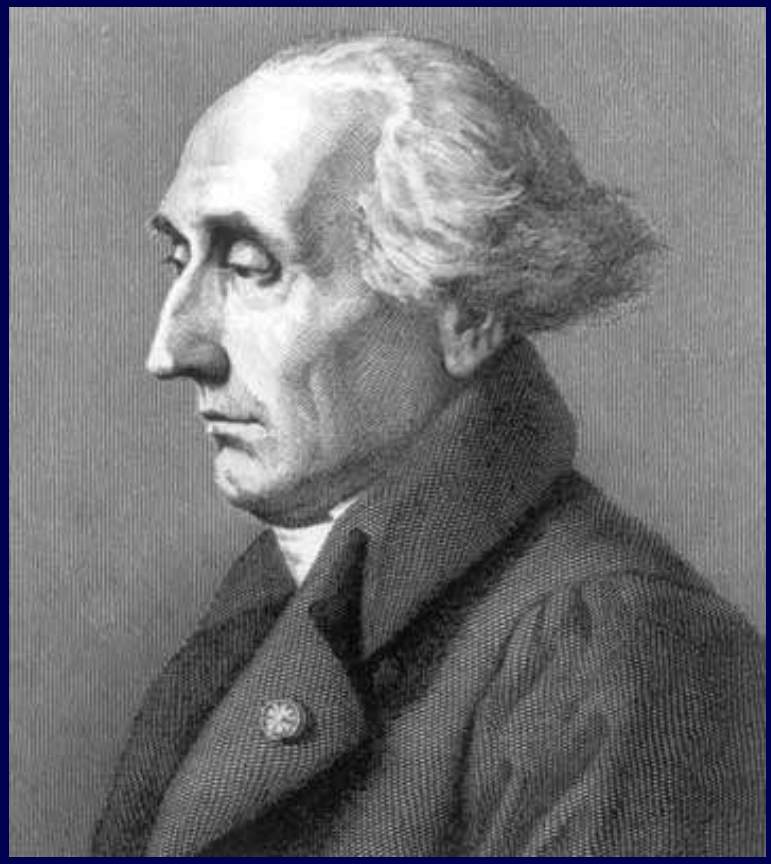

Joseph-Louis Lagrange
(1736 – 1813)

Variational Methods & PDEs

$$E(u) = \int_\Omega (u - f)^2 + \lambda \, |\nabla u| \, dx = \int_\Omega \mathcal{L}(u, u') \, dx$$

Euler-Lagrange equation as necessary condition:

$$\frac{dE}{du} = \frac{\partial \mathcal{L}}{\partial u} - \frac{d}{dx}\left(\frac{\partial \mathcal{L}}{\partial u'}\right) = 0$$

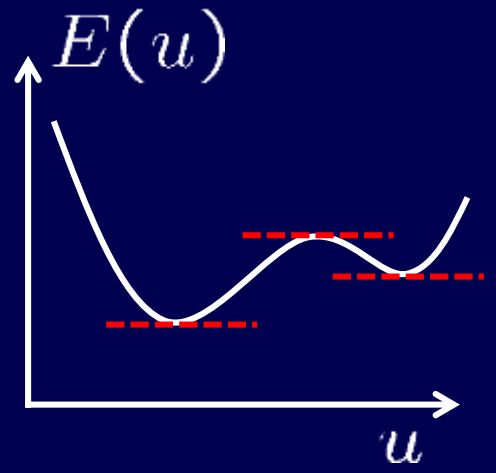

Gradient descent:

$$\frac{\partial u}{\partial t} = -\frac{dE}{du}$$

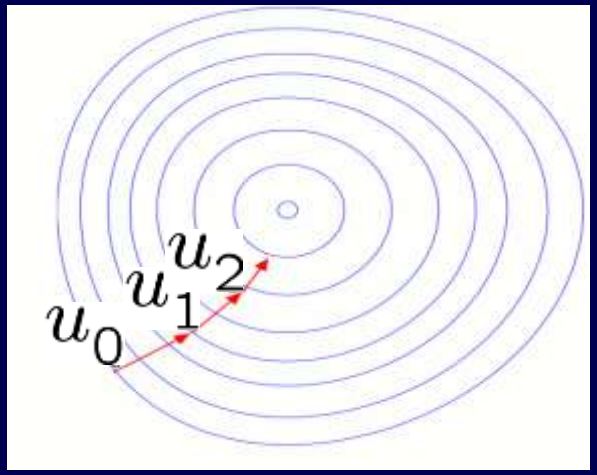

Overview

Multiview reconstruction

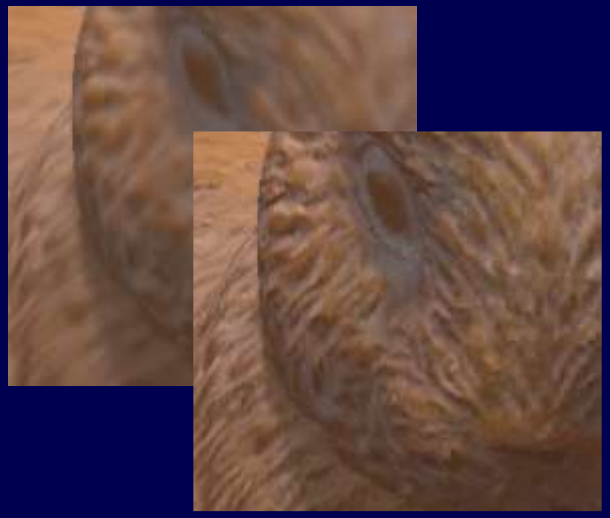

Super-res.textures

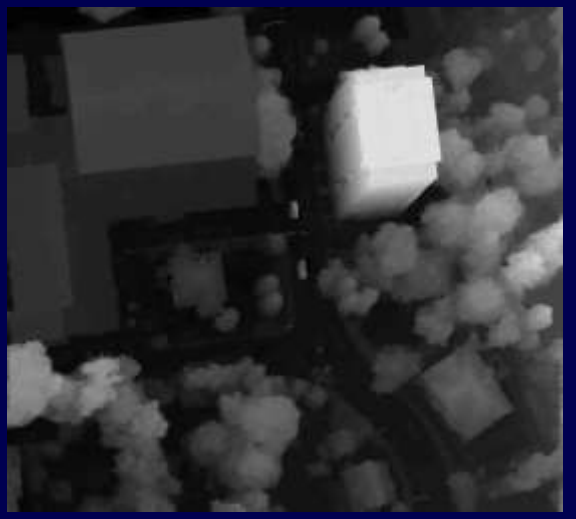

Stereo reconstruction

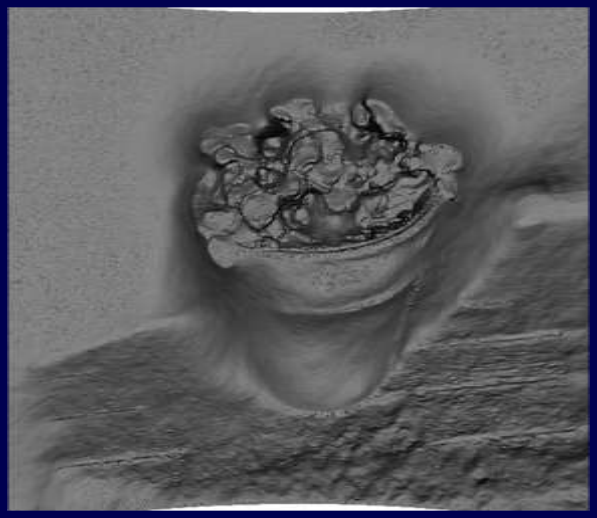

Realtime dense geometry

RGB-D cameras

Reconstruction on the fly

Overview

Multiview reconstruction

Super-res.textures

Stereo reconstruction

Realtime dense geometry

RGB-D cameras

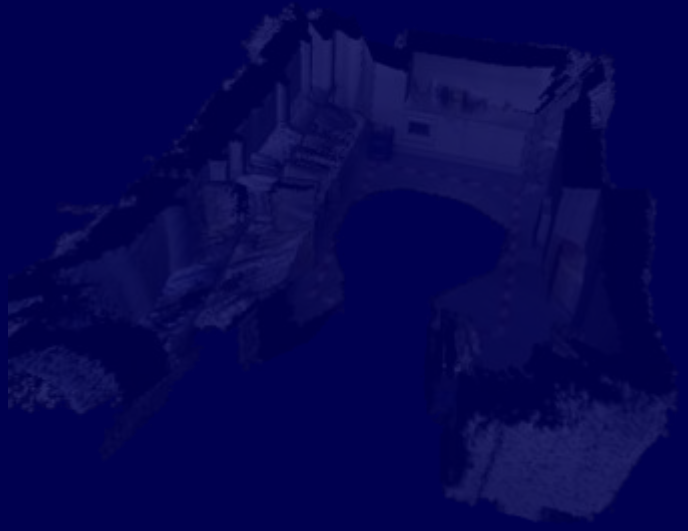

Reconstruction on the fly

3D Reconstruction from Multiple Views

Kolev, Klodt, Brox, Cremers, Int. J. of Computer Vision '09:

Theorem: Globally optimal surfaces can be computed by convex optimization.

Evolution to Global Optimum

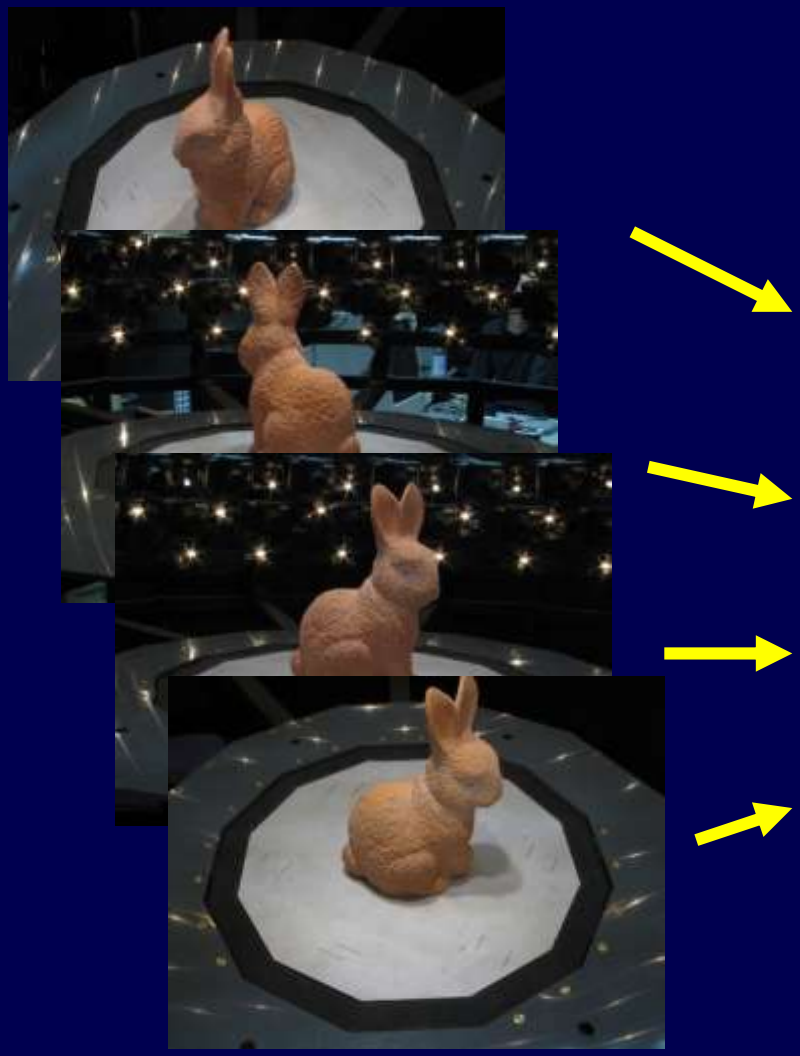

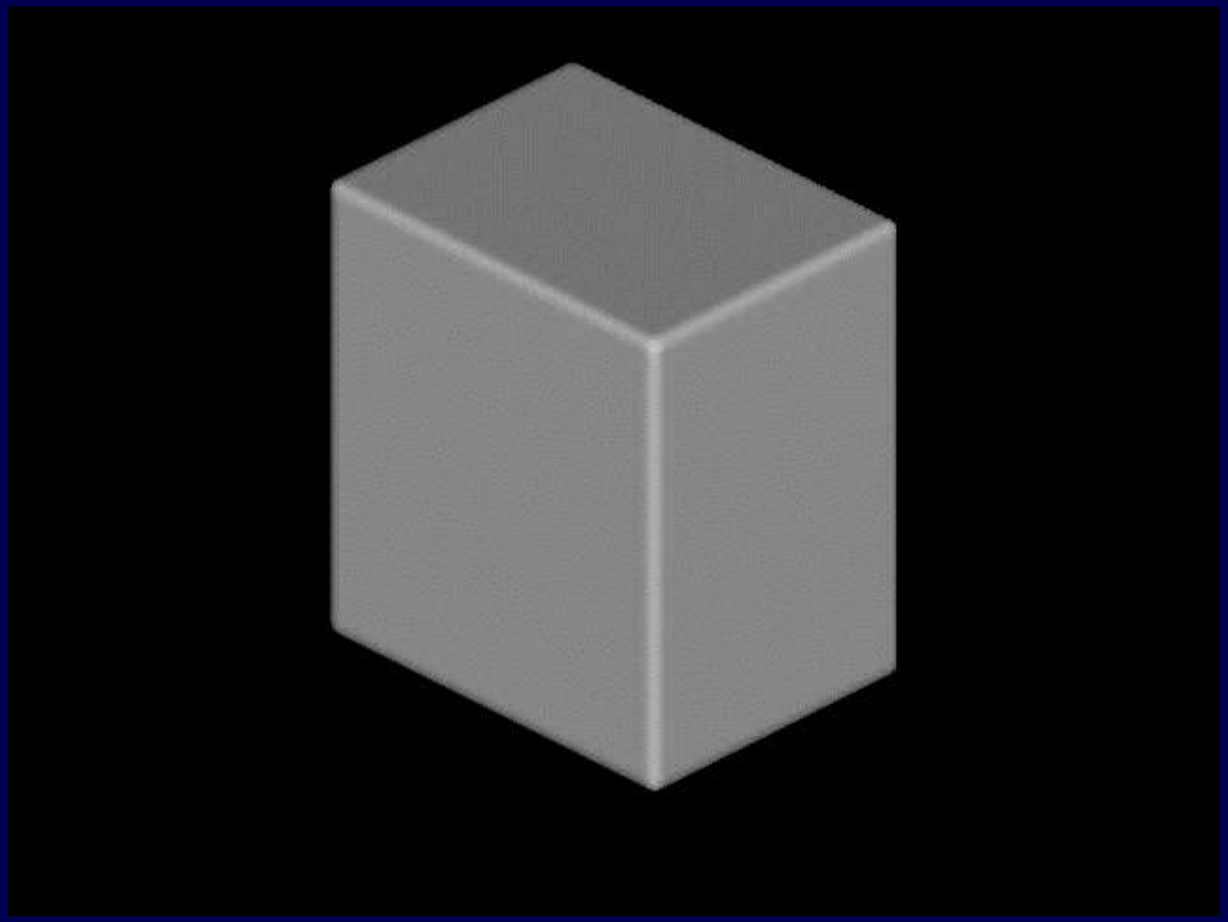

Kolev, Klodt, Brox, Cremers, Int. J. of Computer Vision '09:

Theorem: Globally optimal surfaces can be computed by convex optimization.

Reconstruction of Fine-scale Structures





Image data courtesy of Yasutaka Furukawa.

Reconstructing the Niobids Statues

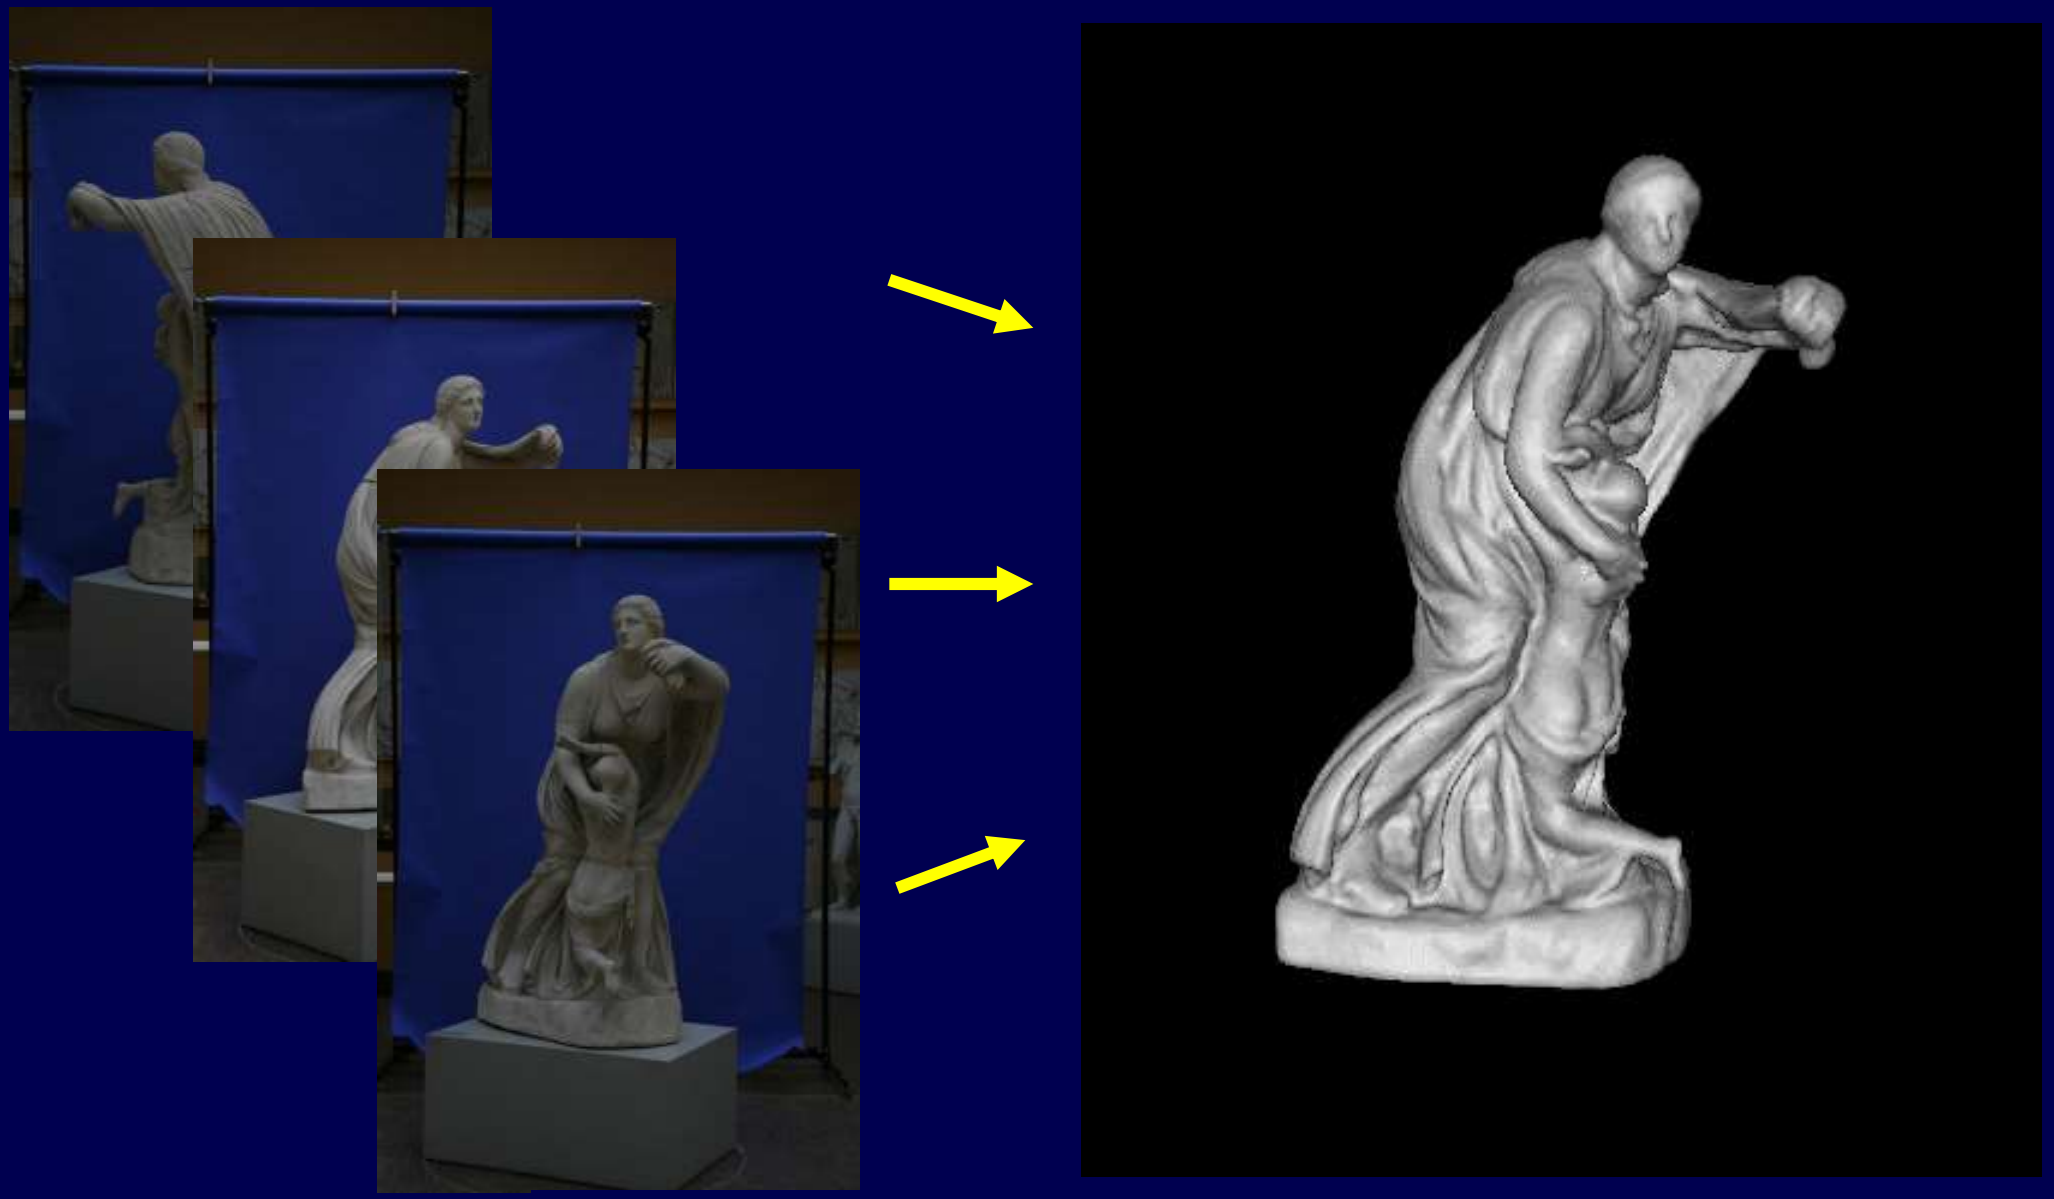

Kolev, Cremers, ECCV '08, PAMI '11

Multiview Reconstruction

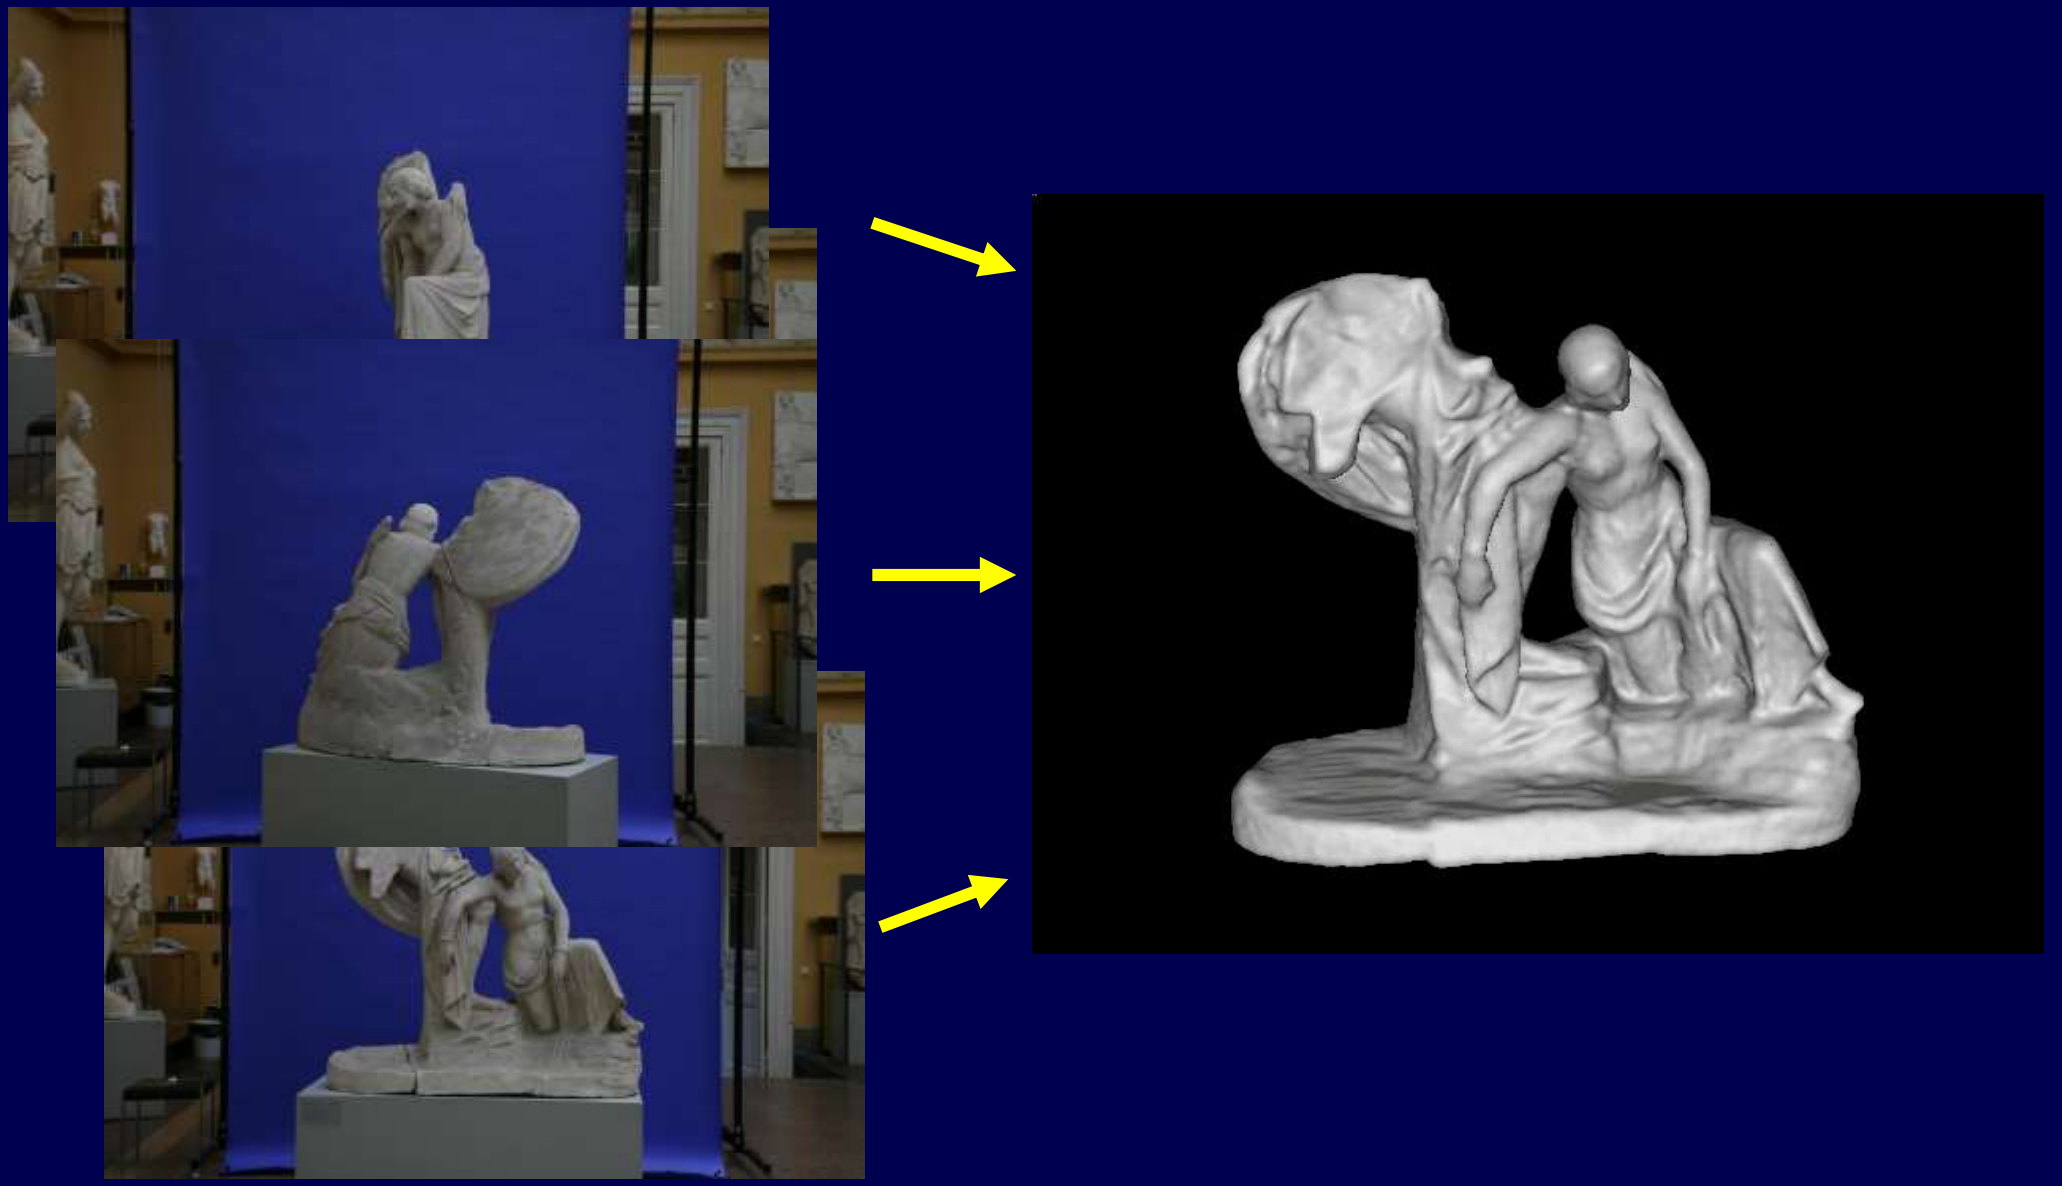

Kolev, Cremers, ECCV '08, PAMI '11

Reconstructing Dynamic Scenes

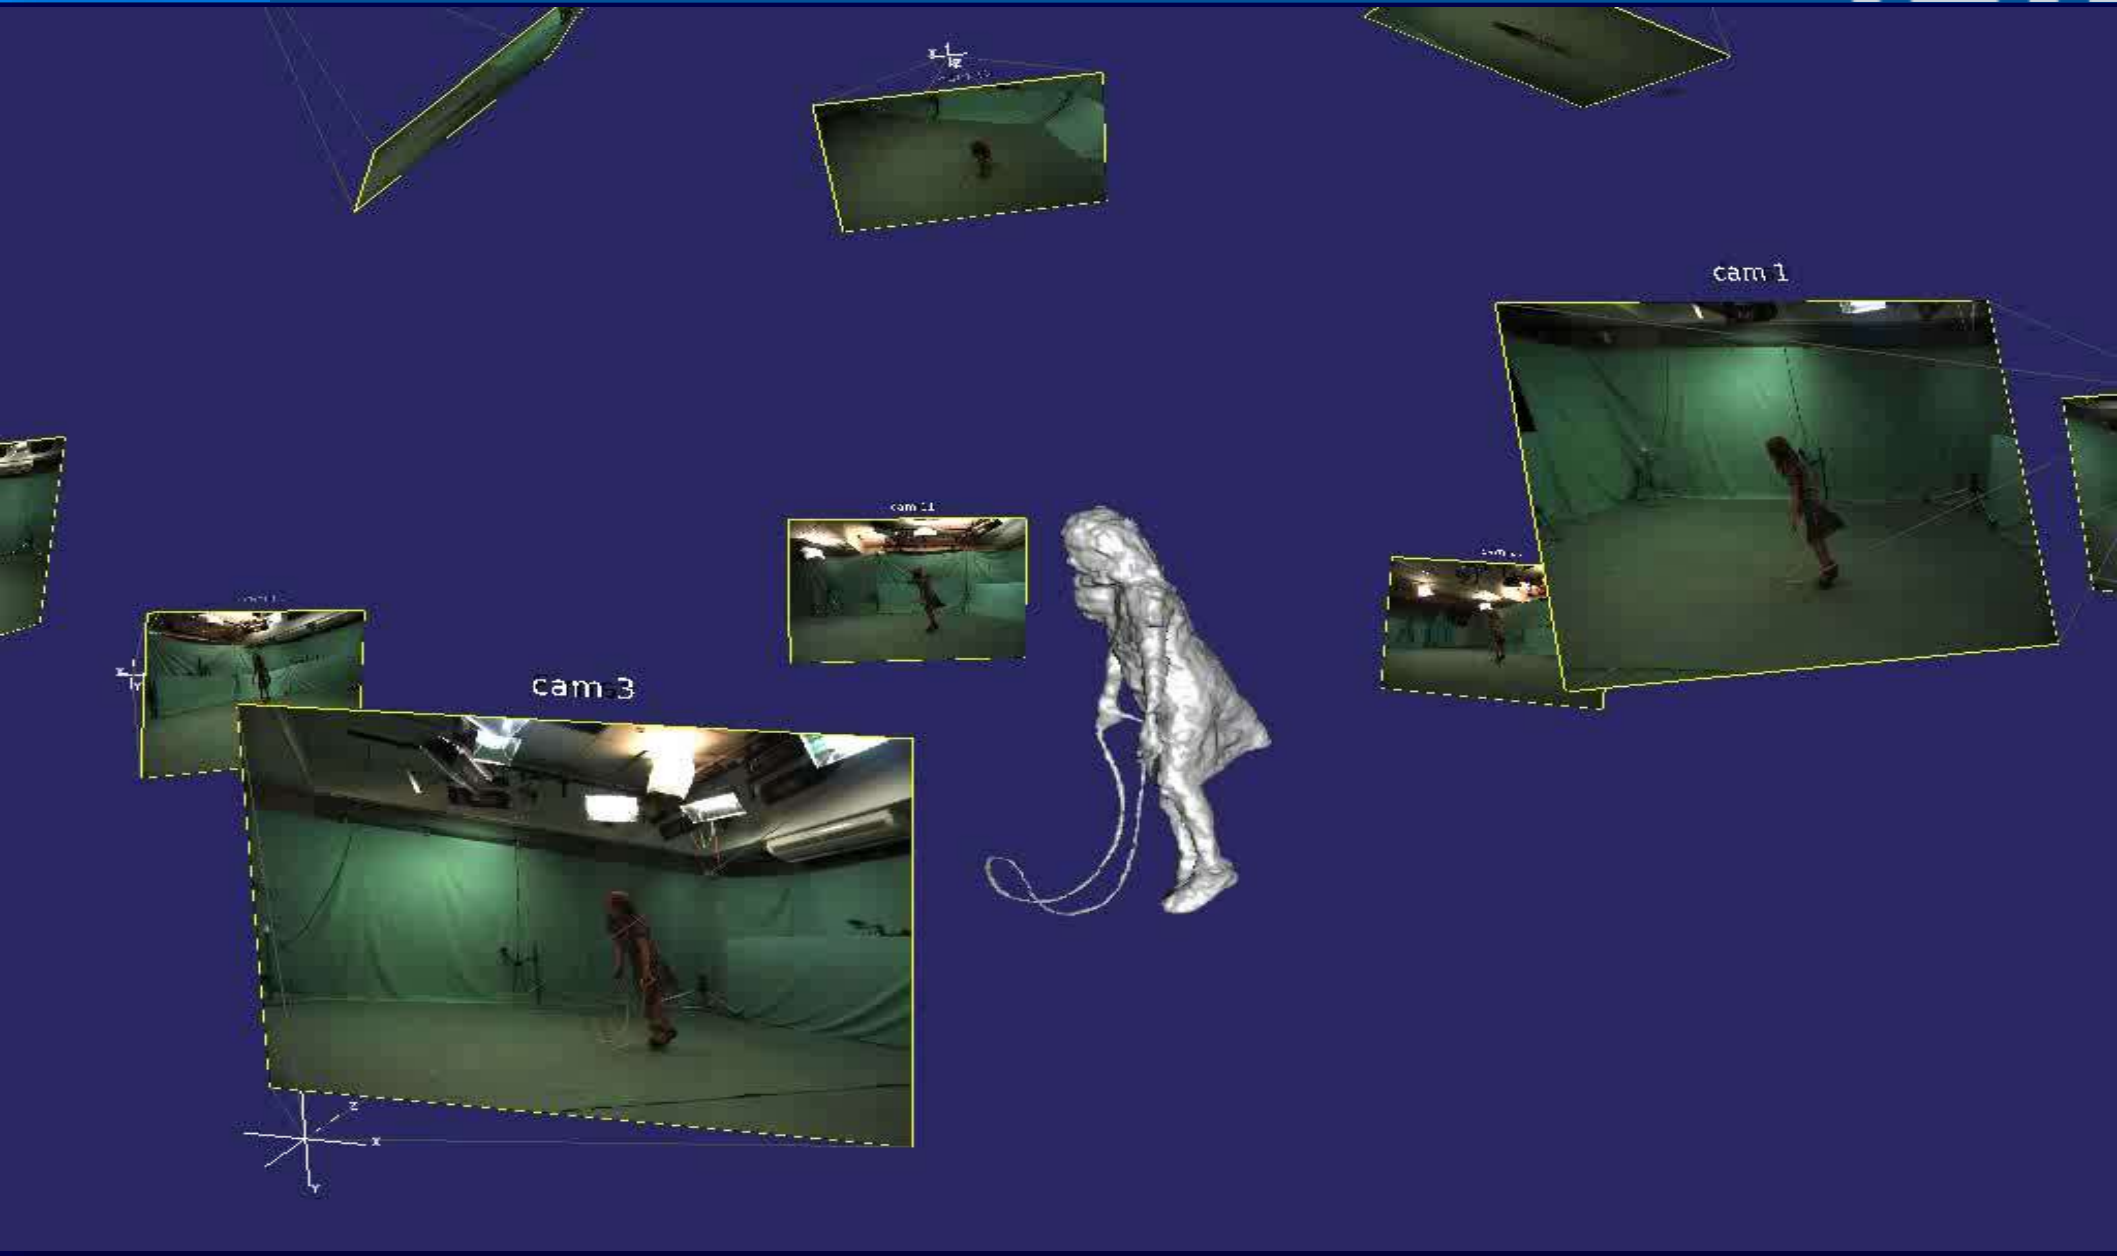

Oswald, Stühmer, Cremers, ECCV ‘14

Action Reconstruction

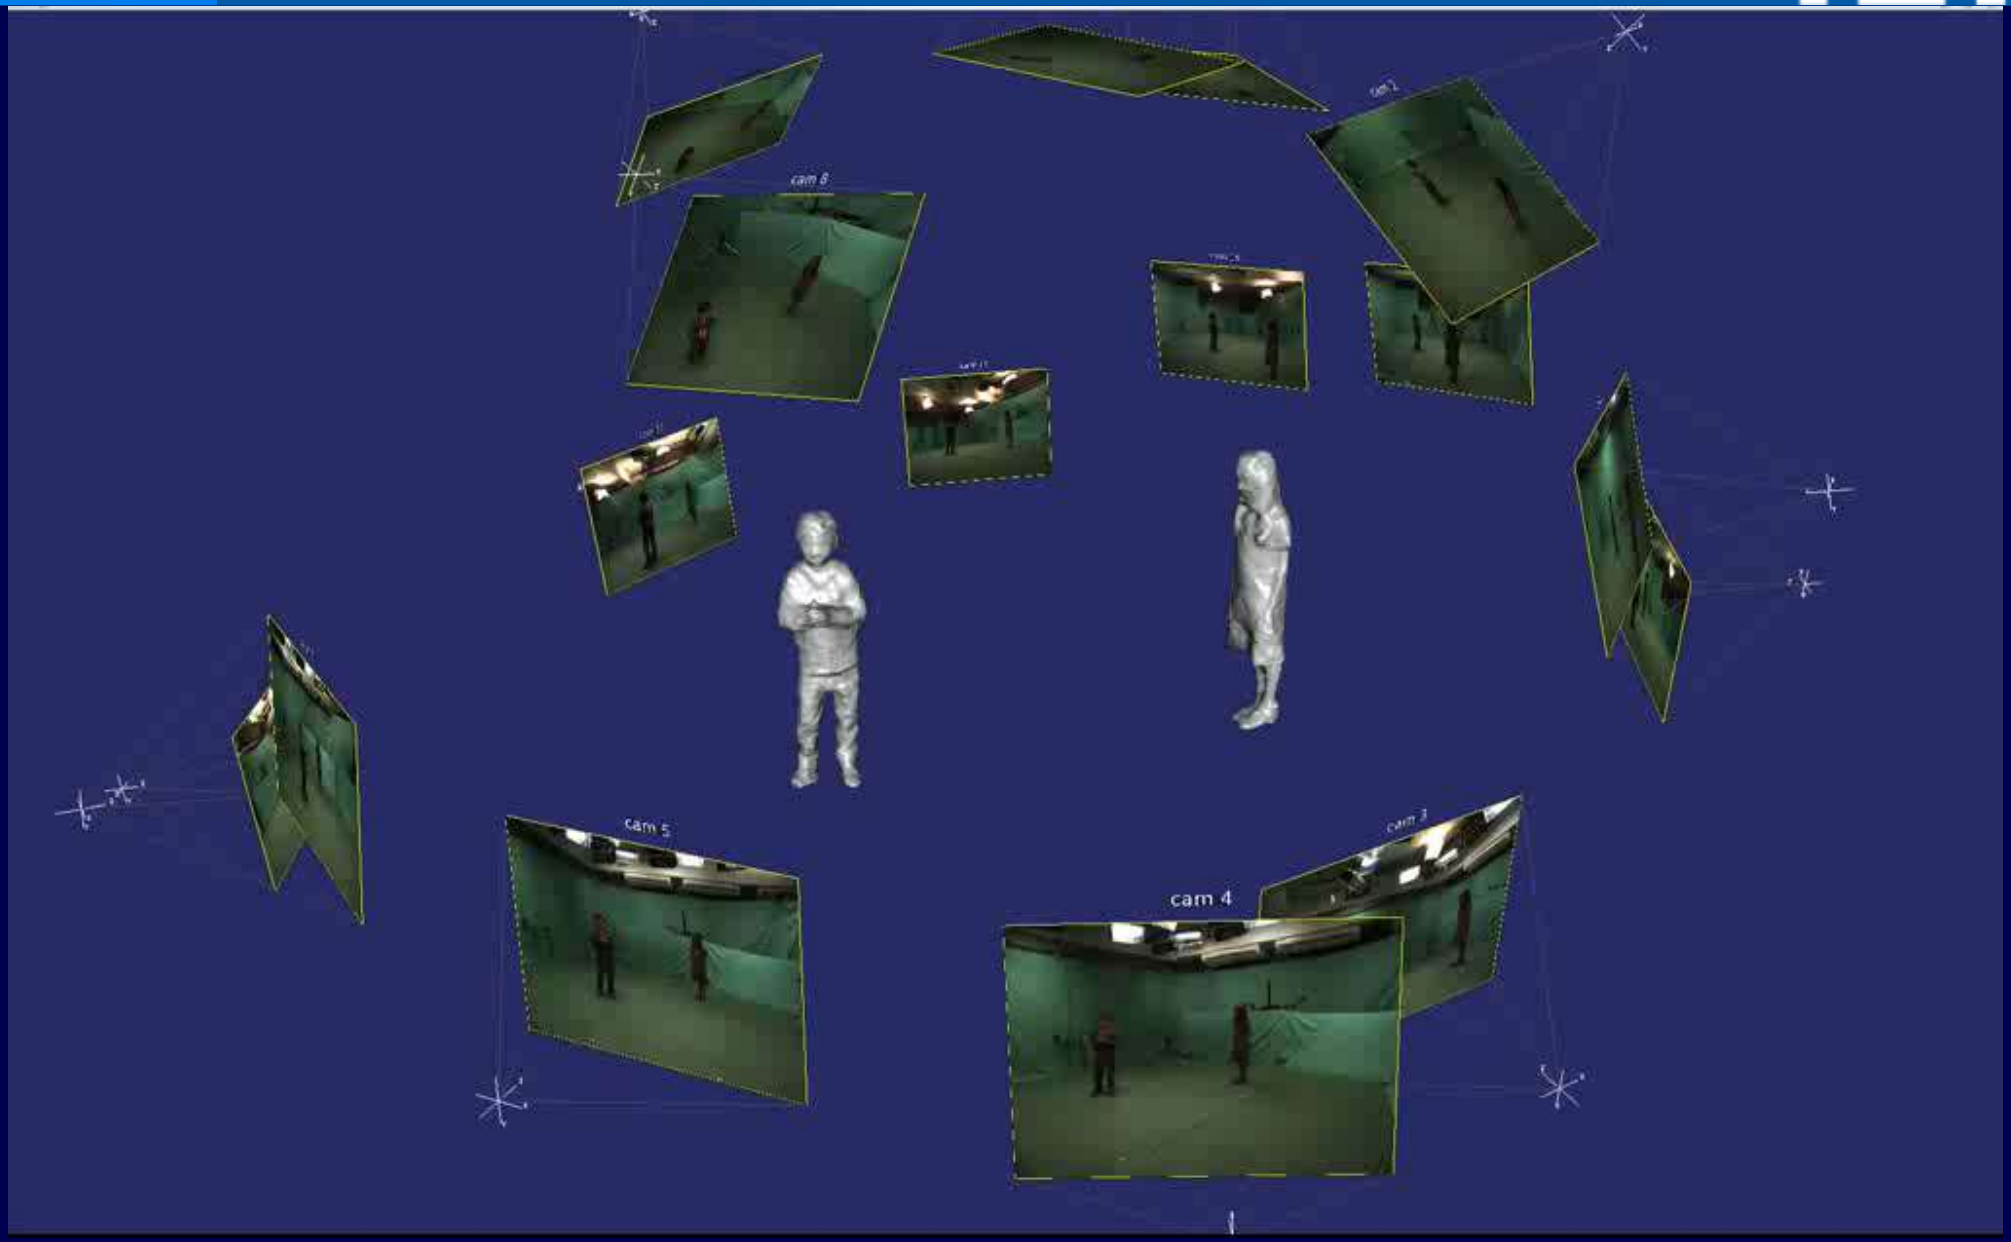

Oswald, Stühmer, Cremers, ECCV '14

Single View Reconstruction

Can we recover geometry from a single image?

Yes: Shape-from-shading, shape-from-focus, shape from symmetry,…

Solution: Fixed-volume silhouette-consistent minimal surface.

$$\min_S |S| \quad \text{s.t. } \text{Vol}(S) = V_0,\ \pi(S) = S_0$$

Single View Reconstruction

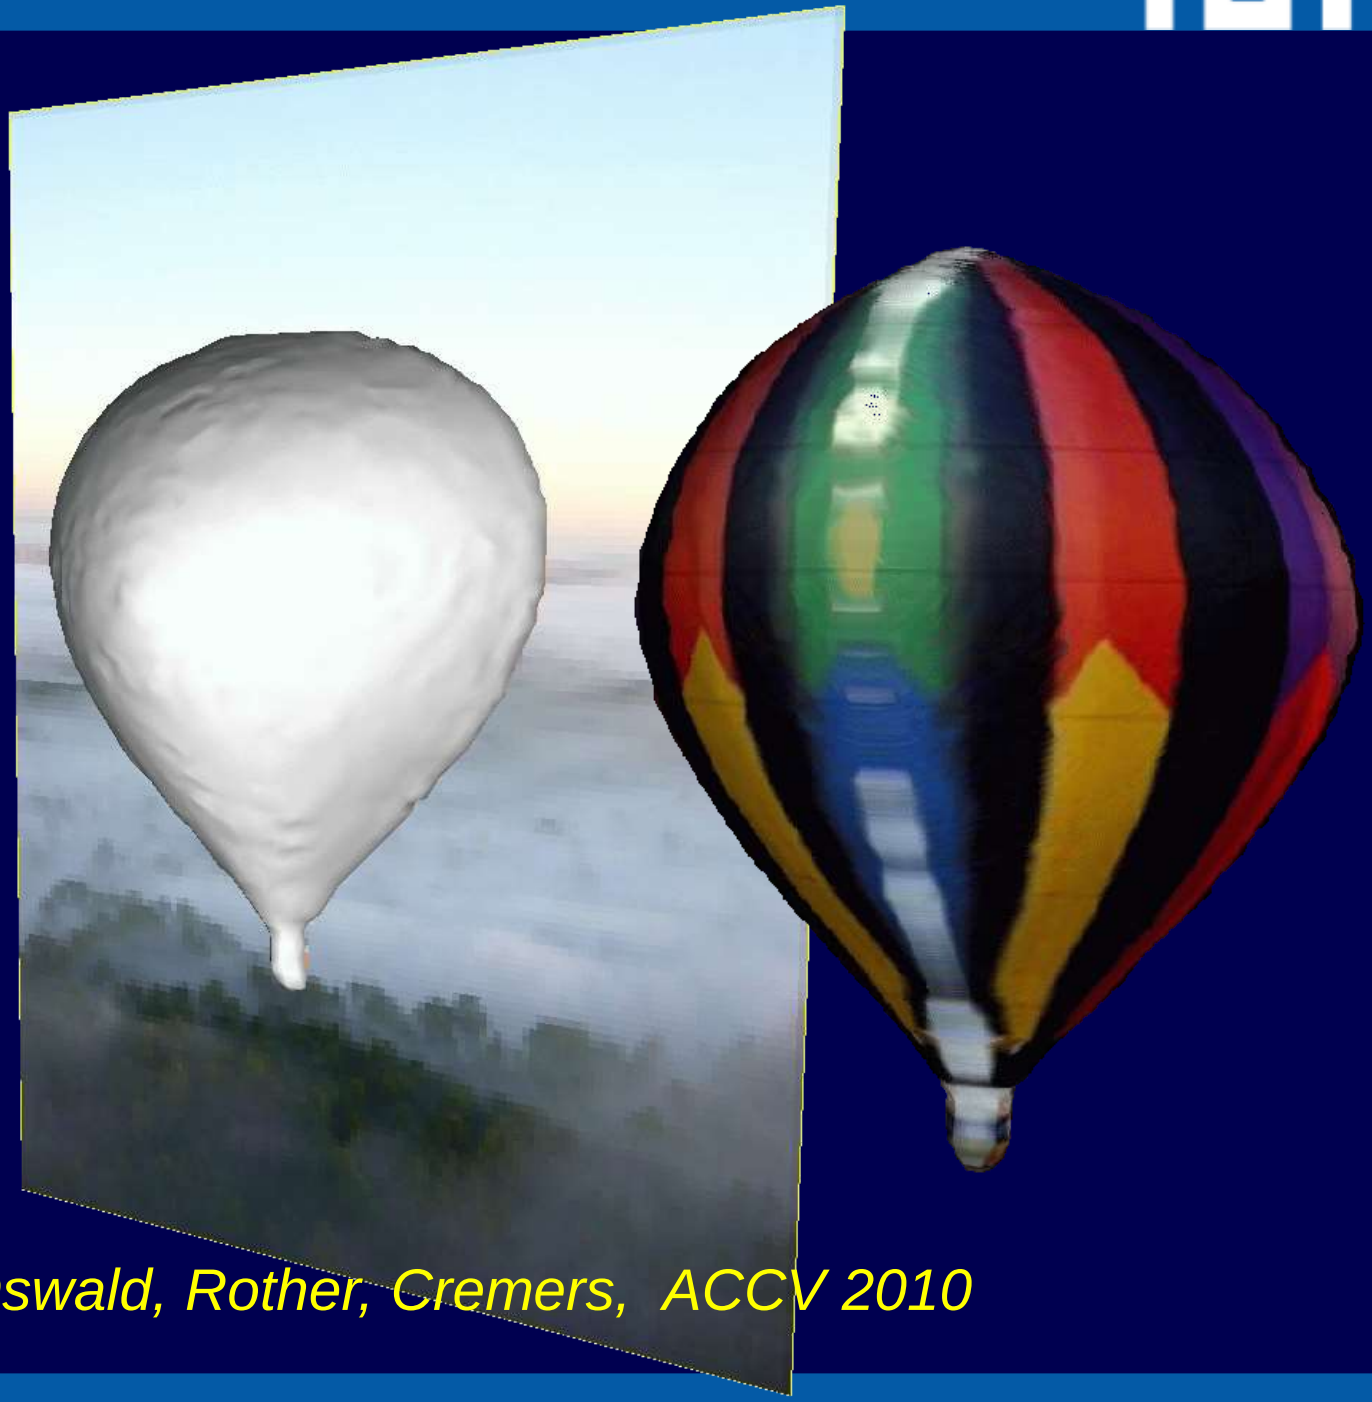

Toeppe, Oswald, Rother, Cremers, ACCV 2010

Single View Reconstruction

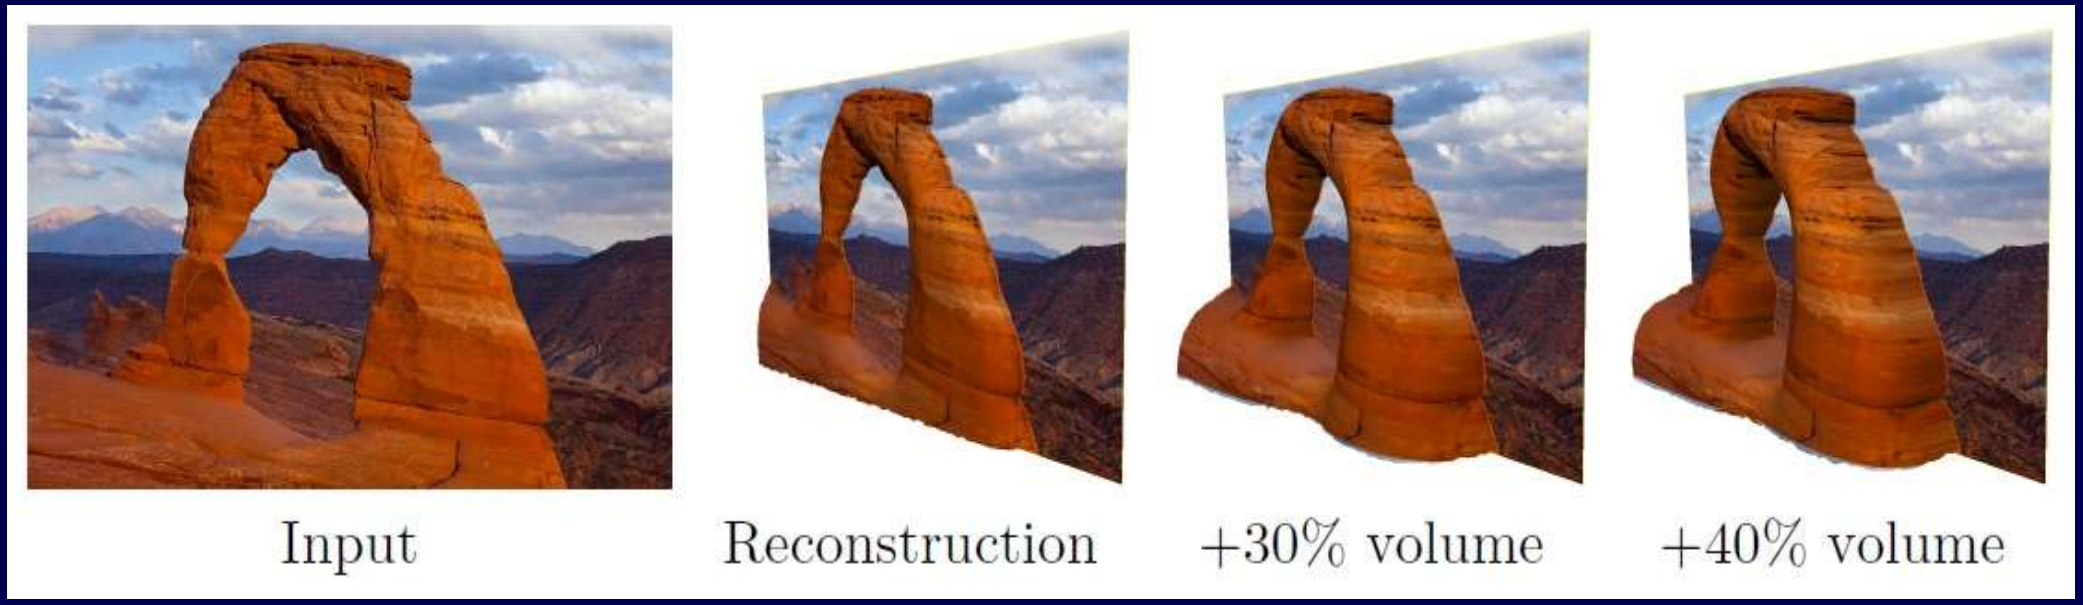

Reconstruction computed in fractions of a second on GPU

Toeppe et al. ACCV 2010, Oswald et al. CVPR 2012

Single View Reconstruction

Toeppe, Oswald, Rother, Cremers, ACCV 2010

Modifying the Material Properties

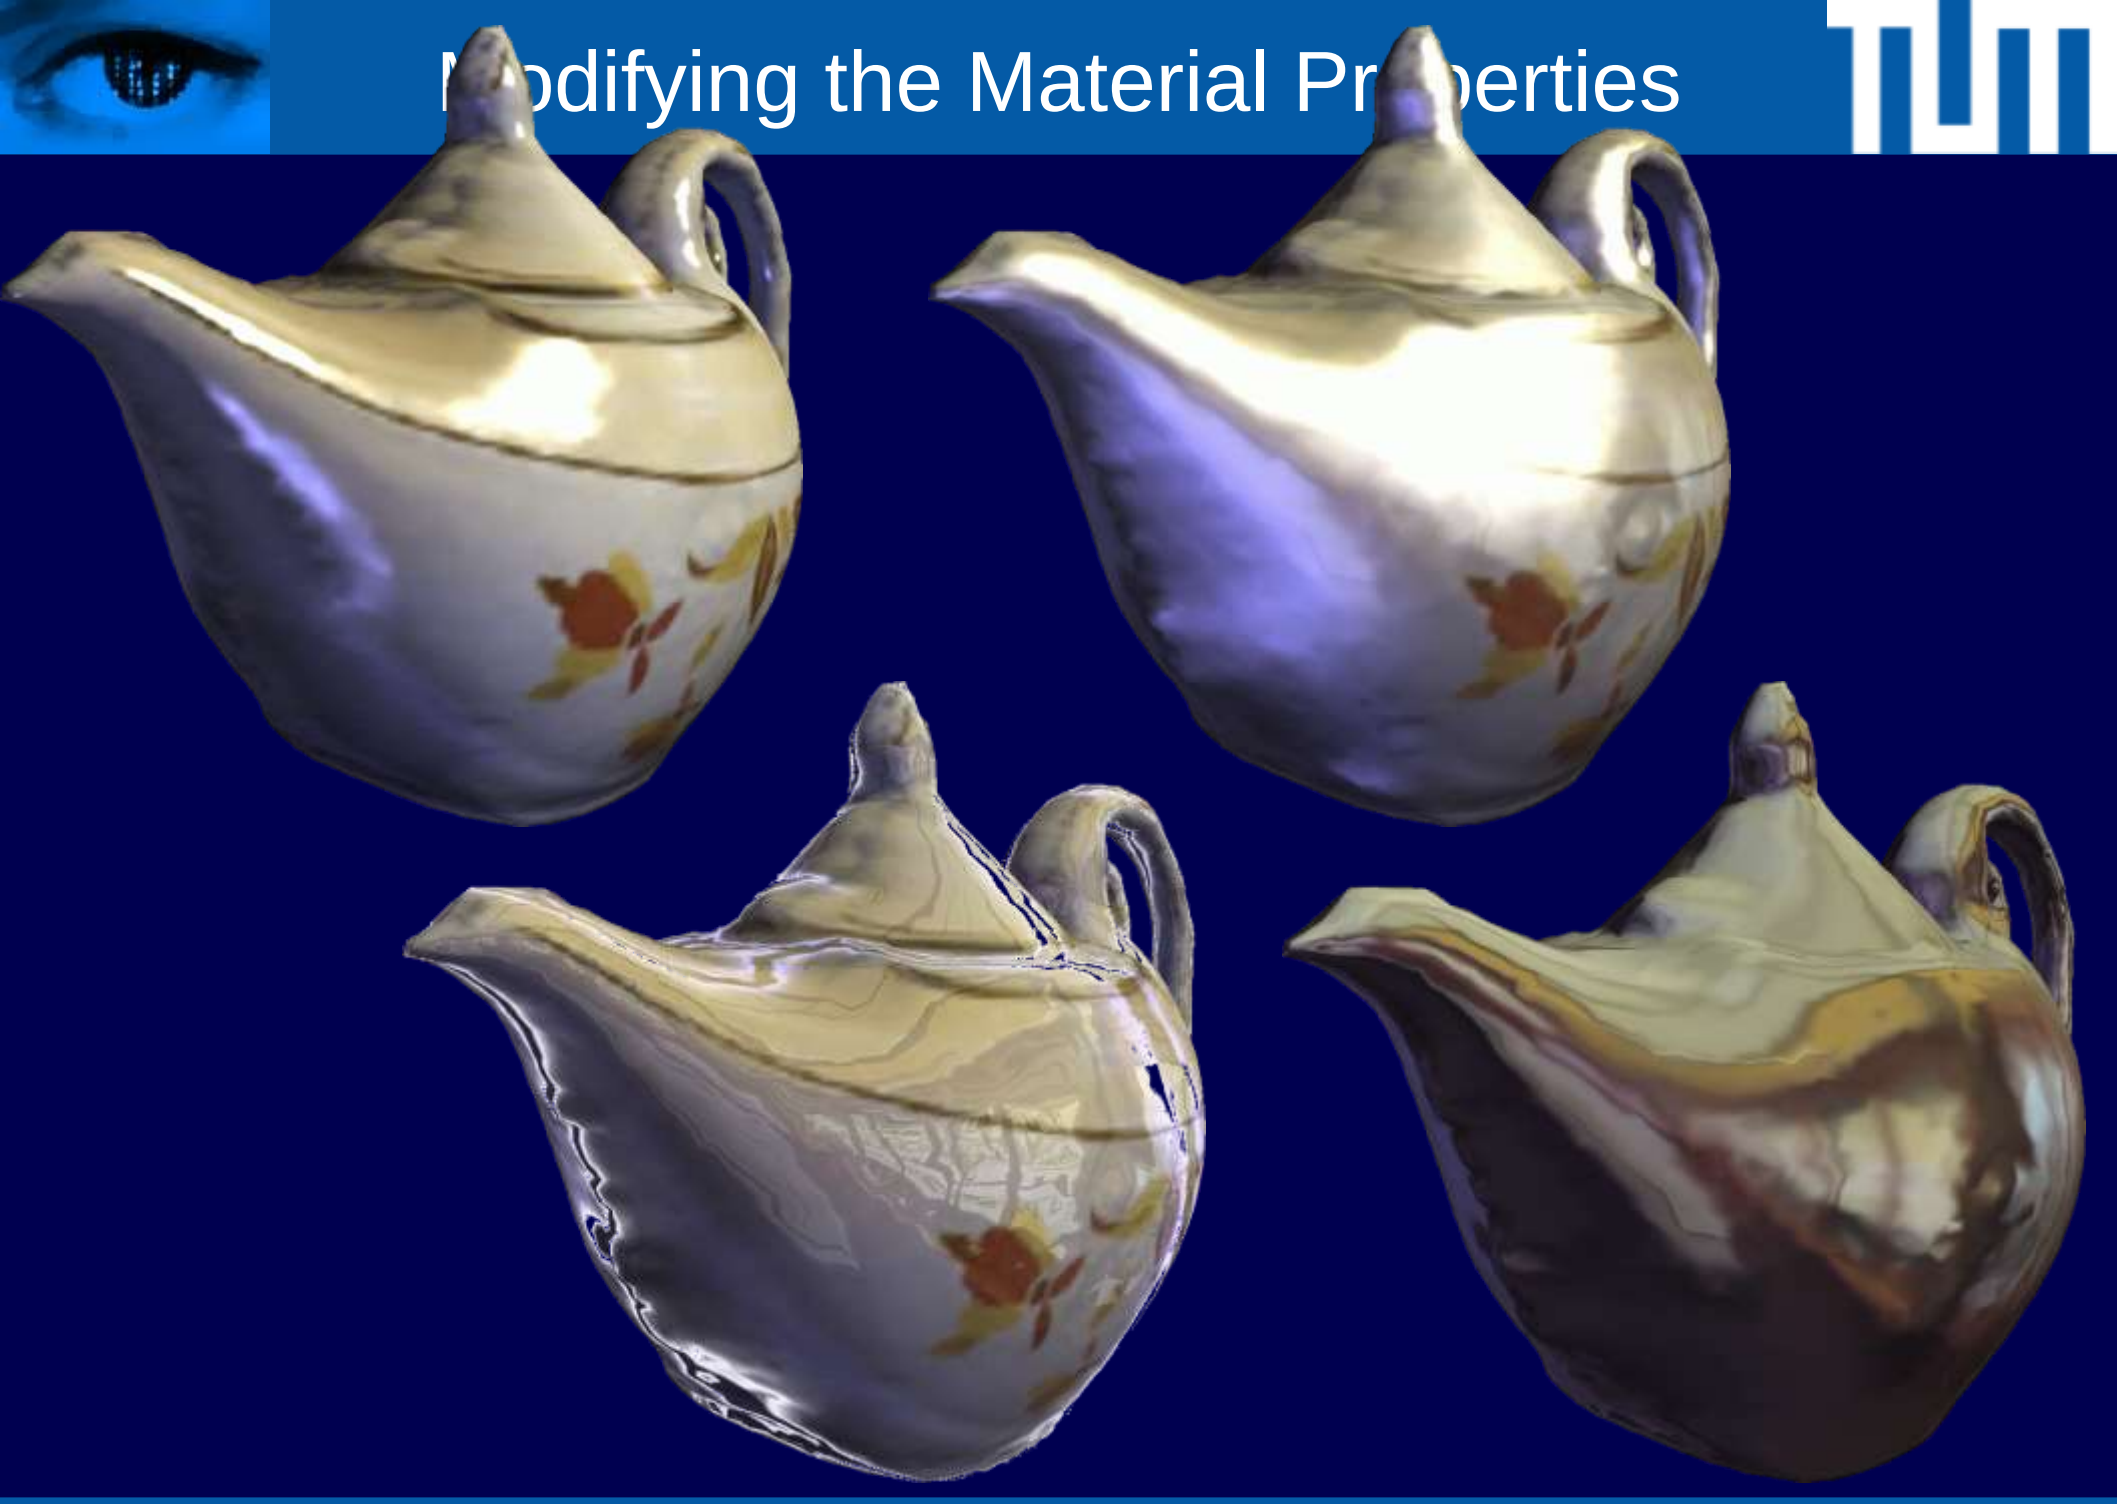

Single View Reconstruction

Toeppe, Oswald, Rother, Cremers, ACCV 2010

Single View Reconstruction

*Toeppe, Oswald, Rother, Cremers, ACCV 2010**

** Best Paper Honorable Mention*

In collaboration with Microsoft Research

Overview

Multiview reconstruction

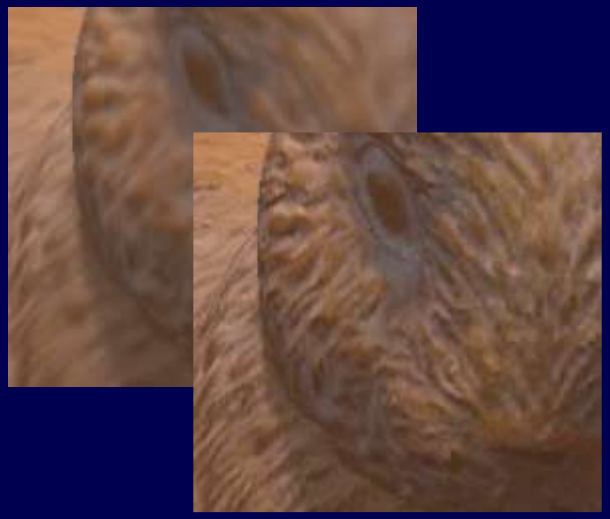

Super-res.textures

Stereo reconstruction

Realtime dense geometry

RGB-D cameras

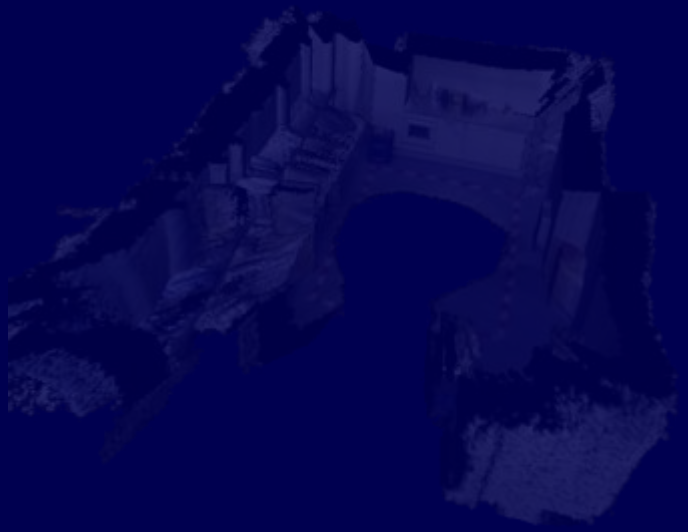

Reconstruction on the fly

Evolution to Global Optimum

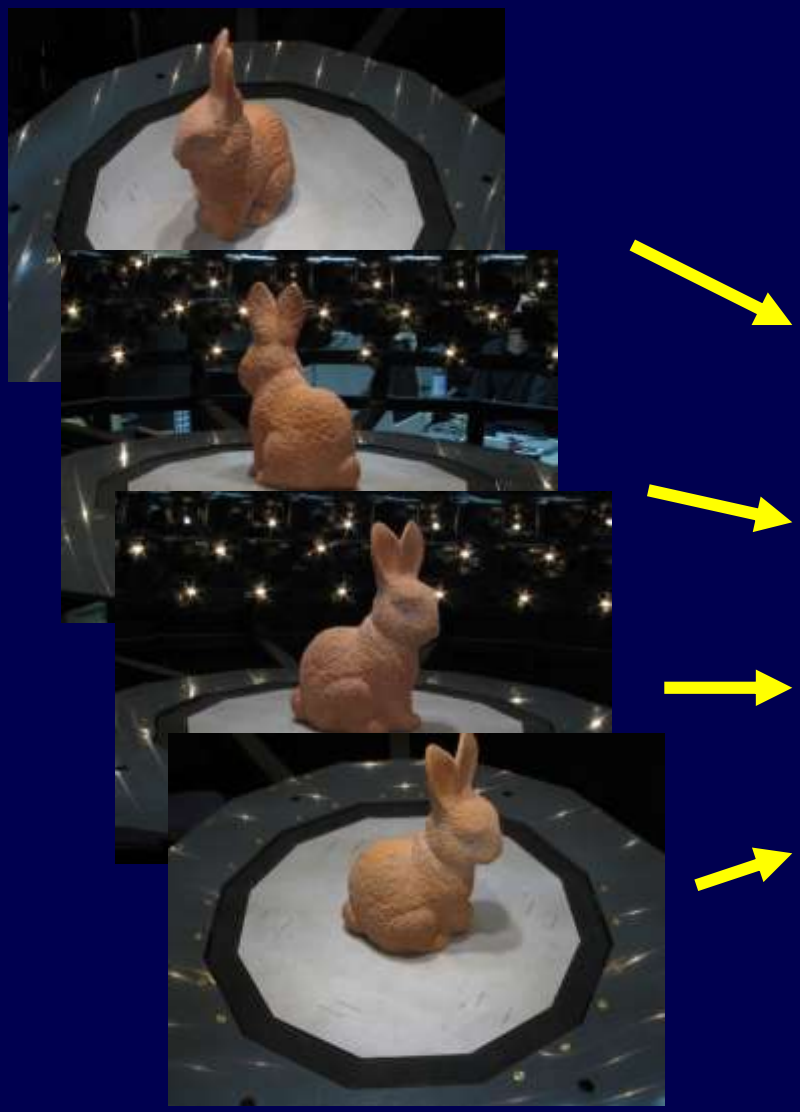

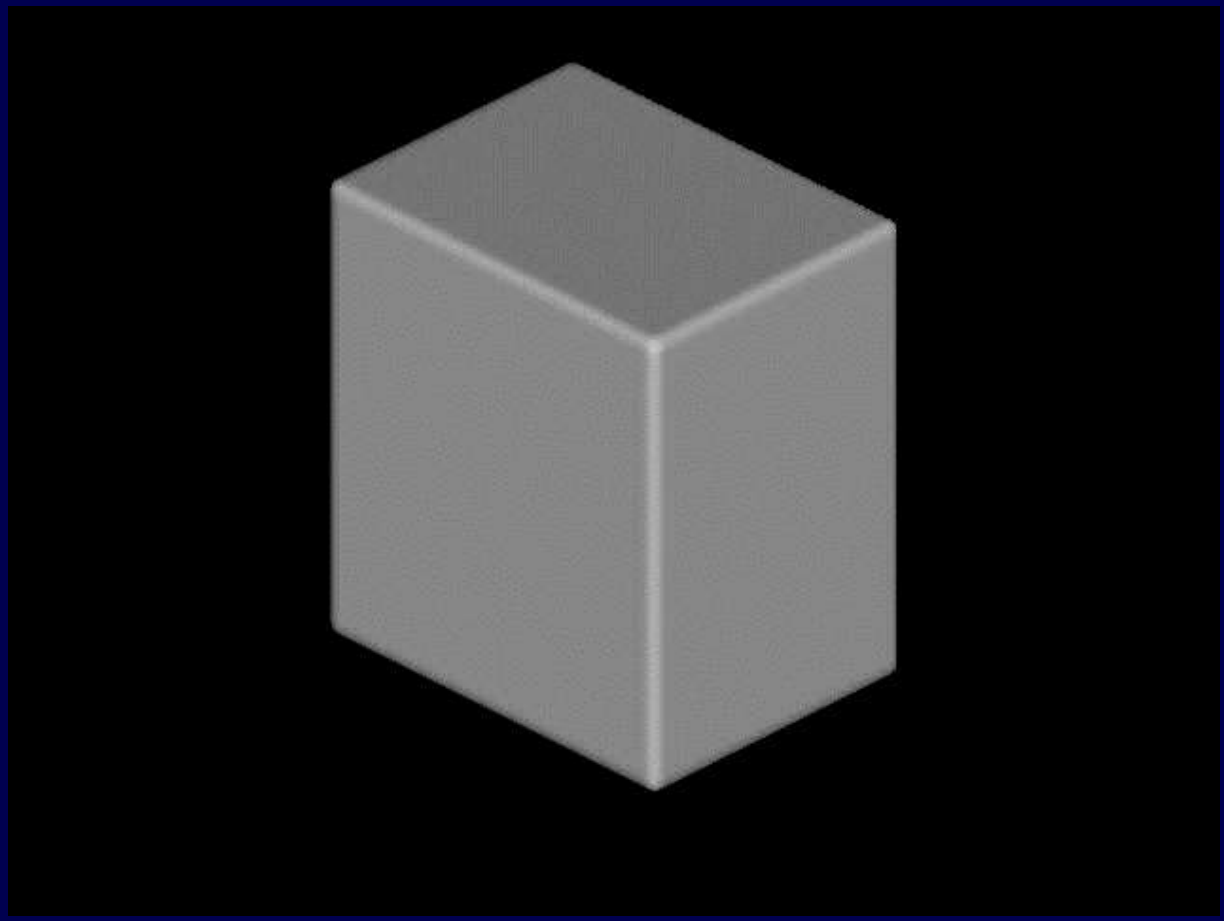

Kolev, Klodt, Brox, Cremers, IJCV 2009

Super-Resolution Texture Map

Given all images $\mathcal{I}_i : \Omega_i \rightarrow \mathbb{R}^3$, determine the surface color $T : S \rightarrow \mathbb{R}^3$

$$\min_T \sum_{i=1}^{n} \int_{\Omega_i} \left(b * (T \circ \beta_i) - \mathcal{I}_i \right)^2 dx + \lambda \int_S \|\nabla_S T\| ds$$

blur & downsample

back-projection

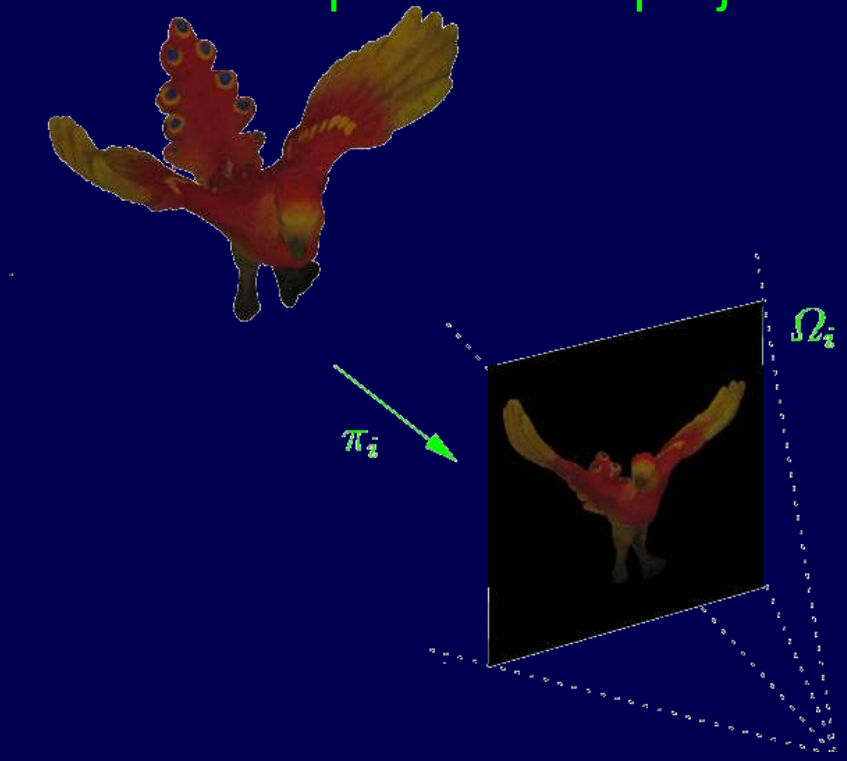

Goldlücke, Cremers, ICCV '09, DAGM '09, IJCV '13*

** Best Paper Award*

Super-Resolution Texture Map

Goldlücke, Cremers, ICCV ’09, DAGM ’09, IJCV ‘13*

** Best Paper Award*

Super-Resolution Texture Map

Weighted average

Super-resolution texture

Goldlücke, Cremers, ICCV '09, DAGM '09, IJCV '13*

** Best Paper Award*

Super-Resolution Texture Map

Closeup of input image

Super-resolution texture

Goldlücke, Cremers, ICCV ’09, DAGM ’09, IJCV ‘13*

** Best Paper Award*

Overview

Multiview reconstruction

Super-res.textures

Stereo reconstruction

Realtime dense geometry

RGB-D cameras

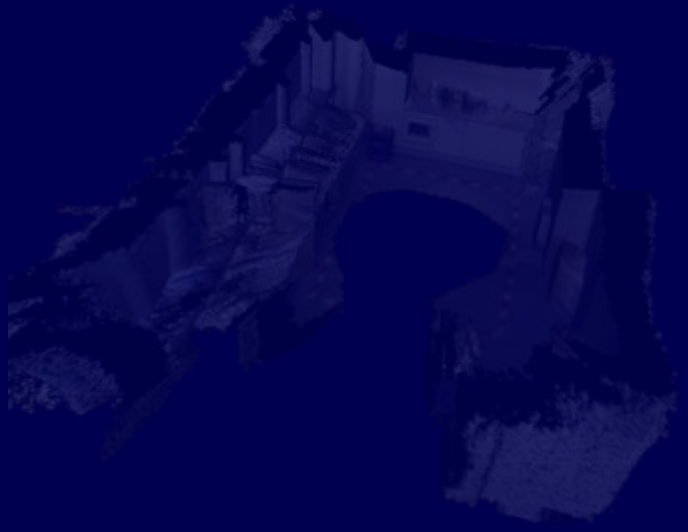

Reconstruction on the fly

From Binary to Multilabel Optimization

$$u : \Omega \rightarrow \Gamma = [\gamma_{min}, \gamma_{max}]$$

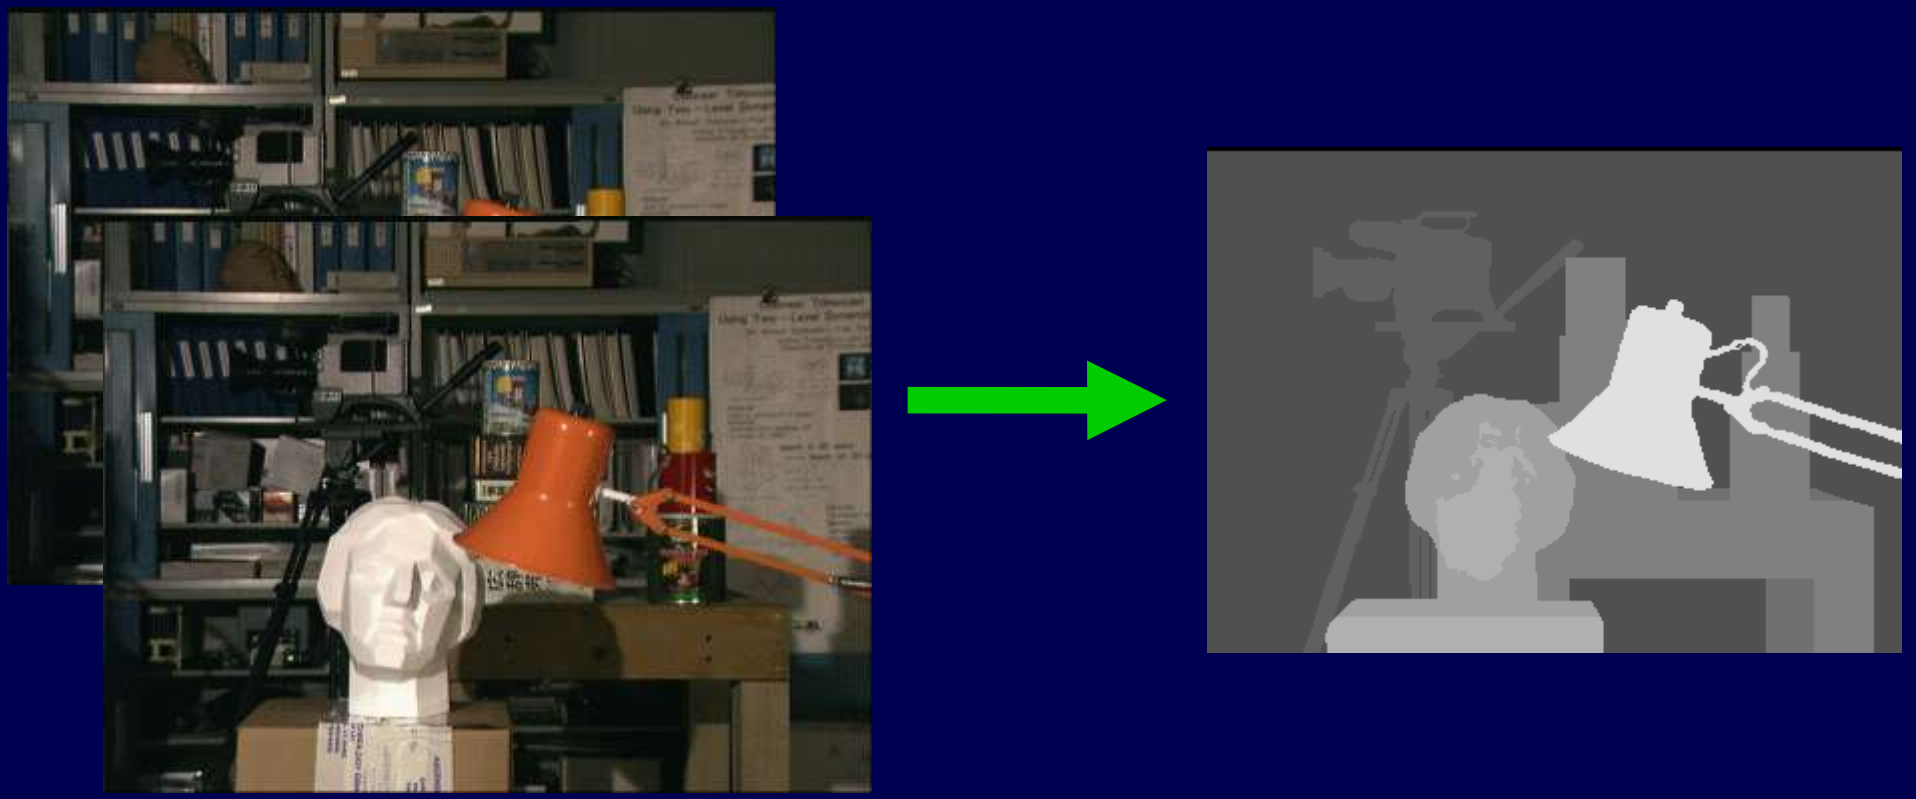

Example: Stereo Reconstruction

Pock, Schoenemann, Bischof, Cremers, Europ. Conf. on Computer Vision '08:
Theorem: Stereo reconstruction can be solved by convex optimization.

Evolution to Global Minimum

Reconstruction from Aerial Images

1/2 input images (6 Mpixel)
Courtesy of H. Hirschmüller

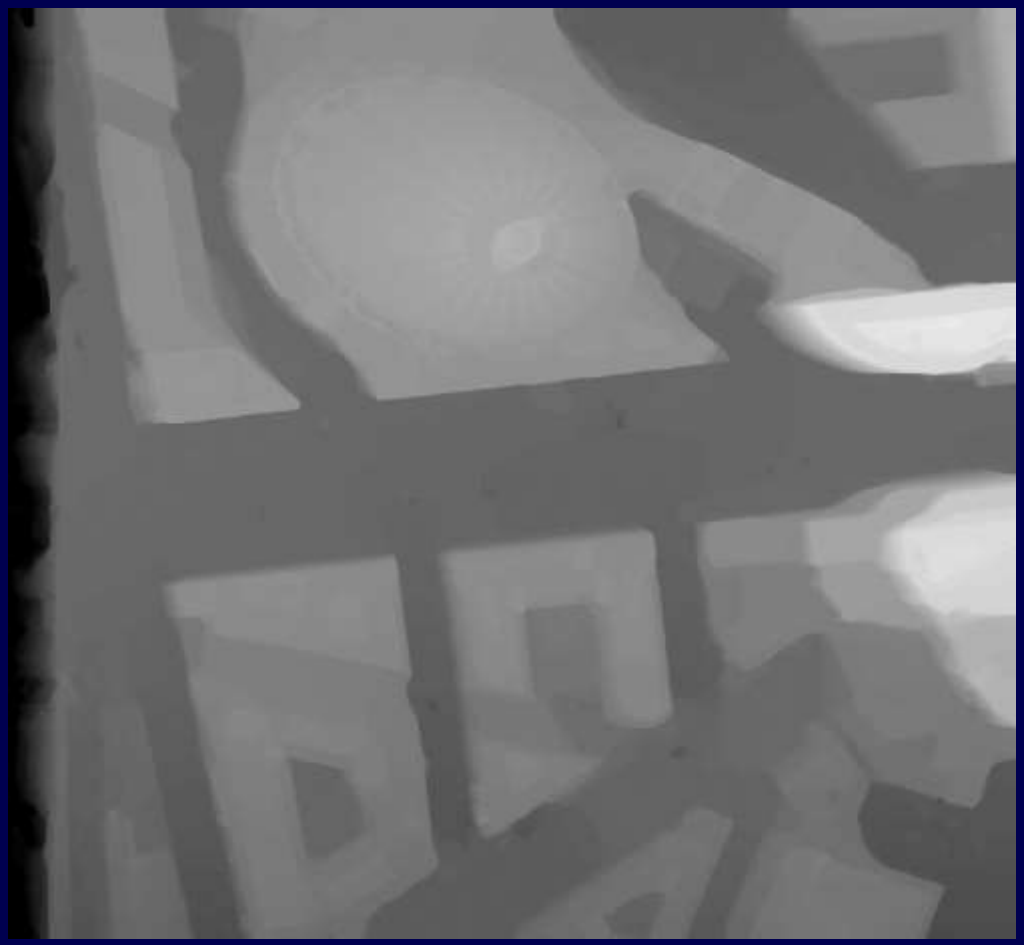

Depth reconstruction
77 seconds

Stangl, Souiai, Cremers, GCPR '13

Reconstruction from Aerial Images

One of two input images

Courtesy of Microsoft

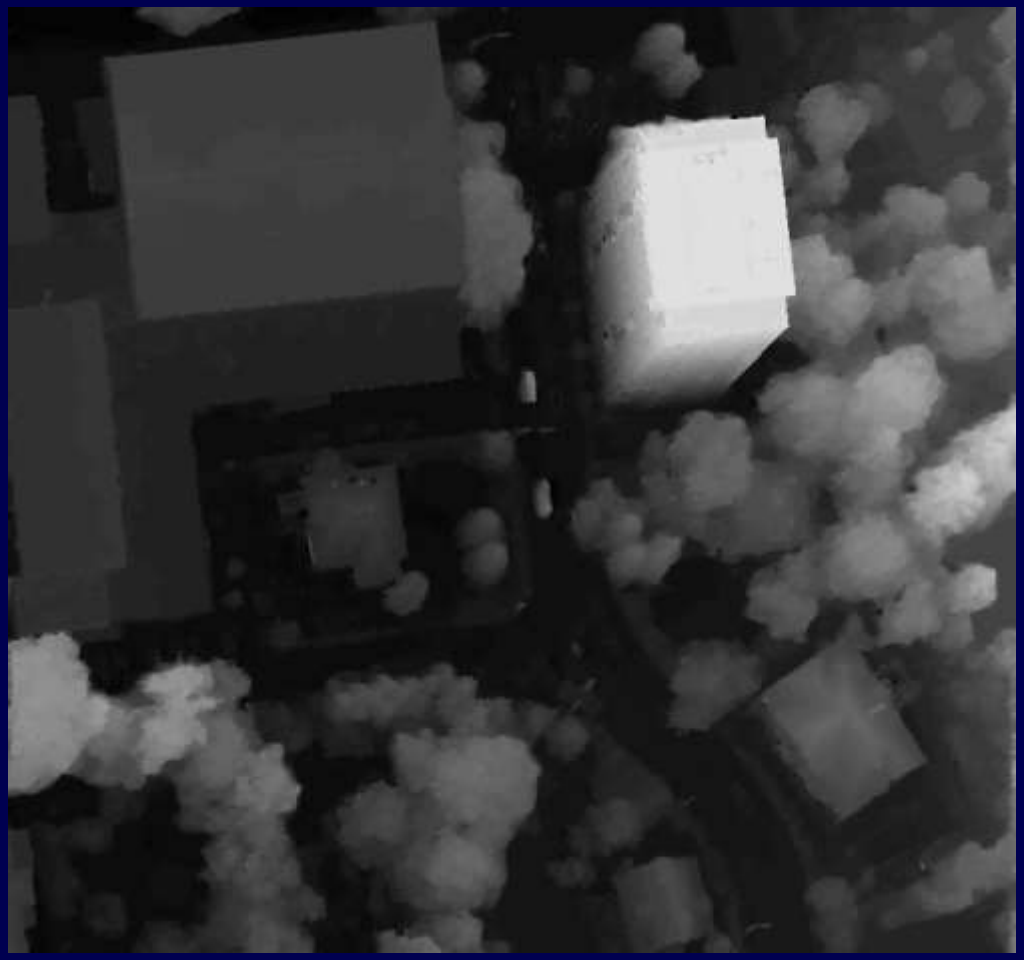

Depth reconstruction

Reconstruction from Aerial Images

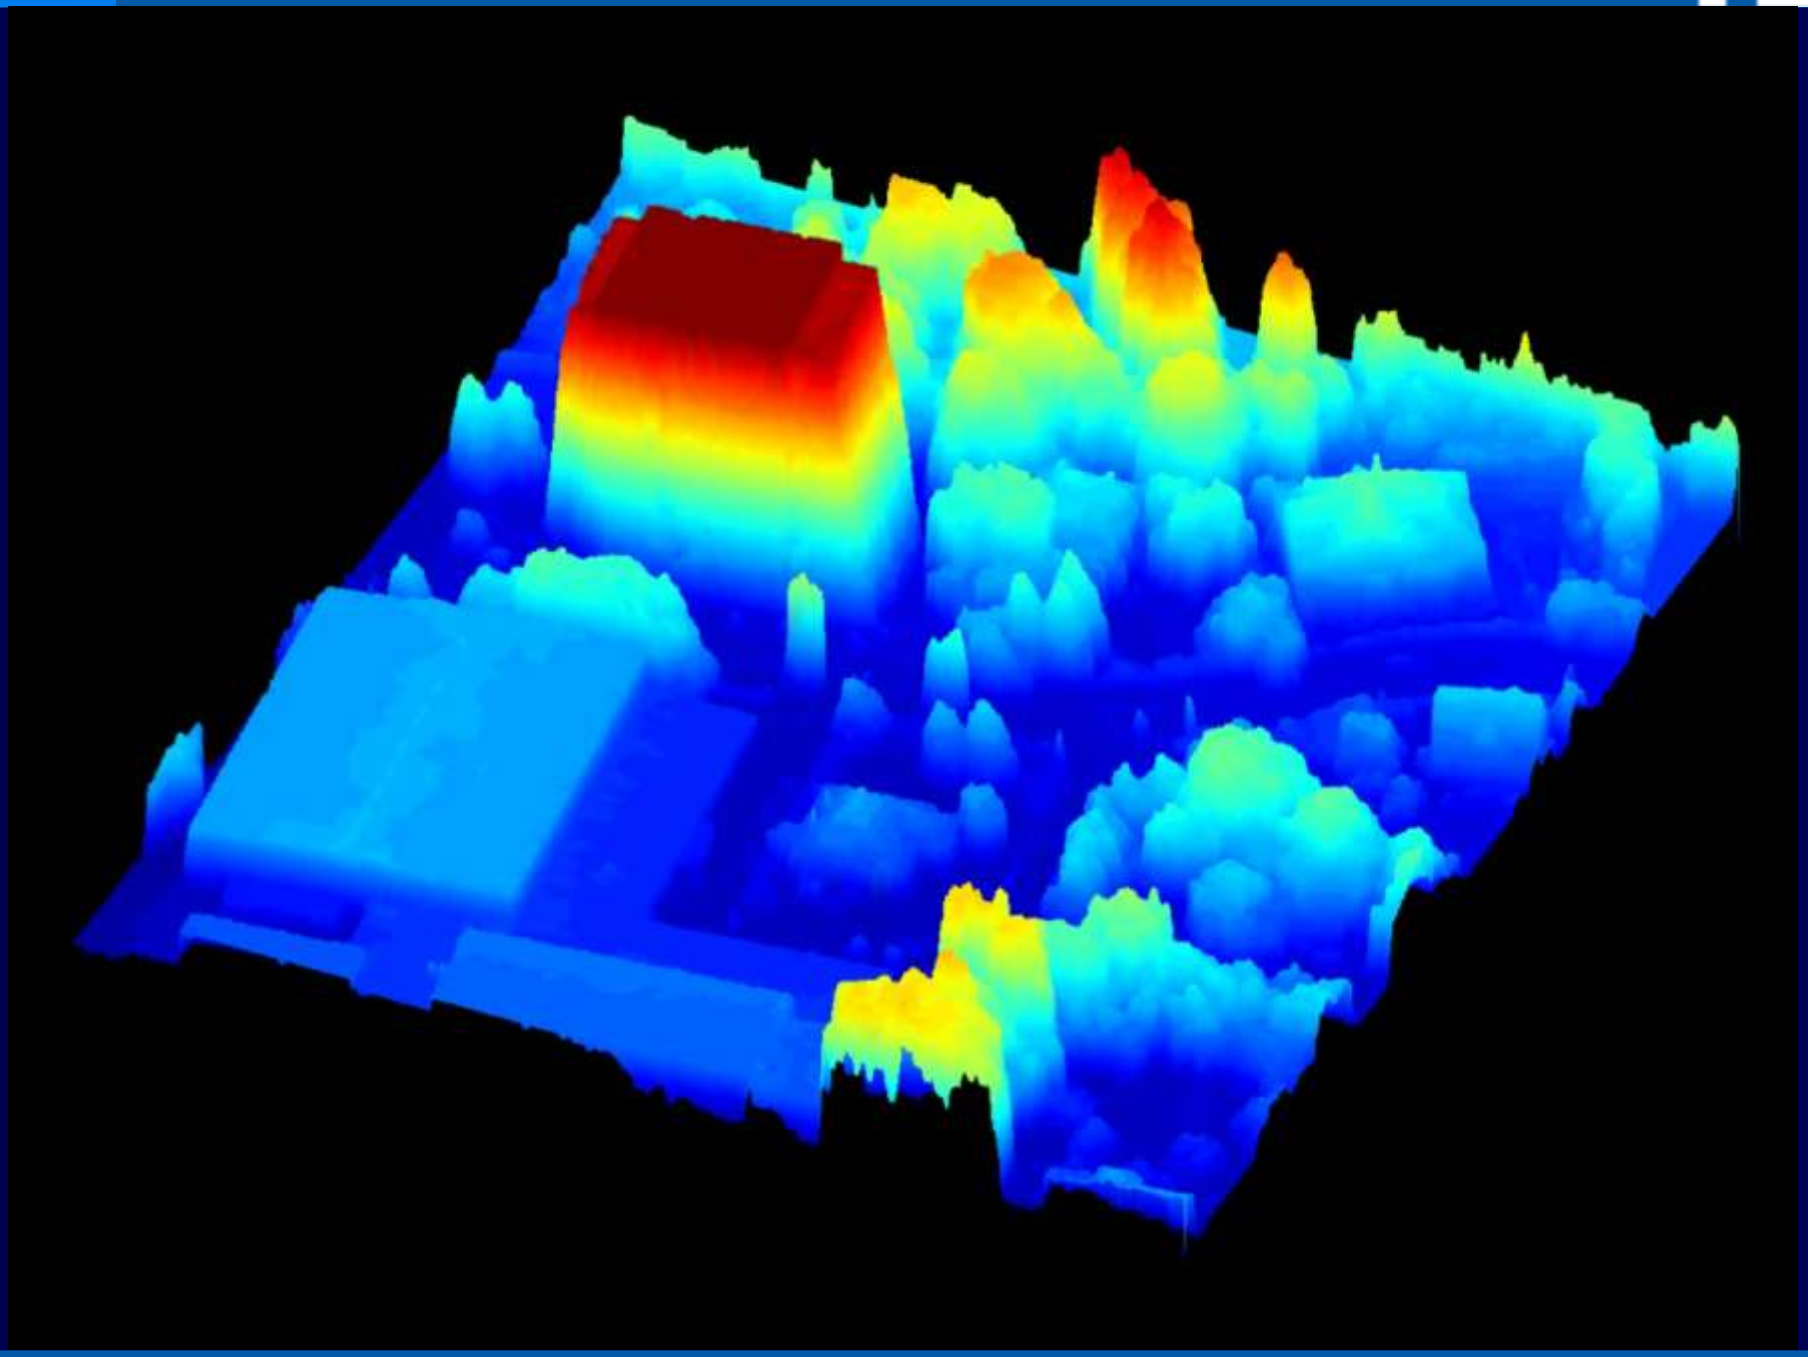

Highly accurate Stereo Reconstruction

Stereo input

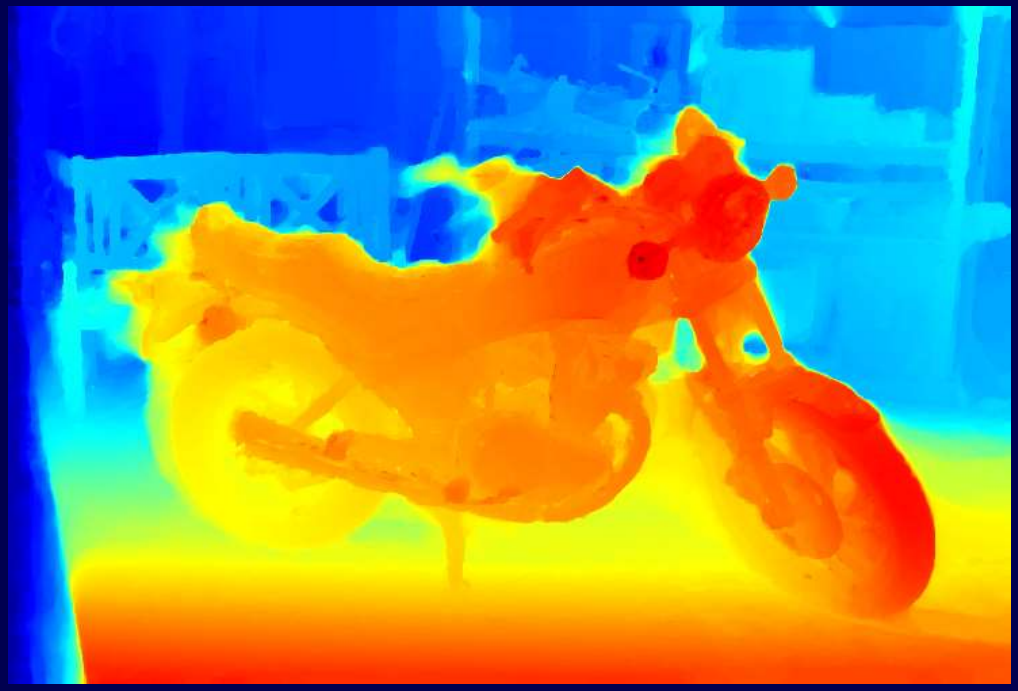

Depth reconstruction

*Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR ’16 **

* Best Paper Honorable Mention

Munich from the Air

1/2 input images (1000x1000)

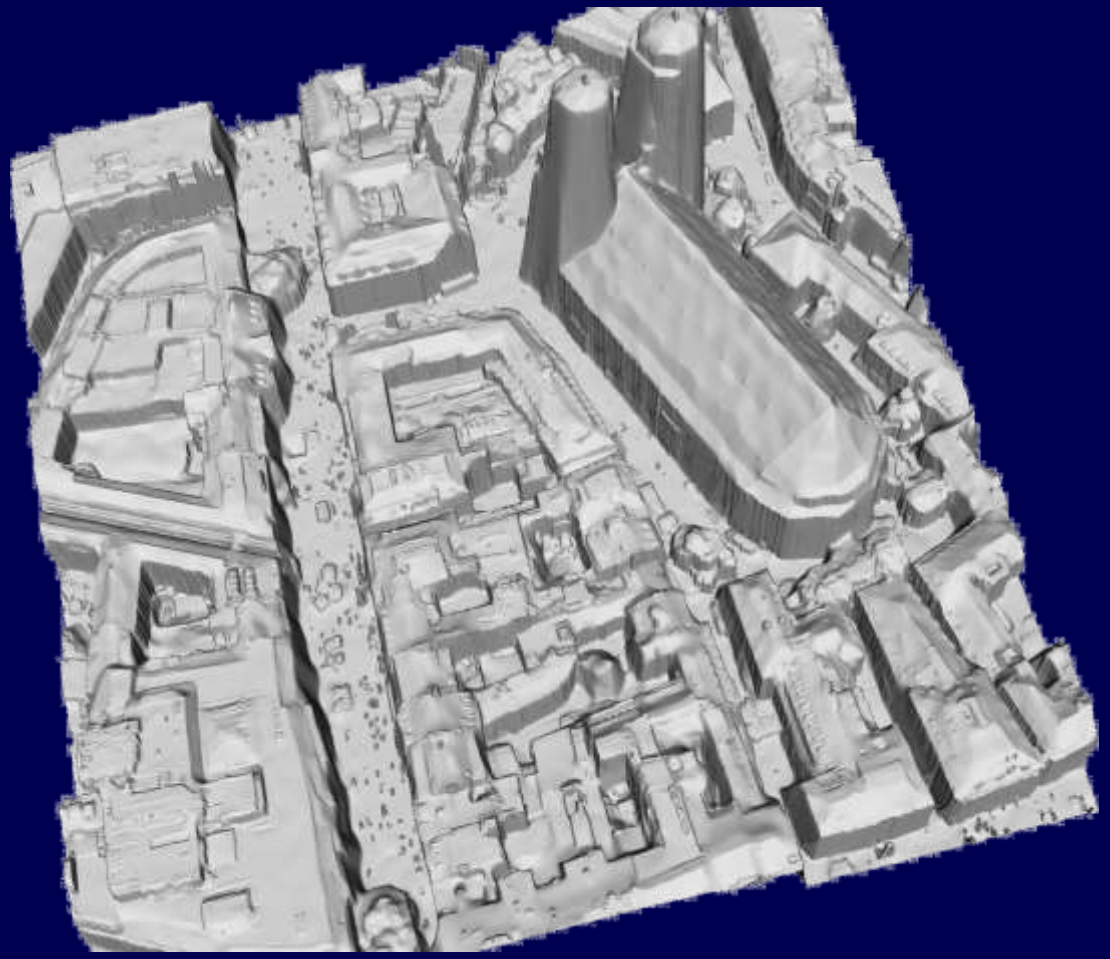

Depth reconstruction

Kuschk, Cremers, ICCV Big Data Workshop 2013

Overview

Multiview reconstruction

Super-res.textures

Stereo reconstruction

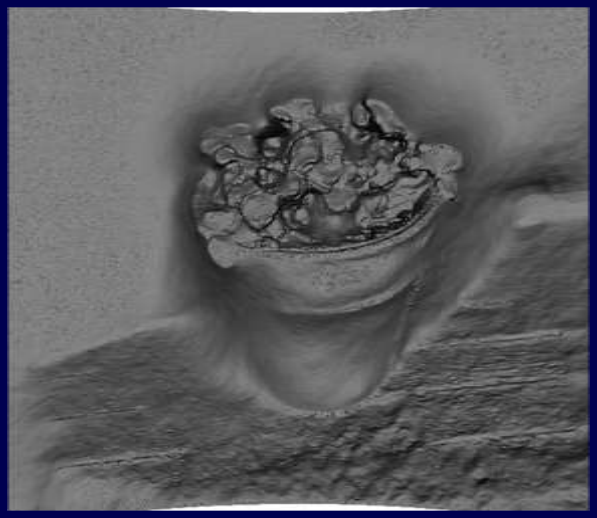

Realtime dense geometry

RGB-D cameras

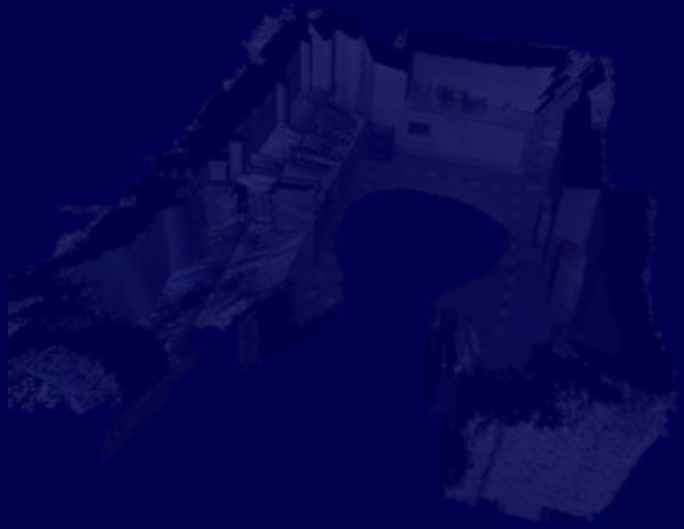

Reconstruction on the fly

From Dense Flow to Dense Geometry

Input video

Optical flow field

$$\min_{u:\Omega\to\mathbb{R}^2} \int_\Omega \left| I_1(x) - I_2(x+u) \right| dx \; + \; J(u)$$

Horn & Schunck '81, Zach et al. DAGM '07, Wedel et al. ICCV '09

From Dense Flow to Dense Geometry

Input video

Optical flow field*

* 60 fps @ 640x480

$$\min_{u:\Omega\to\mathbb{R}^2} \int_\Omega \left| I_1(x) - I_2(x+u) \right| dx \; + \; J(u)$$

Horn & Schunck '81, Zach et al. DAGM '07, Wedel et al. ICCV '09

Dense geometry from hand-held camera

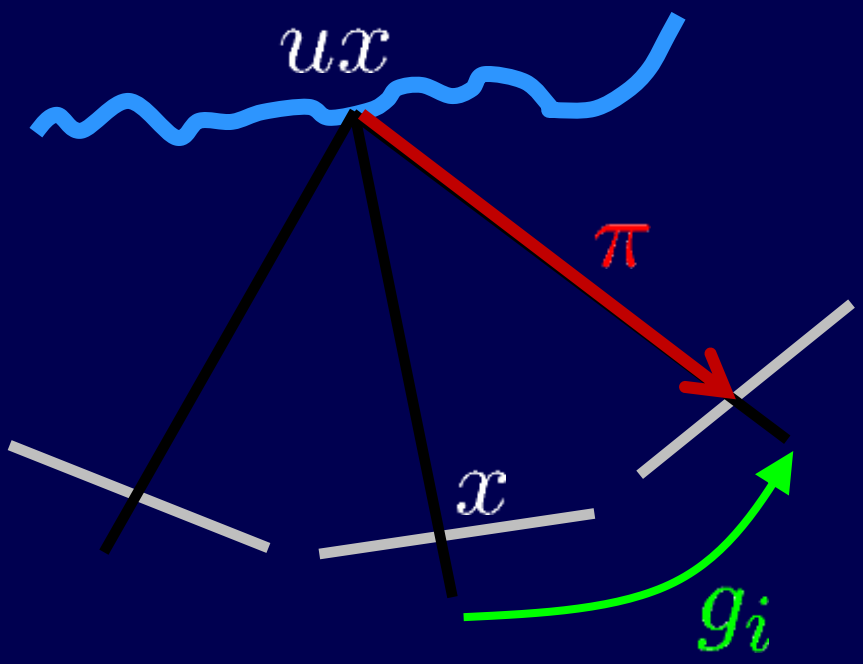

Brightness constancy:

$$I_0(x) \stackrel{!}{=} I_i\Big(\pi\big(g_i(ux)\big)\Big)$$

$$\min_u \sum_i \int_\Omega \Big|I_0(x) - I_i\Big(\pi\big(g_i(u \cdot x)\big)\Big|\, dx + \int_\Omega |\nabla u(x)|\, dx$$

Stuehmer , Gumhold, Cremers, DAGM '10

Dense geometry from hand-held camera

Stuehmer, Gumhold, Cremers, DAGM ’10

Dense geometry from hand-held camera

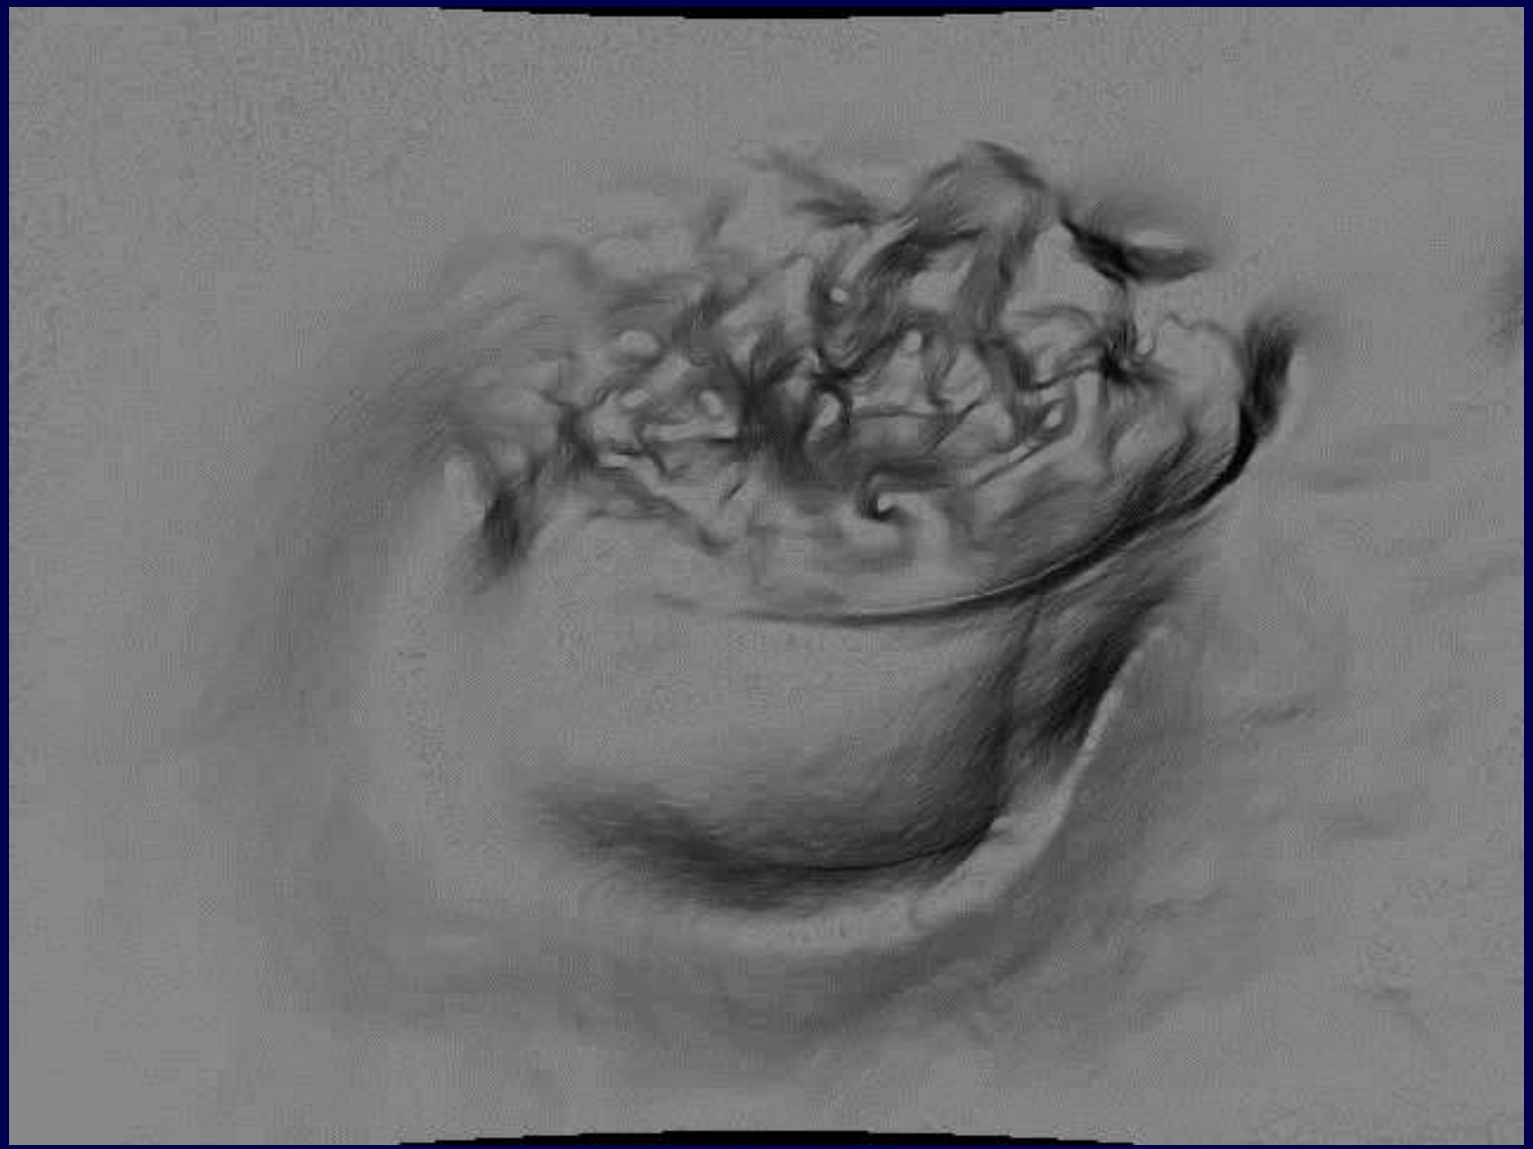

Stuehmer, Gumhold, Cremers, DAGM ’10

Realtime Dense Reconstruction

16.0 fps

22.0 fps

41.1 fps

Stuehmer, Gumhold, Cremers, DAGM ’10

Overview

Multiview reconstruction

Super-res.textures

Stereo reconstruction

Realtime dense geometry

RGB-D cameras

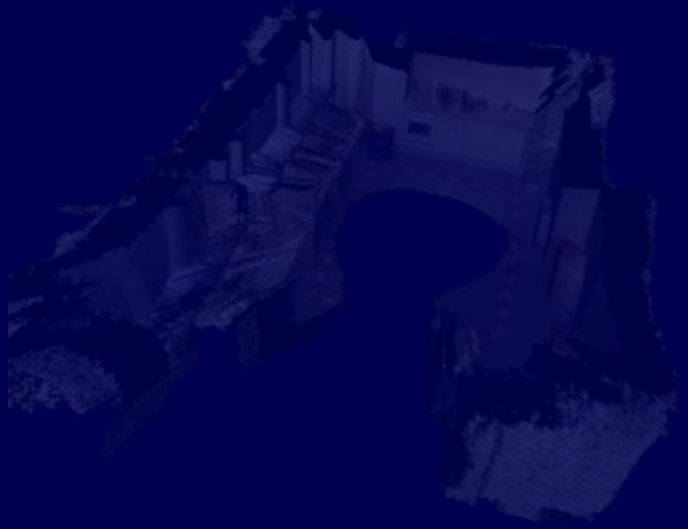

Reconstruction on the fly

RGB-D Camera Tracking

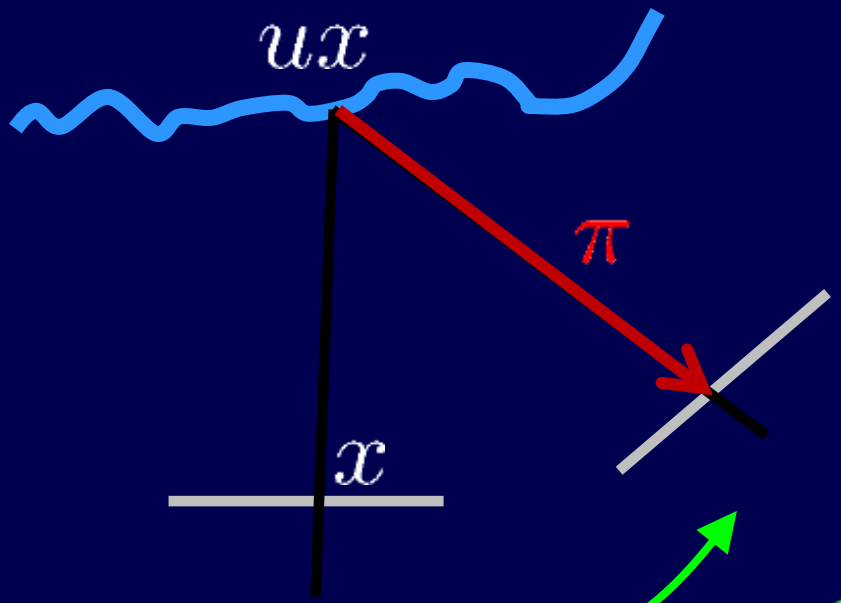

$$g_\xi \in SE(3), \quad \xi \in \mathbb{R}^6$$

Optimize dense photo-consistency:

$$\min_{\xi \in \mathbb{R}^6} \int_\Omega \left| I_0(x) - I_i\Big(\pi\big(g_\xi(u \cdot x)\big)\Big) \right| dx$$

Steinbruecker et al. ICCV '11, Kerl et al., ICRA '13

Realtime 3D Modeling

Color input

Depth input

Realtime 3D Modeling

Realtime 3D Modeling

Realtime 3D Modeling

Realtime 3D Modeling

Realtime 3D Modeling

Realtime 3D Modeling

Realtime 3D Modeling

Full-Body Scanner

Overview

Multiview reconstruction

Super-res.textures

Stereo reconstruction

Realtime dense geometry

RGB-D cameras

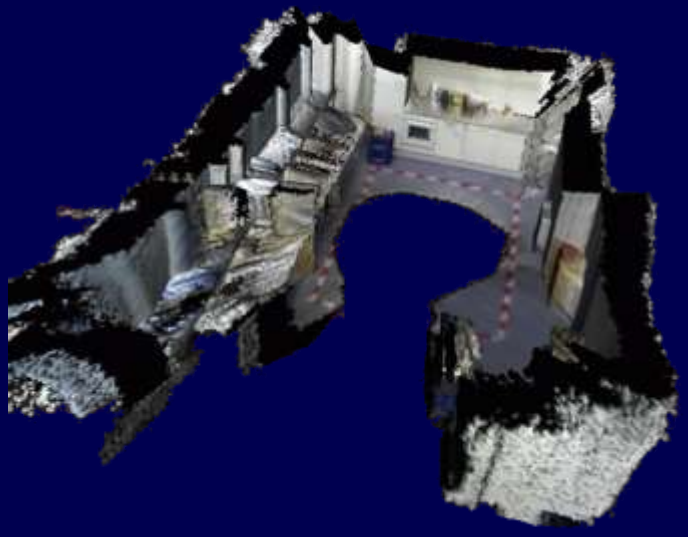

Reconstruction on the fly

Reconstruction on the Fly

Bylow, Sturm, Kerl, Kahl, Cremers RSS '13

Large Scale: Loop Closure

Kerl, Sturm, Cremers ICRA '13

Large Scale: Octrees

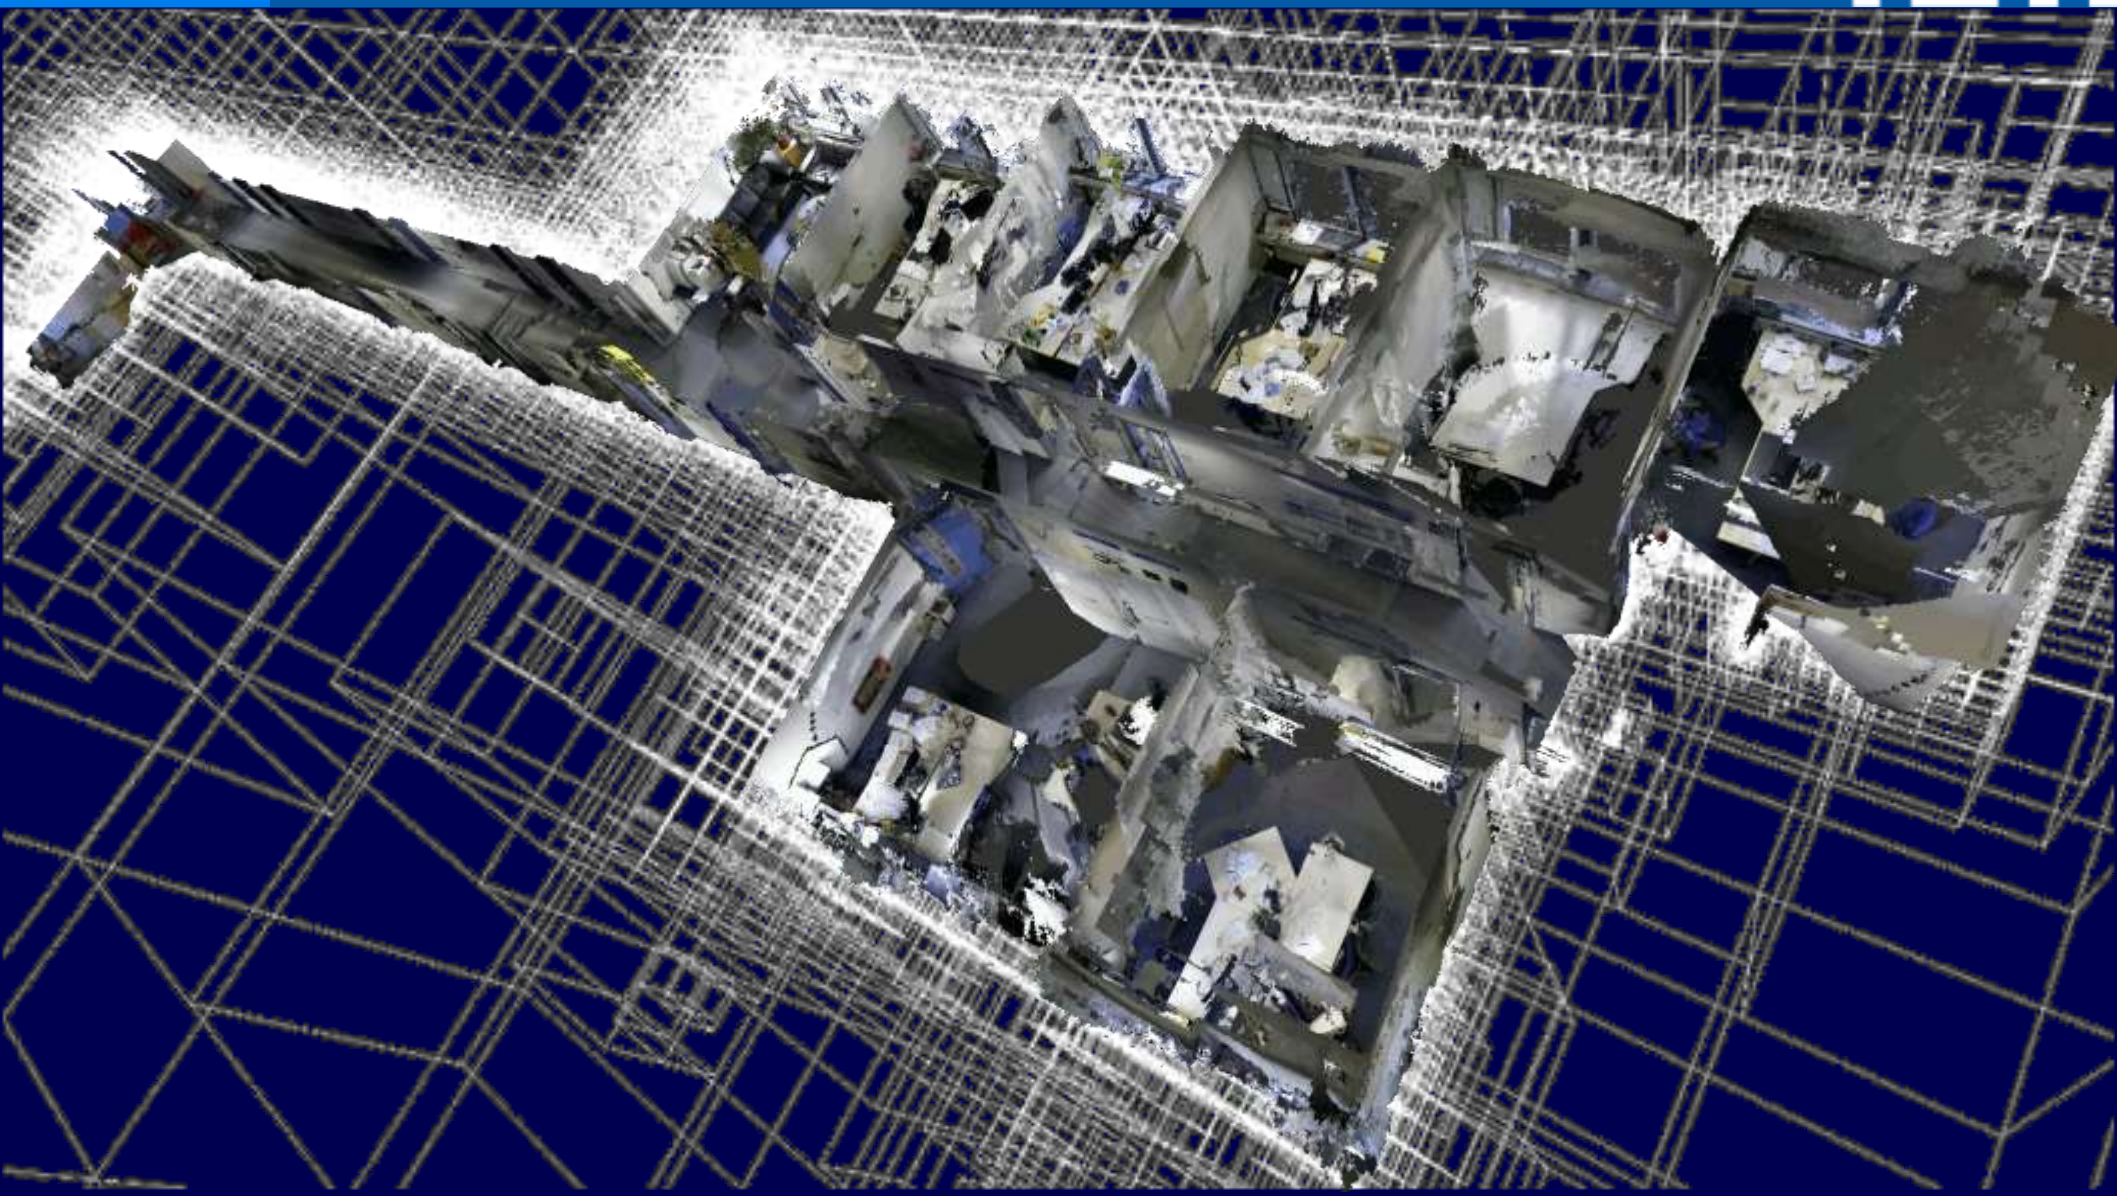

Steinbrücker, Kerl, Sturm, Cremers ICCV '13

Realtime Large-Scale Reconstruction

Steinbrücker, Kerl, Sturm, Cremers ICCV '13, ICRA '14

Summary

multiview reconstruction

super-res. textures

action reconstruction

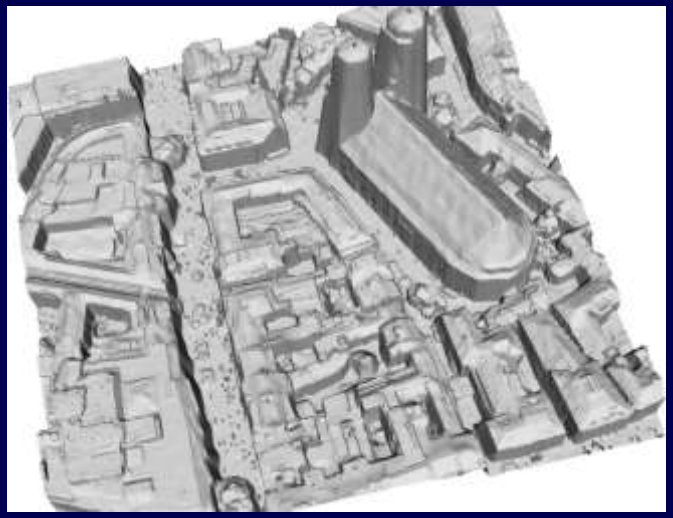

stereo reconstruction

RGB-D modeling

3D on the fly