Variational Methods for Computer Vision
ICCV Tutorial, 6.11.2011

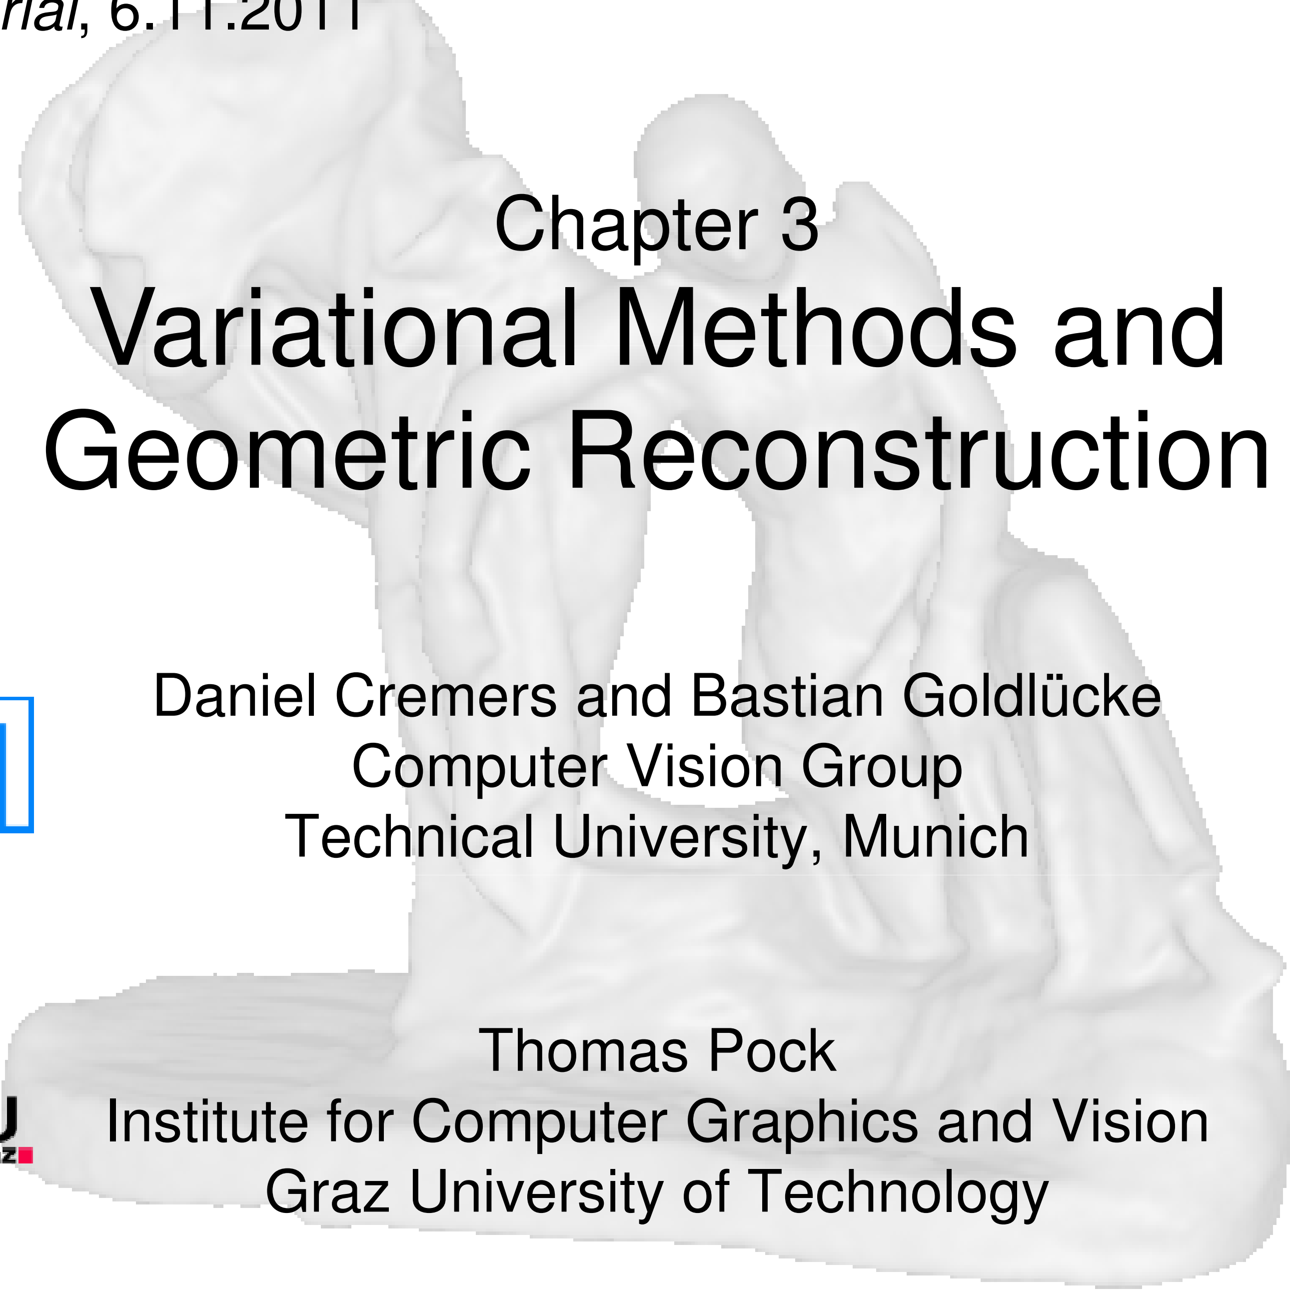

Chapter 3

Variational Methods and Geometric Reconstruction

Daniel Cremers and Bastian Goldlücke
Computer Vision Group
Technical University, Munich

Thomas Pock
Institute for Computer Graphics and Vision
Graz University of Technology

Mutliple view reconstruction

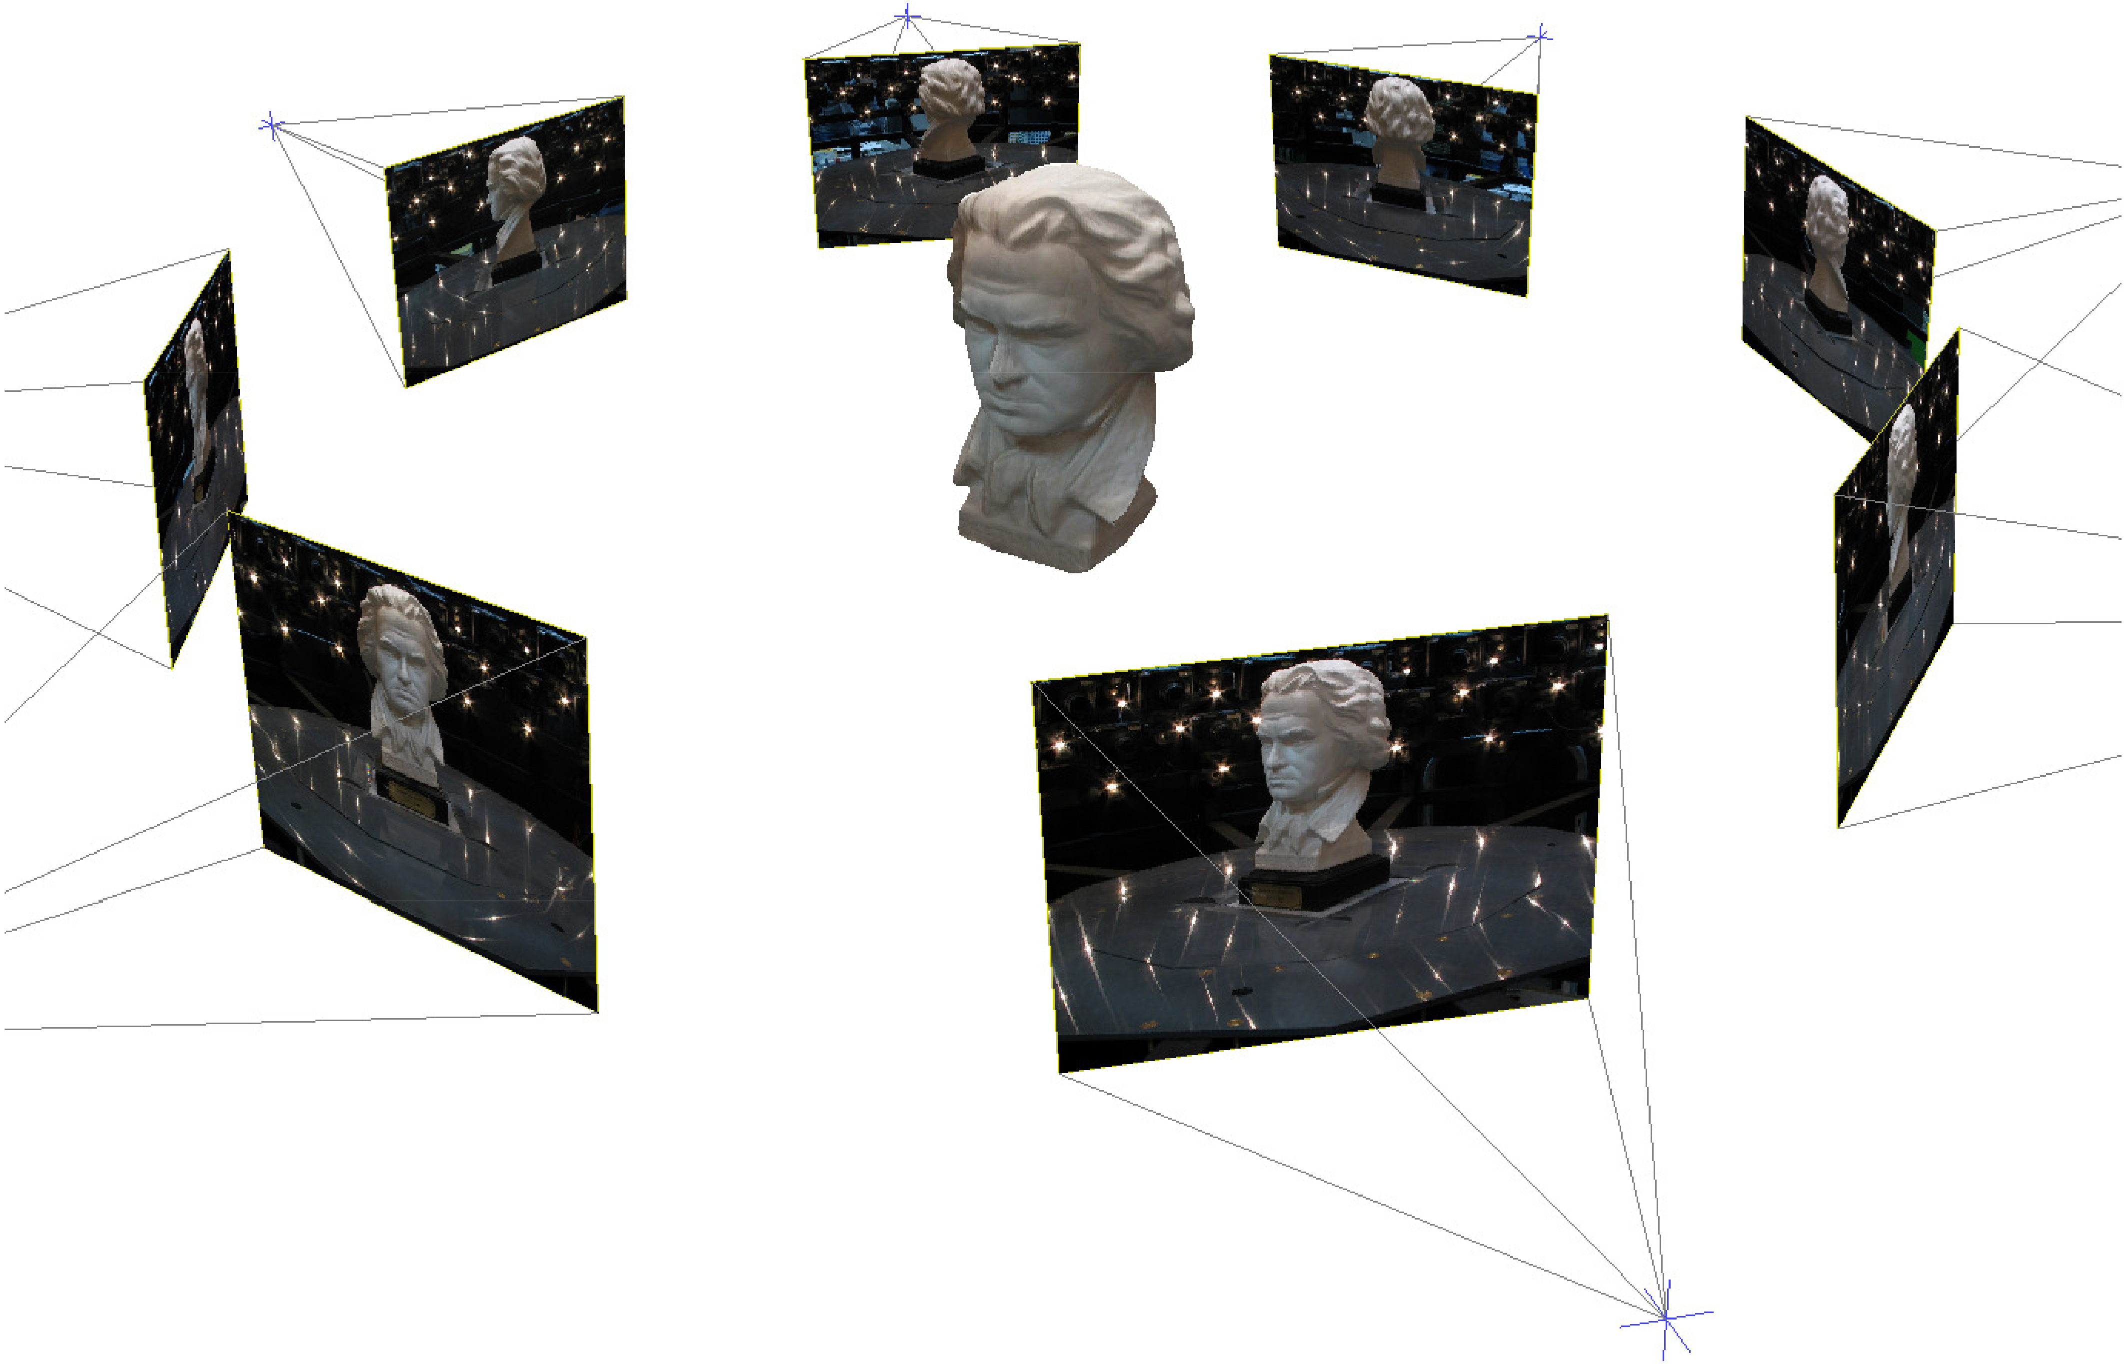

Variational methods and convexity

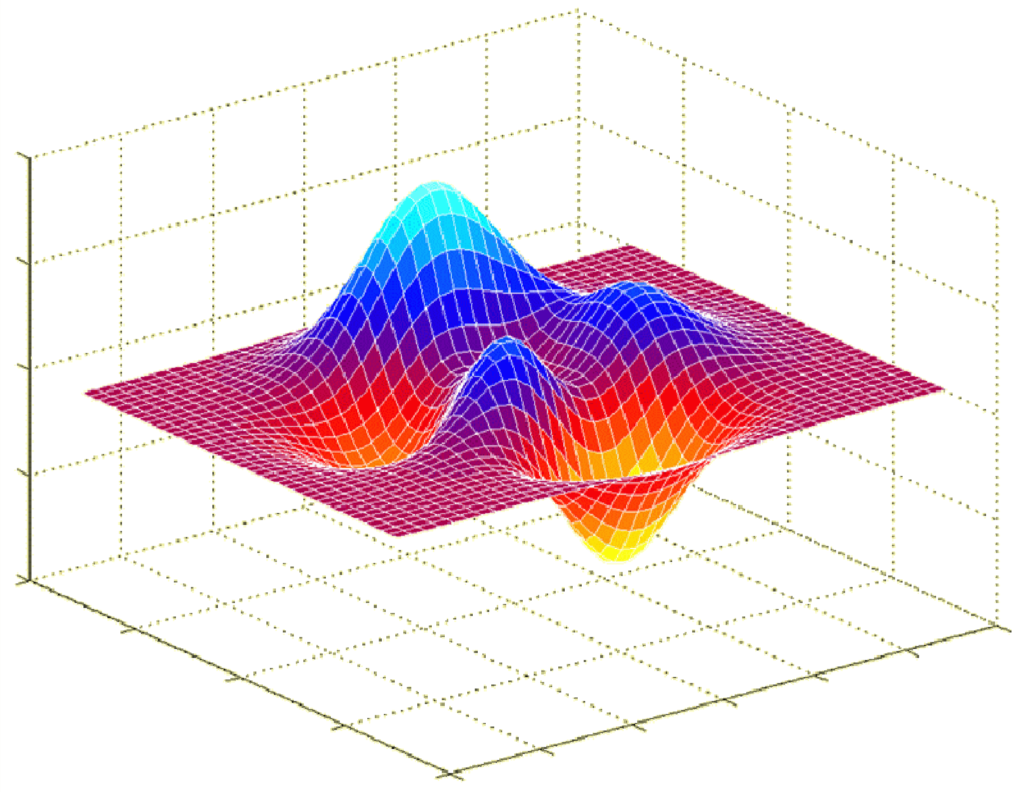

Non-convex energy

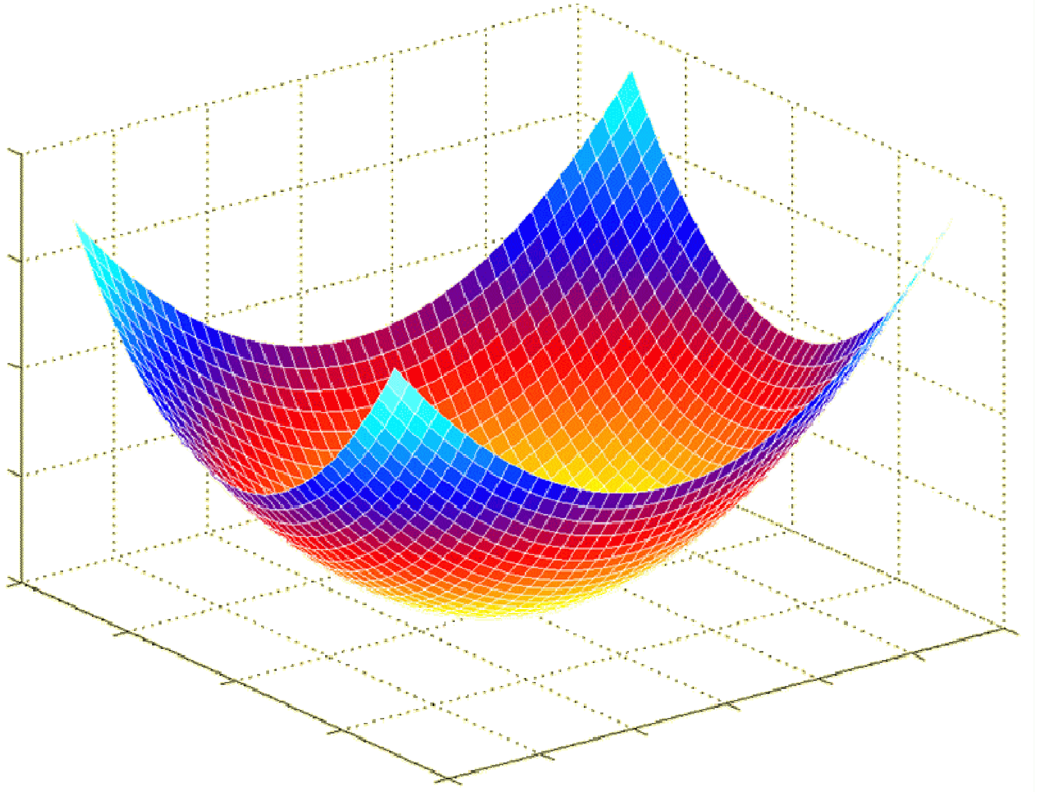

Convex energy

Overview

Multiview reconstruction

Super-resolution textures

Stereo & silhouettes

Single view reconstruction

Overview

Multiview reconstruction

Super-resolution textures

Stereo & silhouettes

Single view reconstruction

Stereo-weighted minimal surfaces

$$\rho : \left(V \subset \mathbb{R}^3\right) \to [0,1]$$

$$E(S) = \int_S \rho(s)\, ds$$

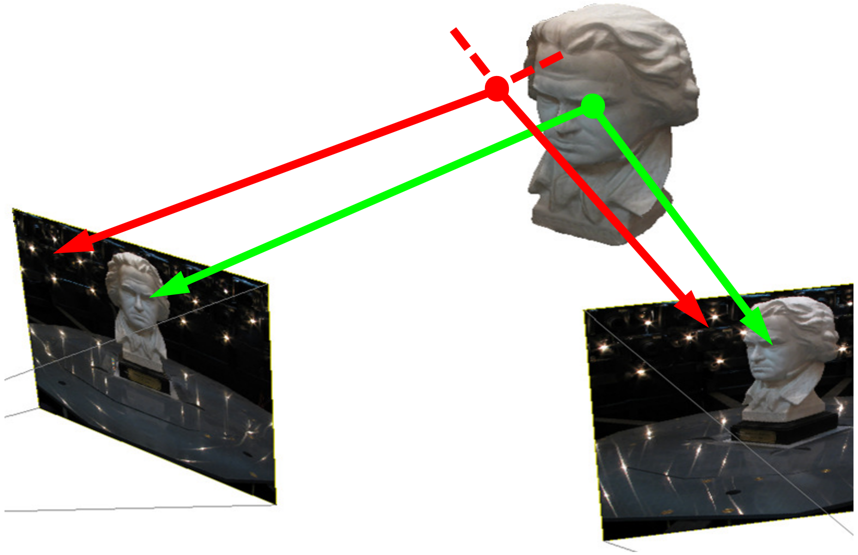

3D Reconstruction: *Faugeras, Keriven '98,* *Duan et al. '04*

Segmentation: *Kichenassamy et al. '95,* *Caselles et al. '95*

Optimal solution is the empty set: $\arg\min_S E(S) = \emptyset$

Resort:
- Local optimization: *Faugeras, Keriven TIP '98*
- Generative object/background modeling: *Yezzi, Soatto '03,…*
- Constrain search space: *Vogiatsis, Torr, Cipolla CVPR '05*
- Intelligent ballooning: *Boykov, Lempitsky BMVC '06*

Solution 1: Volumetric photoconsistency

$$E(S) = \int_{S_{in}} \rho_{obj}(x)dx \; + \int_{S_{out}} \rho_{bck}(x)dx \; + \int_S \rho(s)ds$$

$$\rho_{obj},\ \rho_{bck} : V \to [0, 1]$$

$$S_{in},\ S_{out} \subset V$$

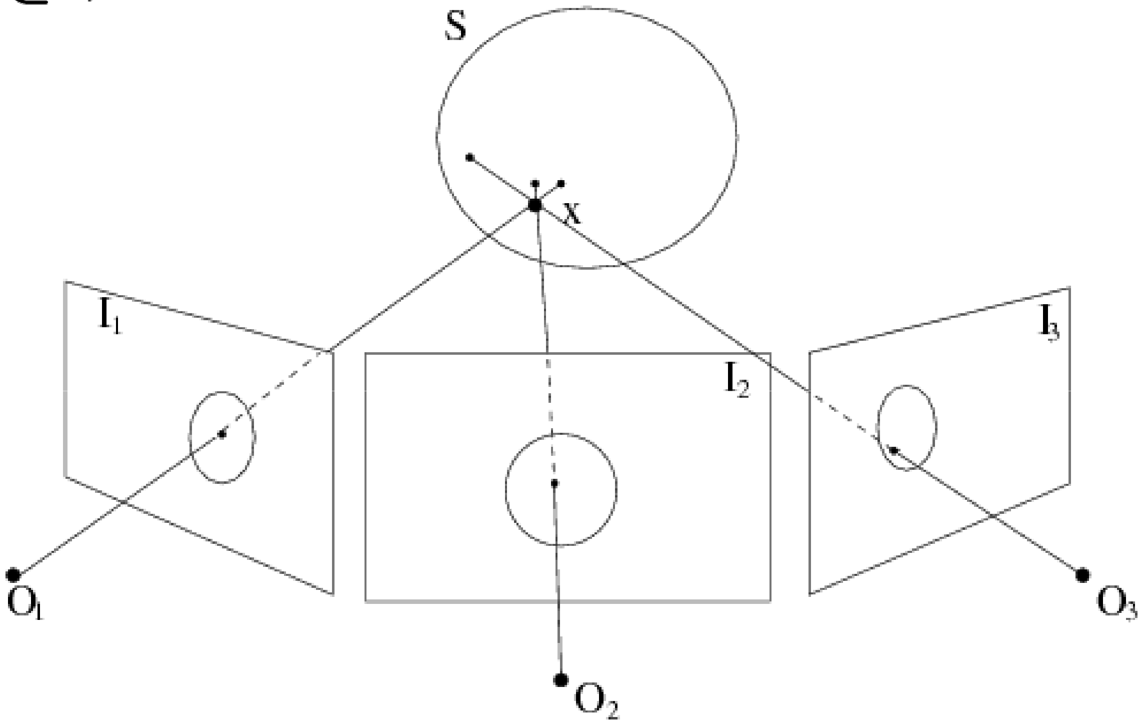

Kolev, Klodt, Brox, Cremers, IJCV '09

Global optima via convex relaxation

$$E(S) = \int_{S_{in}} \rho_{obj}\, dx \; + \int_{S_{out}} \rho_{bck}\, dx \; + \int_S \rho\, ds$$

implicit representation

$$E(u) = \int_V \rho_{obj}(x)(1 - u(x)) \; + \; \rho_{bck}(x) u(x)\, dx \; + \int_V \rho(x) |\nabla u|\, dx, \quad \text{s. t. } u : V \to \{0, 1\}$$

relaxation

$$E(u) = \int_V \rho_{obj}(x)(1 - u(x)) \; + \; \rho_{bck}(x) u(x)\, dx \; + \int_V \rho(x) |\nabla u|\, dx, \quad \text{s. t. } u : V \to [0, 1]$$

Global optima via convex relaxation

$$E(u) = \int_V \rho_{obj}\,(1 - u(x)) \; + \; \rho_{bck}\, u(x)\, dx \; + \; \int_V \rho\, |\nabla u|\, dx, \qquad (*)$$
$$\text{s. t. } u : V \to [0, 1]$$

<u>Theorem:</u> Thresholding a minimizer u^* of the relaxed problem $(*)$ leads to an optimal solution of the original binary problem:

$$u_{opt}(x) = 1_{u^* \geq \mu}(x) = \begin{cases} 1, & \text{if } u^*(x) \geq \mu \\ 0, & \text{if } u^*(x) < \mu \end{cases}$$

for any threshold $\mu \in (0, 1)$.

Chan, Esedoglu, Nikolova, TIP ’06

Kolev, Klodt, Brox, Esedoglu, Cremers, EMMCVPR ’07, IJCV ’09

A thresholding theorem

Let

$$u^* : \Omega \to [0,1]$$

be a (real-valued) minimizer of

$$E(u) = \int_\Omega f\,u + |\nabla u|\,dx.$$

Then for any threshold $\mu \in (0,1)$, the binary function

$$1_{u^* \geq \mu}(x)$$

is a global minimizer of the original binary problem.

A thresholding theorem

Proof: 1) $$u(x) = \int_0^1 1_{u \geq \mu}(x)\, d\mu$$ (layer cake formula)

2) $$\int_\Omega |\nabla u|\, dx = \int_0^1 \int_\Omega |\nabla 1_{u \geq \mu}(x)|\, dx\, d\mu$$ (coarea formula)

$$E(u) = \int_\Omega fu + |\nabla u|\, dx = \int_0^1 \int_\Omega f 1_{u \geq \mu} + |\nabla 1_{u \geq \mu}|\, dx\, d\mu = \int_0^1 E\big(1_{u \geq \mu}\big) d\mu$$

If $1_{u^* \geq \mu}$ is not minimizer, i.e. there exists a set $\Sigma \subset \Omega$

$$E(1_\Sigma) < E(1_{u^* \geq \mu})$$

$$\Rightarrow E(1_\Sigma) = \int_0^1 E(1_\Sigma)\, d\mu < \int_0^1 E(1_{u^* \geq \mu})\, d\mu = E(u^*)$$ (i.e. u^* not minimizer)

Reconstruction results

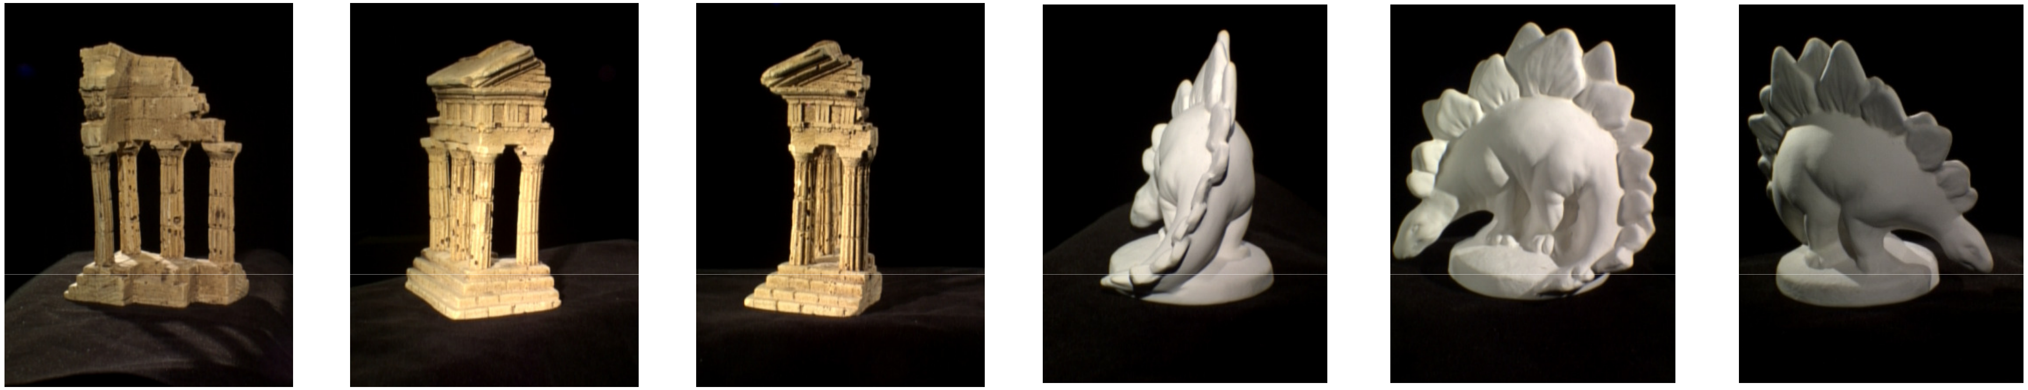

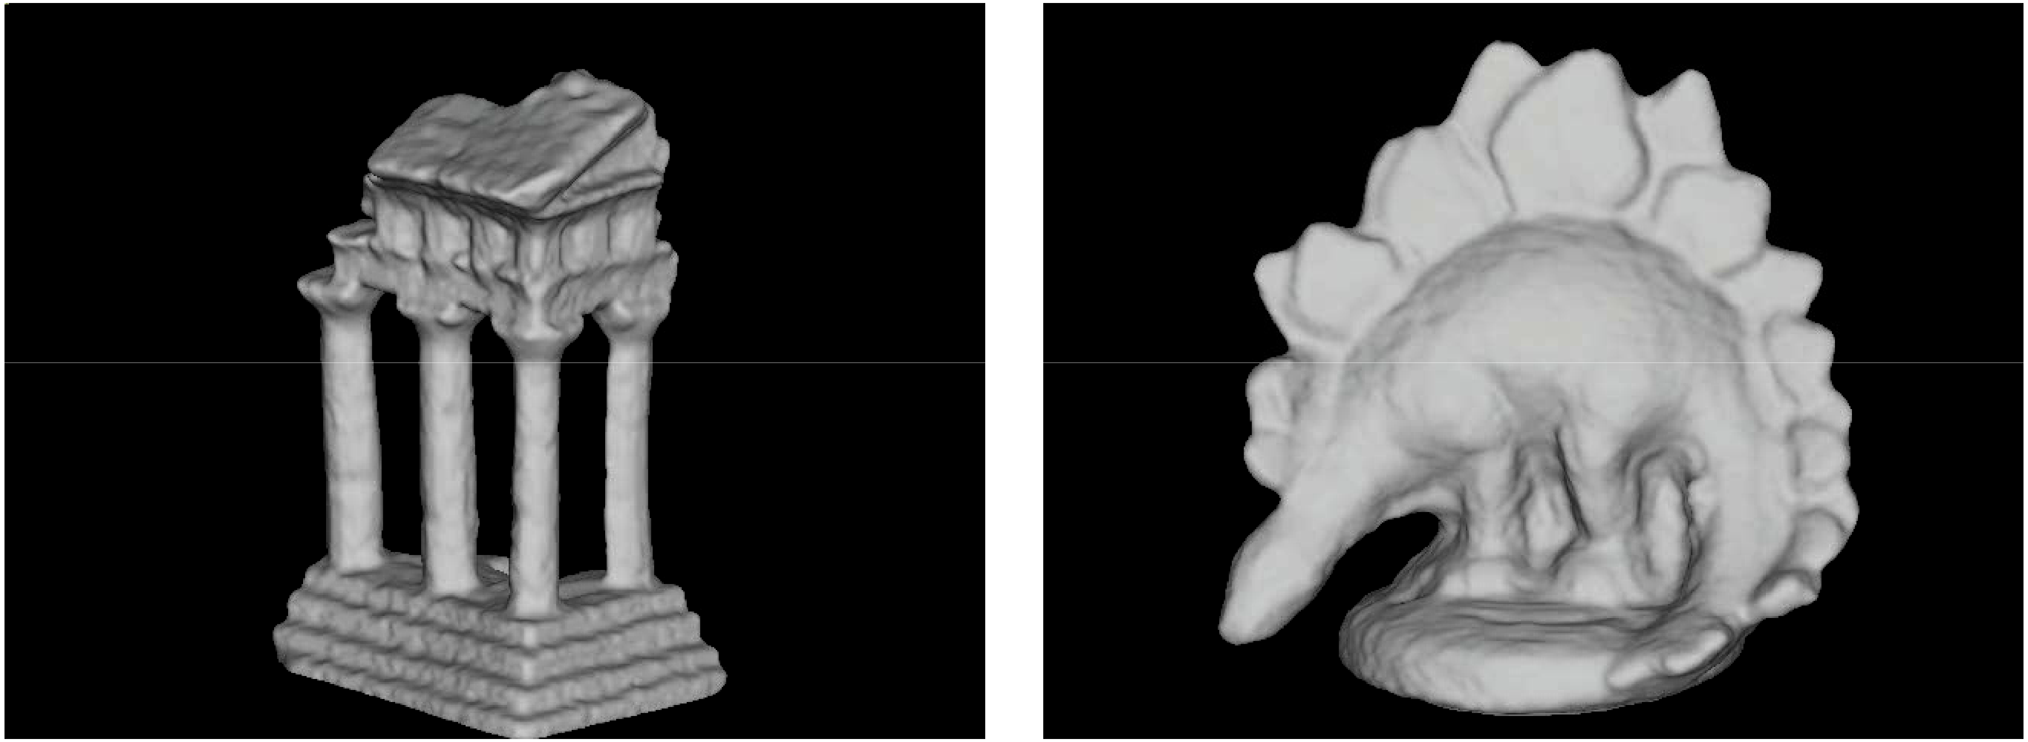

Discrete versus continuous optimization

Level sets	**Convex formulation**	**Graph cuts**
Continuous Local optima	Continuous Global optima*	Discrete Global optima*

* for certain functionals only

Klodt, Schoenmann, Kolev, Schikora, Cremers, ECCV '08

Discrete versus continuous optimization

Level sets	**Convex formulation**	**Graph cuts**
Continuous Local optima	Continuous Global optima* Parallel implementations	Discrete Global optima*

* for certain functionals only

Klodt, Schoenmann, Kolev, Schikora, Cremers, ECCV ’08

Discrete versus continuous optimization

Level sets	**Convex formulation**	**Graph cuts**
Continuous Local optima	Continuous Global optima* Parallel implementations	Discrete Global optima* Memory limitations

* for certain functionals only

Klodt, Schoenmann, Kolev, Schikora, Cremers, ECCV ’08

Discrete versus continuous optimization

Level sets	Convex formulation	Graph cuts
Continuous Local optima	Continuous Global optima* Parallel implementations	Discrete Global optima* Memory limitations Metrication errors

* for certain functionals only

Klodt, Schoenmann, Kolev, Schikora, Cremers, ECCV ’08

Metrication errors and consistency

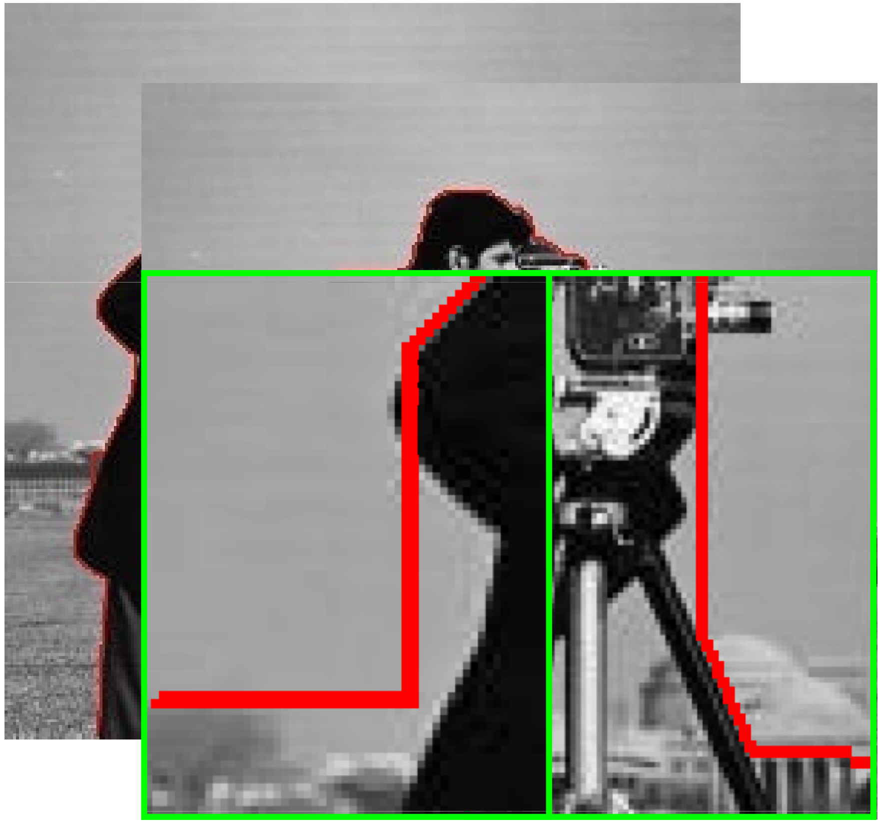

Discrete graph cut optimization
(4-connected grid)

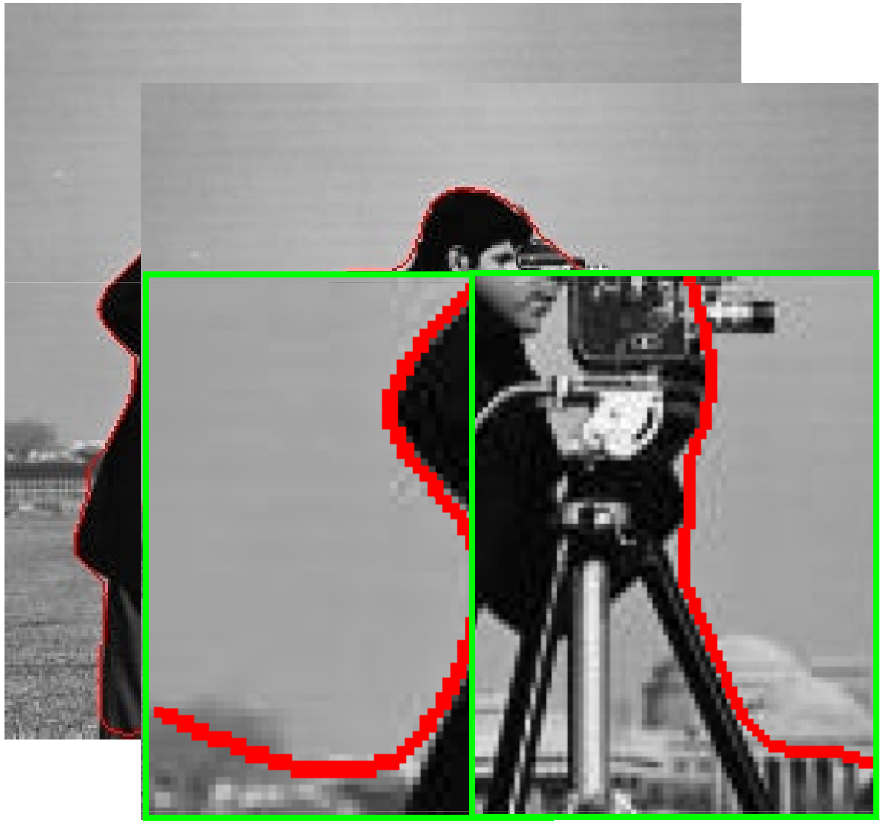

Continuous convex formulation
(4-connected grid)

Improvements:

Larger neighborhoods (*Boykov, Kolmogorov ’03*, *Kirsanov, Gortler ’04*)

Continuous Maximum Flow (*Appleton, Talbot ‘06*)

A minimal surface example: The catenoid

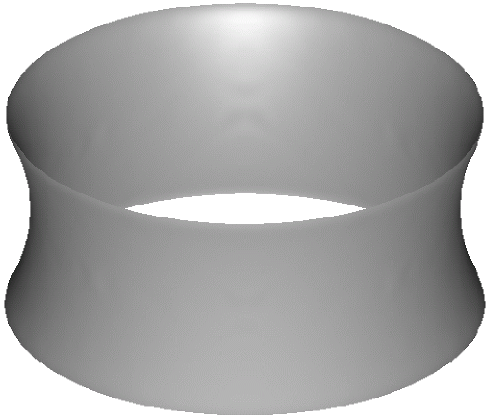

True solution

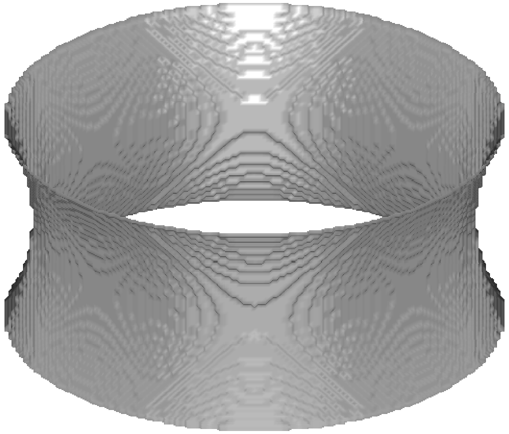

Convex formulation
(6-connected grid)

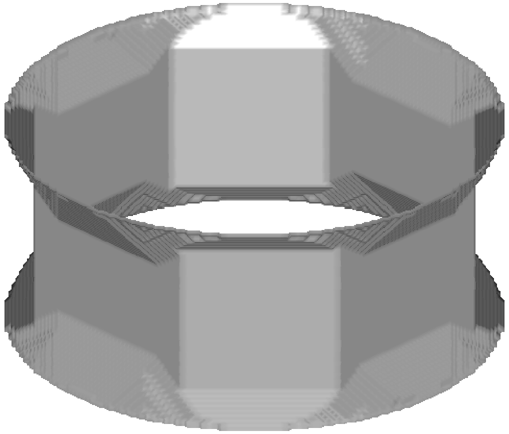

Graph cut
(26-connected grid)

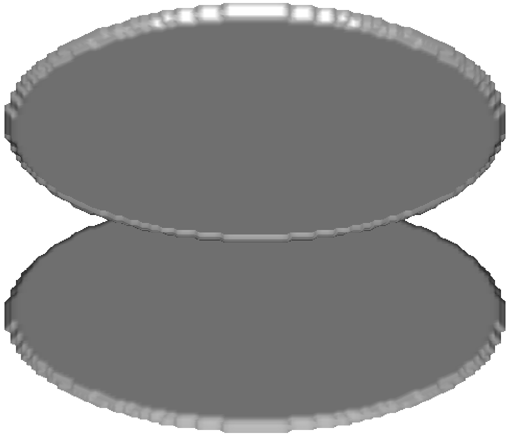

Graph cut
(6-connected grid)

Klodt, Schoenmann, Kolev, Schikora, Cremers, ECCV '08

Metrication errors and consistency

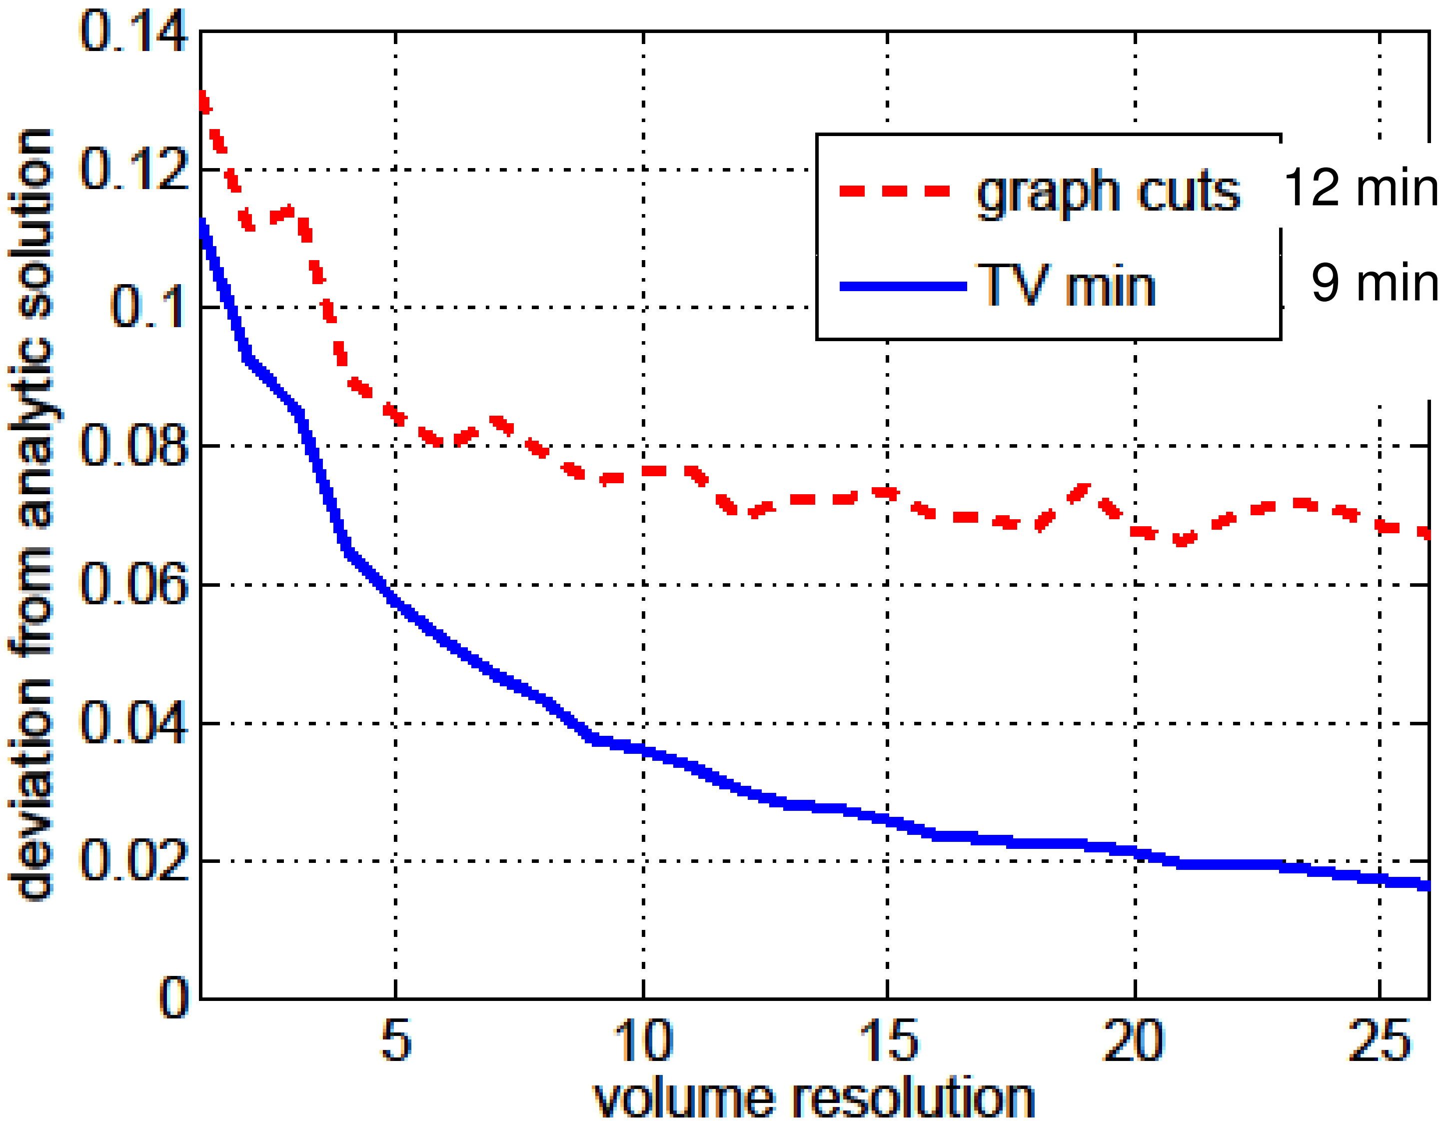

12 min. 32 sec. (CPU)

9 min. 34 sec. (CPU)

8 sec. (GPU)

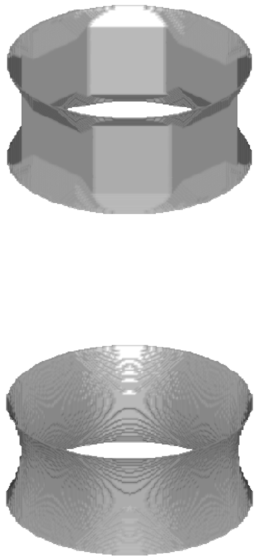

Klodt, Schoenmann, Kolev, Schikora, Cremers, ECCV '08

Overview

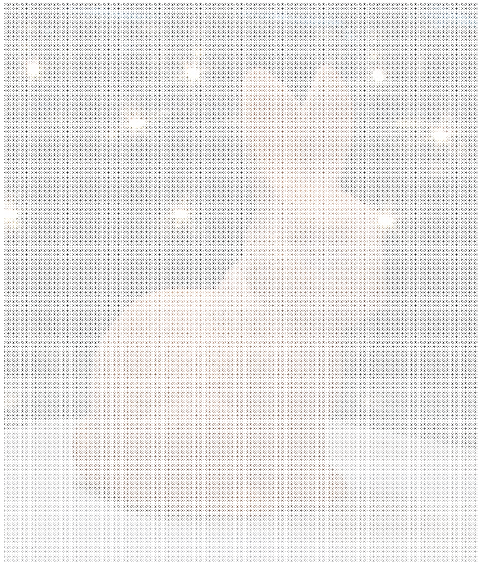

Multiview reconstruction

Super-resolution textures

Stereo & silhouettes

Single view reconstruction

Evolution to global optimum

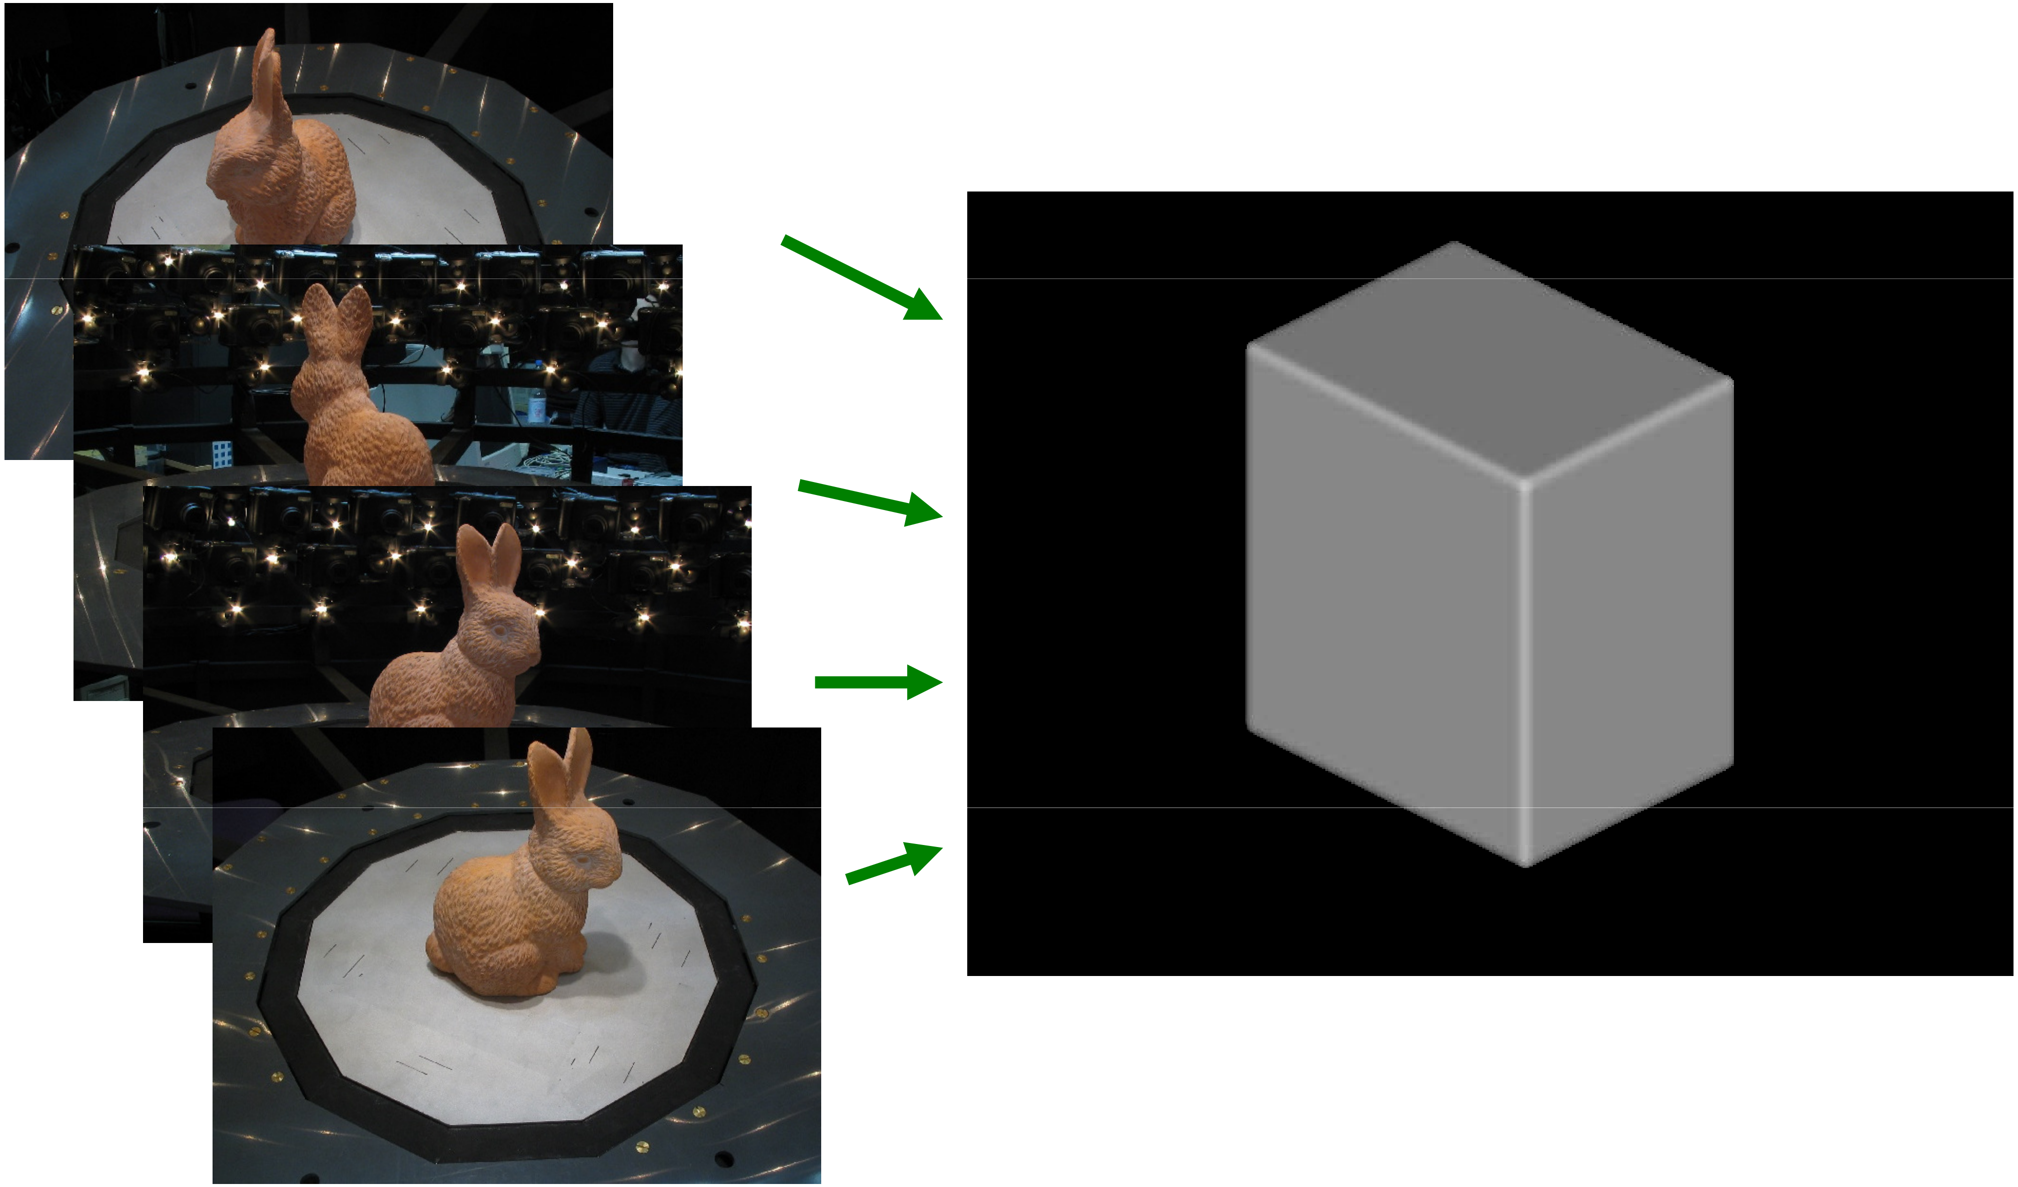

Super-resolution texture estimation

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

$$\min_T \sum_{i=1}^{n} \int_{\Omega_i} \Big(\underbrace{b \,*}_{\text{blur \& downsample}} (T \circ \underbrace{\beta_i}_{\text{back-projection}}) - \mathcal{I}_i \Big)^2 dx + \lambda \int_S \|\nabla_S T\| ds$$

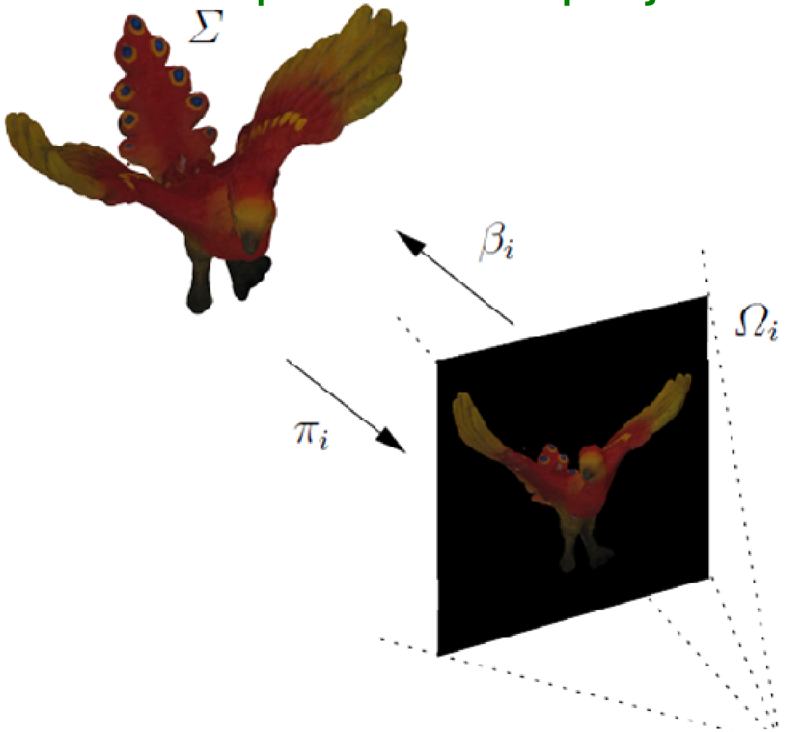

Goldlücke, Cremers, ICCV '09, DAGM '09

Super-resolution texture estimation

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

$$\min_T \sum_{i=1}^{n} \int_{\Omega_i} \Big(\underbrace{b \,*}_{\text{blur \& downsample}} (T \circ \underbrace{\beta_i}_{\text{back-projection}}) - \mathcal{I}_i \Big)^2 dx + \lambda \int_S \|\nabla_S T\| ds$$

Euler-Lagrange equation gives a PDE on the surface:

$$\frac{dE}{dT} = -\mathrm{div}_S \left(\frac{\nabla_S T}{\|\nabla_S T\|_S} \right) + \sum_{i=1}^{n} \frac{v_i}{\lambda} \big((\mathcal{J}_i \mathcal{D}_i) \circ \pi_i \big) = 0$$

where $\mathcal{D}_i = \bar{b} * (b * (T \circ \beta_i) - \mathcal{I})$ and $\mathcal{J}_i = \left\| \frac{\partial \beta_i}{\partial x} \times \frac{\partial \beta_i}{\partial y} \right\|^{-1}$.

Conformal parameterization of the surface → PDE on charts.

Goldlücke, Cremers, ICCV '09, DAGM '09

Joint estimation of geometry & super-resolution texture

Determine normal displacement $D : S \to \mathbb{R}$ and surface color $T : S \to \mathbb{R}^3$

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

geometry-dependent back-projection

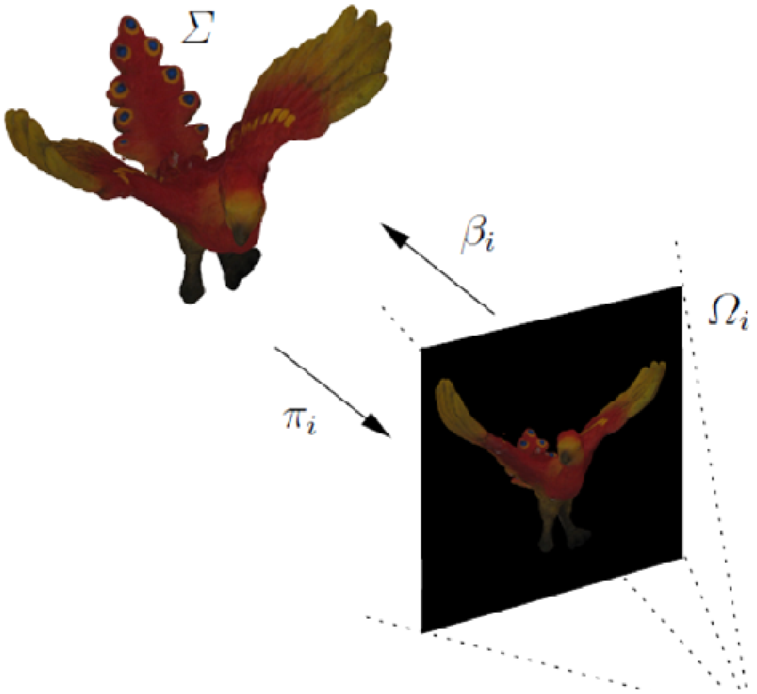

Goldlücke, Cremers, ICCV ’09, DAGM ’09

Super-resolution texture estimation

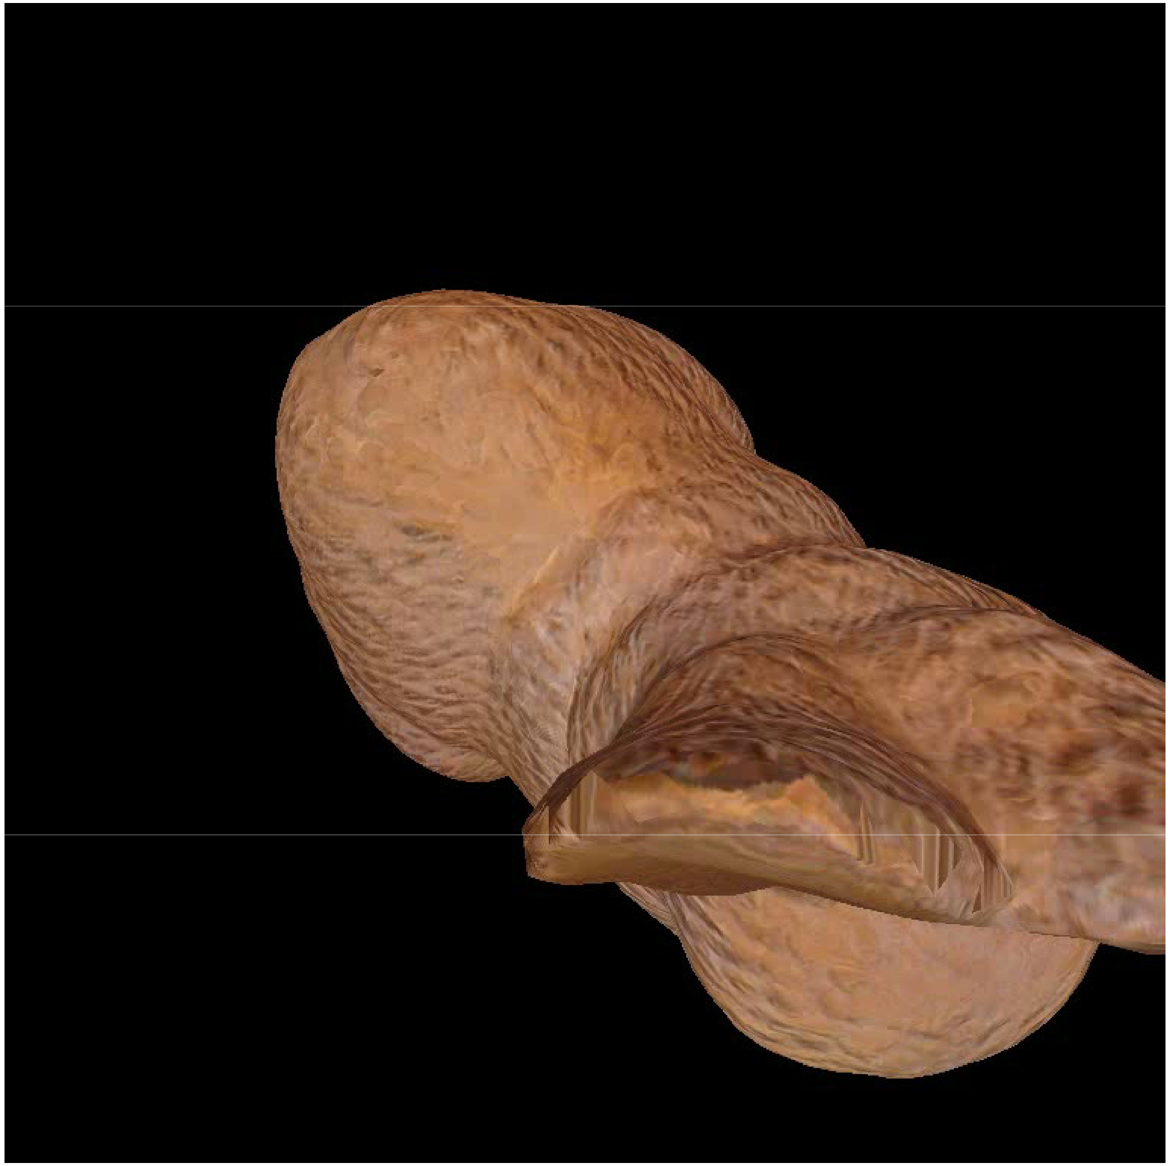

Goldlücke, Cremers, ICCV '09, DAGM '09

Evolution of the super-resolution process

Goldlücke, Cremers, ICCV '09, DAGM '09

Super-resolution texture estimation

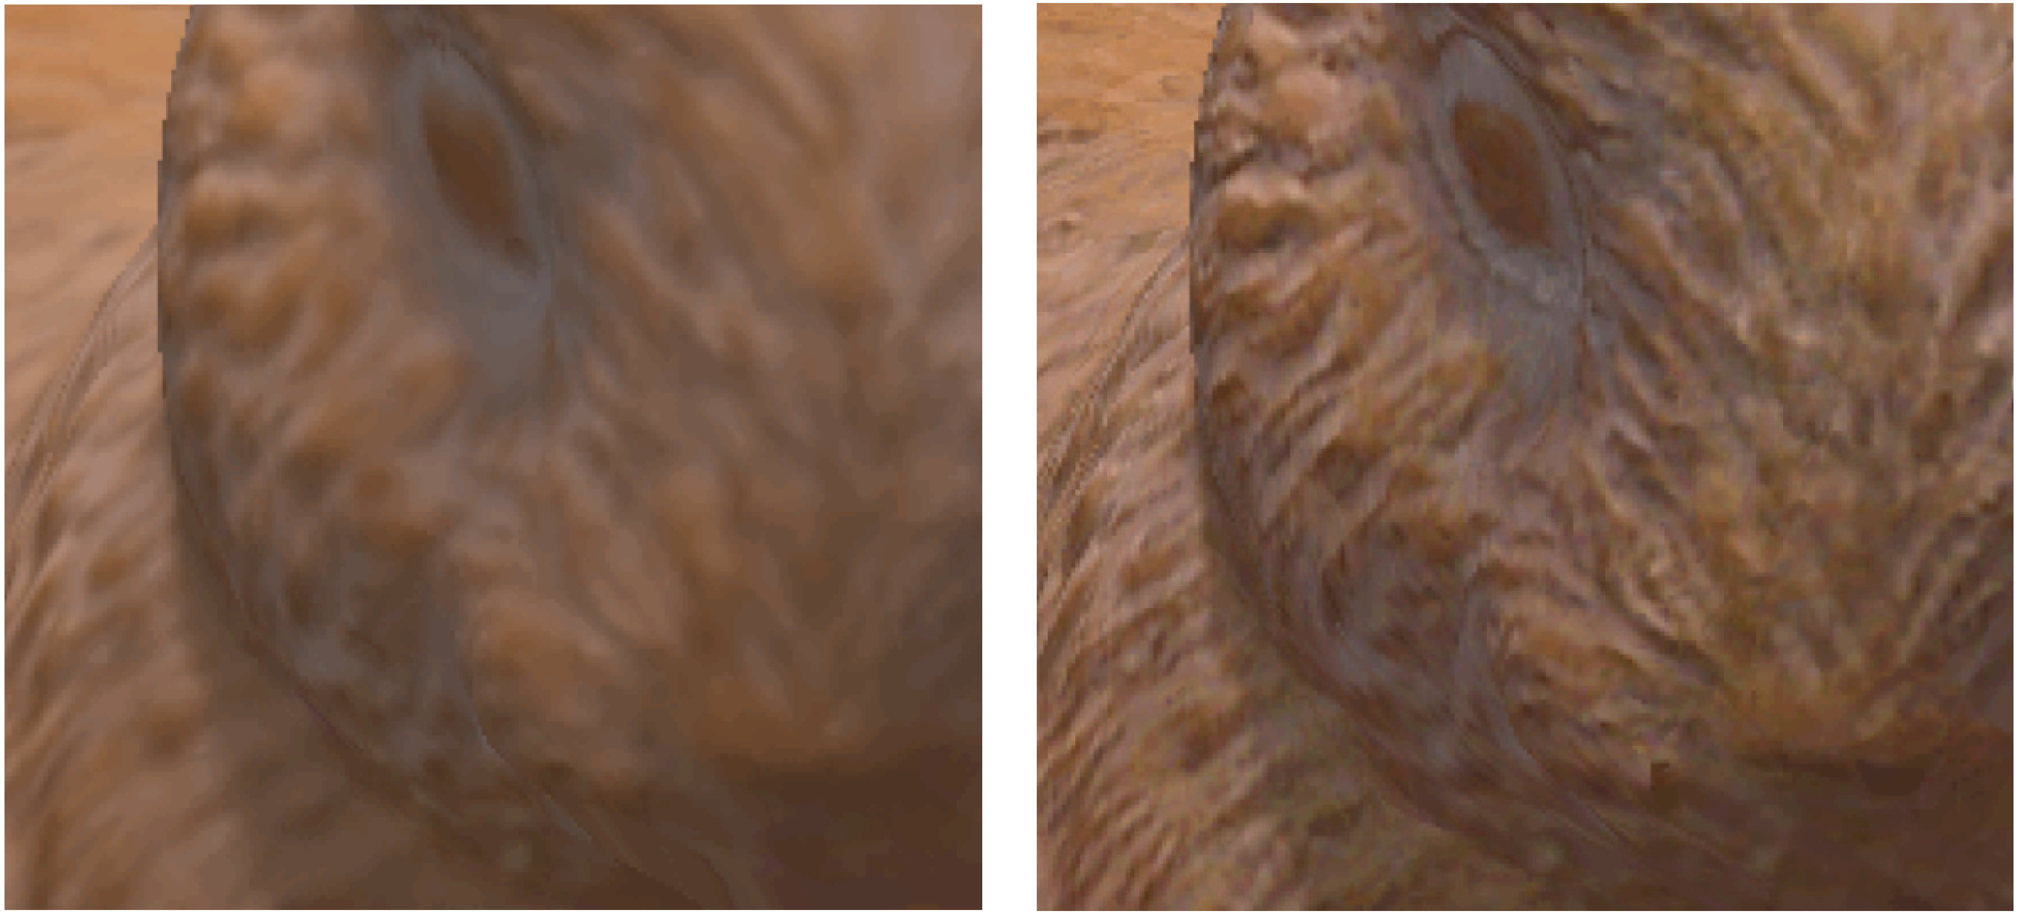

Weighted average

Super-resolution texture

Goldlücke, Cremers, ICCV '09, DAGM '09

Super-resolution texture estimation

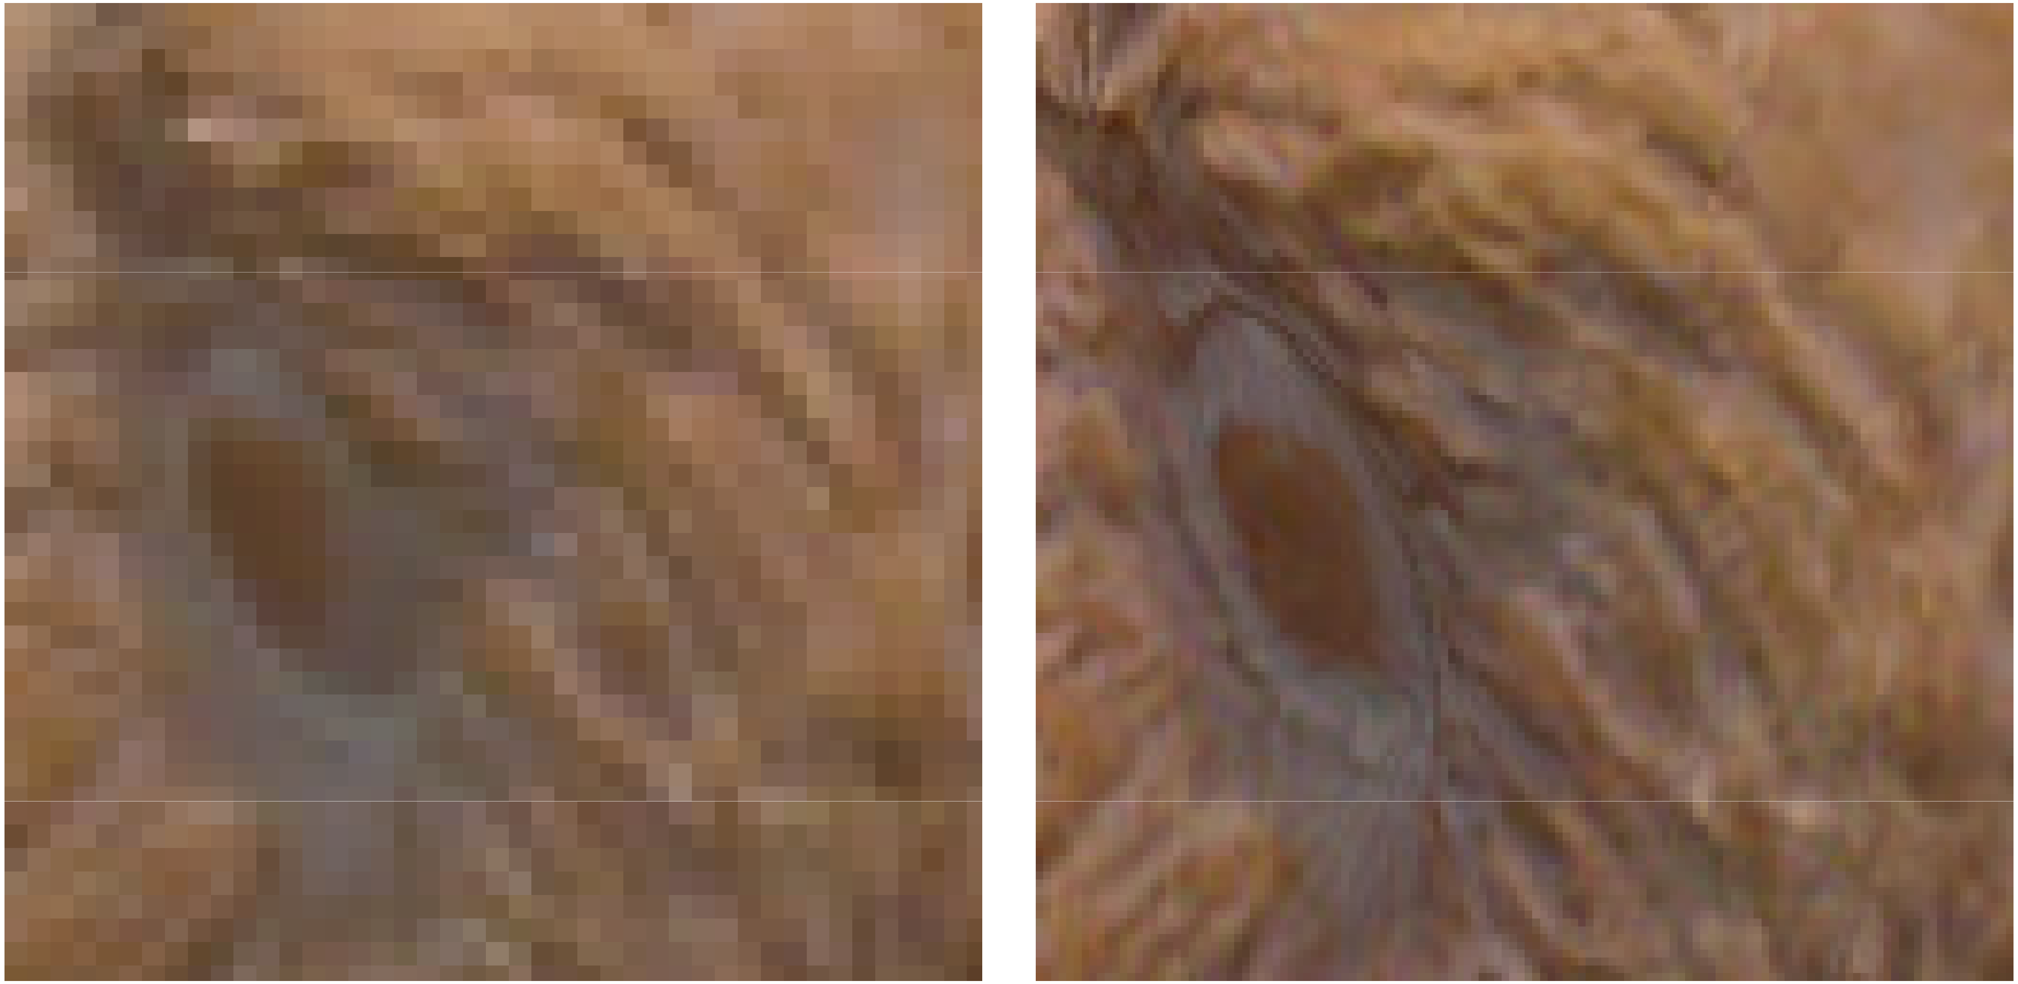

Input image

Super-resolution texture

Goldlücke, Cremers, ICCV ’09, DAGM ’09

Overview

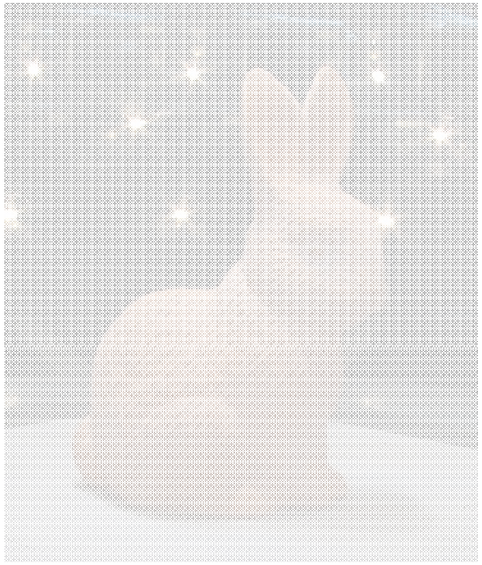

Multiview reconstruction

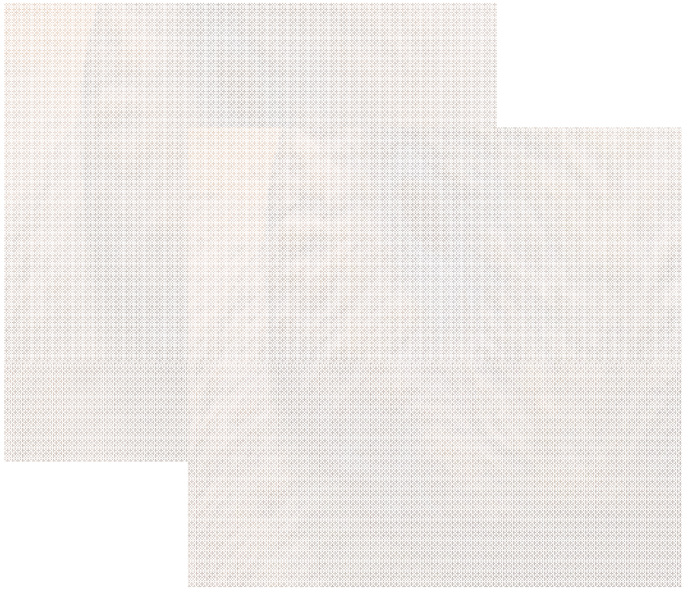

Super-resolution textures

Stereo & silhouettes

Single view reconstruction

Solution 2: Imposing silhouette consistency

$$\min_S \int_S \rho \, dS$$

$$\text{s. t.} \quad \pi_i(S) = S_i \;\; \forall\, i = 1, \dots, n$$

$$\pi_i : V \to \Omega_i$$

$$S_i \subset \Omega_i$$

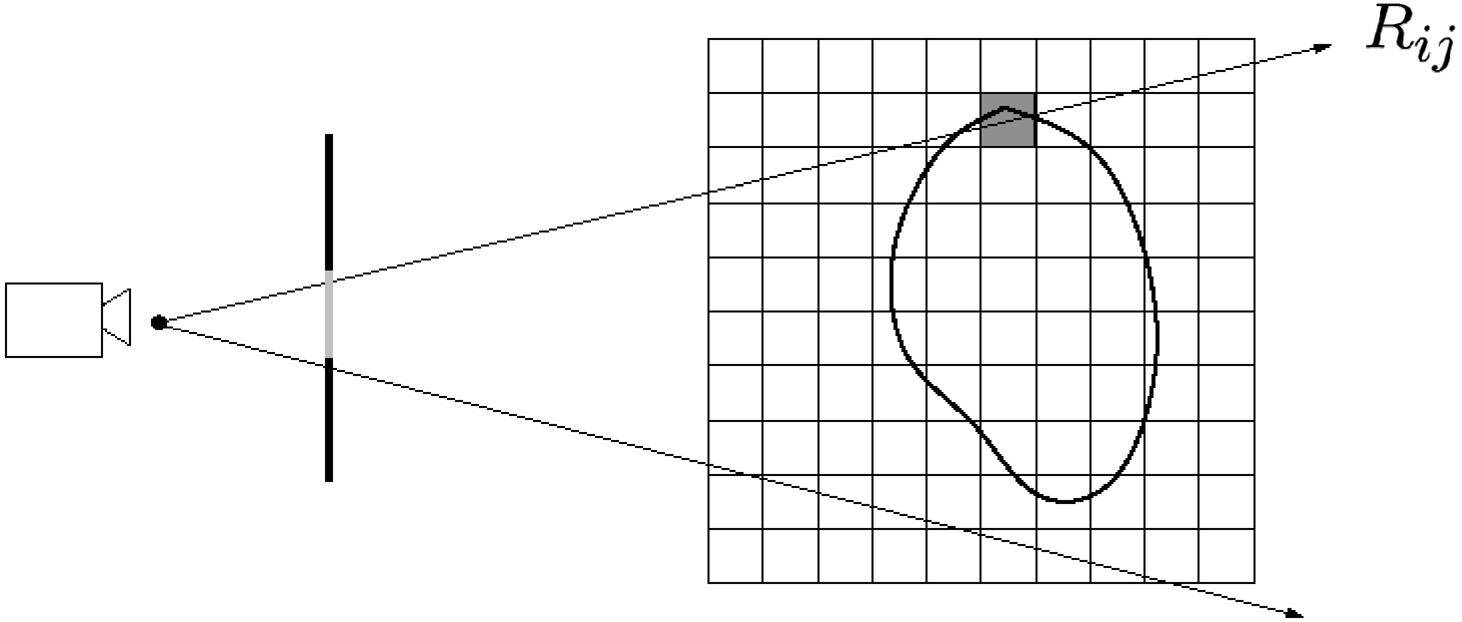

Kolev, Cremers, ECCV '08, PAMI 2010

Silhouette consistency as a convex constraint

$$E(S) = \int_S \rho(x)\, dS,$$

$$\text{s. t.} \qquad \pi_i(S) = S_i \quad \forall\, i = 1, \dots, n$$

$$\Big\downarrow$$

implicit representation & relaxation

$$E(u) = \int_V \rho(x) |\nabla u(x)|\, dx$$

$$\Sigma = \left\{ \begin{array}{l} \text{s. t.} \qquad \xcancel{u : V \to \{0,1\}} \; u : V \to [0,1] \\ \displaystyle\int_{R_{ij}} u(x)\, dx \geq \delta \quad \text{if } j \in S_i \\ \displaystyle\int_{R_{ij}} u(x)\, dx = 0 \quad \text{if } j \notin S_i \end{array} \right.$$

<u>Proposition:</u> The set Σ of silhouette-consistent solutions is convex.

Kolev, Cremers, ECCV ’08, PAMI 2010

Silhouette consistency as a convex constraint

$$E(u) = \int_V \rho(x) |\nabla u(x)| \, dx$$

$$\text{s. t.} \qquad u : V \to [0,1]$$

$$\int_{R_{ij}} u(x) \, dx \geq 1 \quad \text{if } j \in S_i$$

$$\int_{R_{ij}} u(x) \, dx = 0 \quad \text{if } j \notin S_i$$

$$\Big\downarrow$$

thresholding

$$R^S_{obj} = \{x \in V \mid u(x) > \mu\}$$

$$R^S_{bck} = \{x \in V \mid u(x) < \mu\} \text{, where}$$

$$\mu = \min\left\{\left(\min_{i \in \{1,\dots,n\}, j \in S_i} \; \max_{x \in R_{ij}} u^*(x)\right) \text{ , } 0.5\right\}$$

Kolev, Cremers, ECCV '08, PAMI 2010

Bounded optimality

$$u^\star = \arg\min_{u \in \Sigma} E(u) \qquad u' = \arg\min_{u \in \Sigma_{bin}} E(u)$$

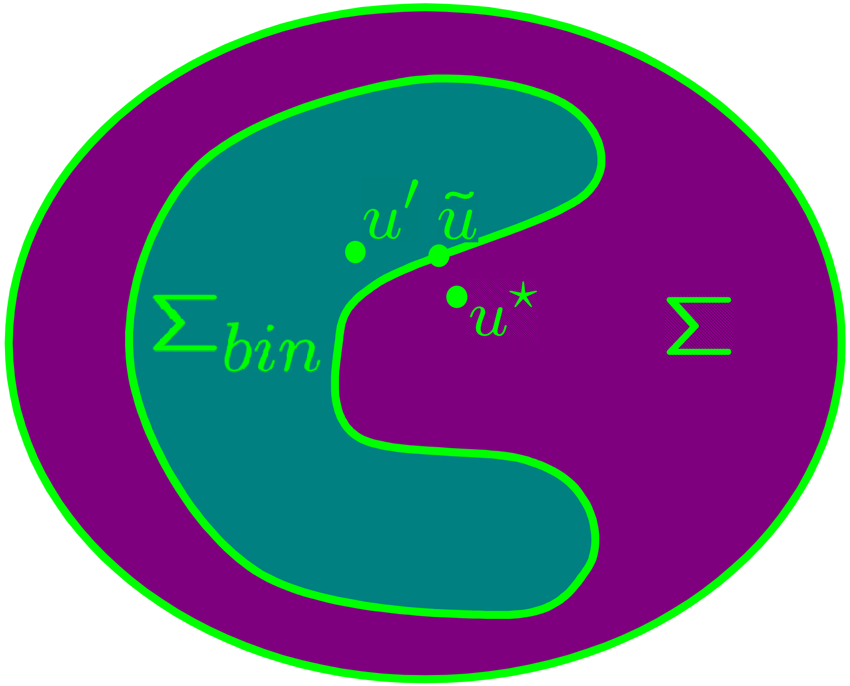

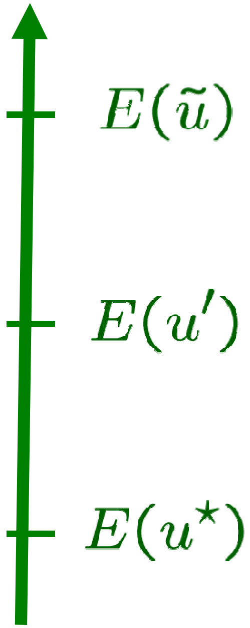

$$E(\tilde{u}) - E(u') \leq E(\tilde{u}) - E(u^\star)$$

Kolev, Cremers, ECCV '08, PAMI 2010

Numerical optimization via lagged diffusivity

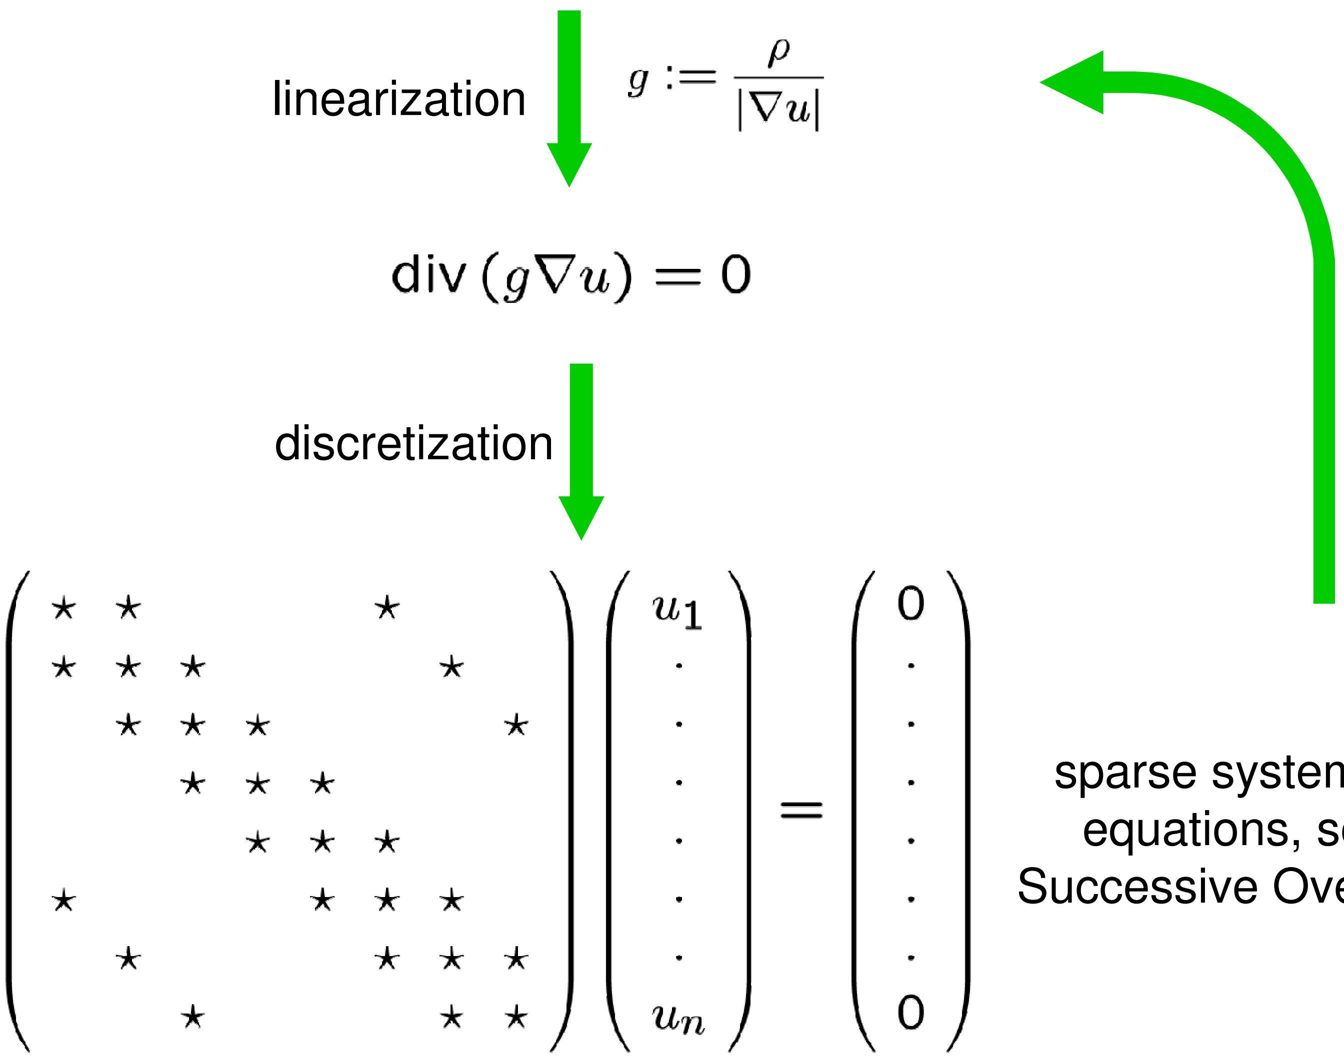

Euler-Lagrange equation

$$\operatorname{div}\left(\rho\frac{\nabla u}{|\nabla u|}\right) = 0$$

linearization

$$g := \frac{\rho}{|\nabla u|}$$

$$\operatorname{div}\left(g\nabla u\right) = 0$$

discretization

$$\begin{pmatrix} \star & \star & & & & \star & & \\ \star & \star & \star & & & & \star & \\ & \star & \star & \star & & & & \star \\ & & \star & \star & \star & & & \\ & & & \star & \star & \star & & \\ \star & & & & \star & \star & \star & \\ & \star & & & & \star & \star & \star \\ & & \star & & & & \star & \star \end{pmatrix} \begin{pmatrix} u_1 \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ u_n \end{pmatrix} = \begin{pmatrix} 0 \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ 0 \end{pmatrix}$$

sparse system of linear equations, solved by Successive Overrelaxation

Reconstruction of non-Lambertian objects

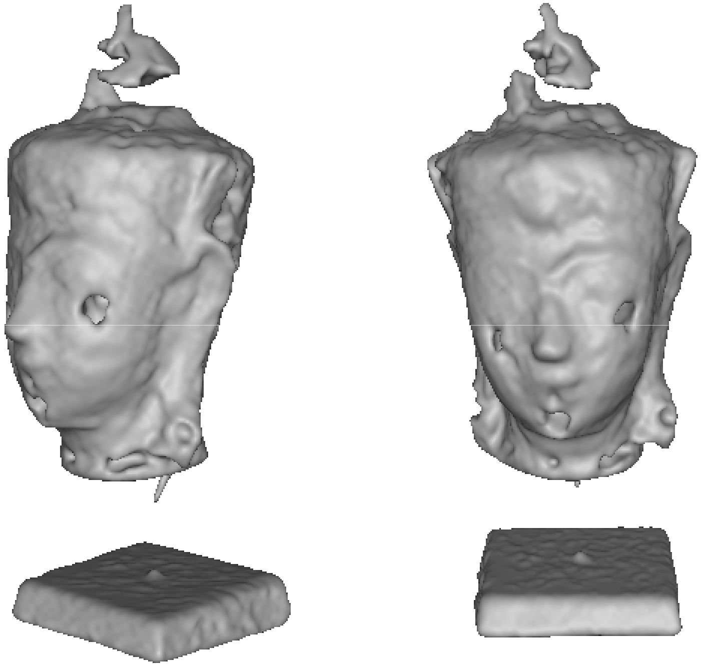

propagation scheme

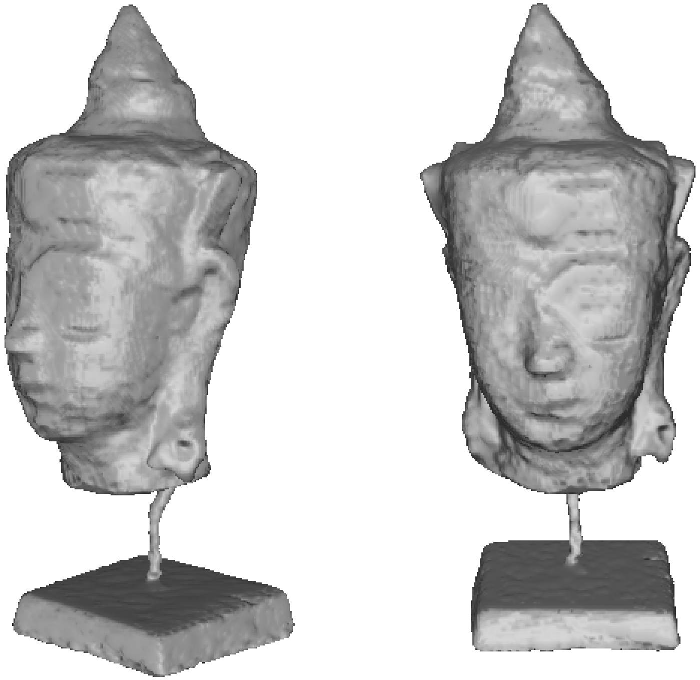

silhouette constraints

Reconstruction of low-textured objects

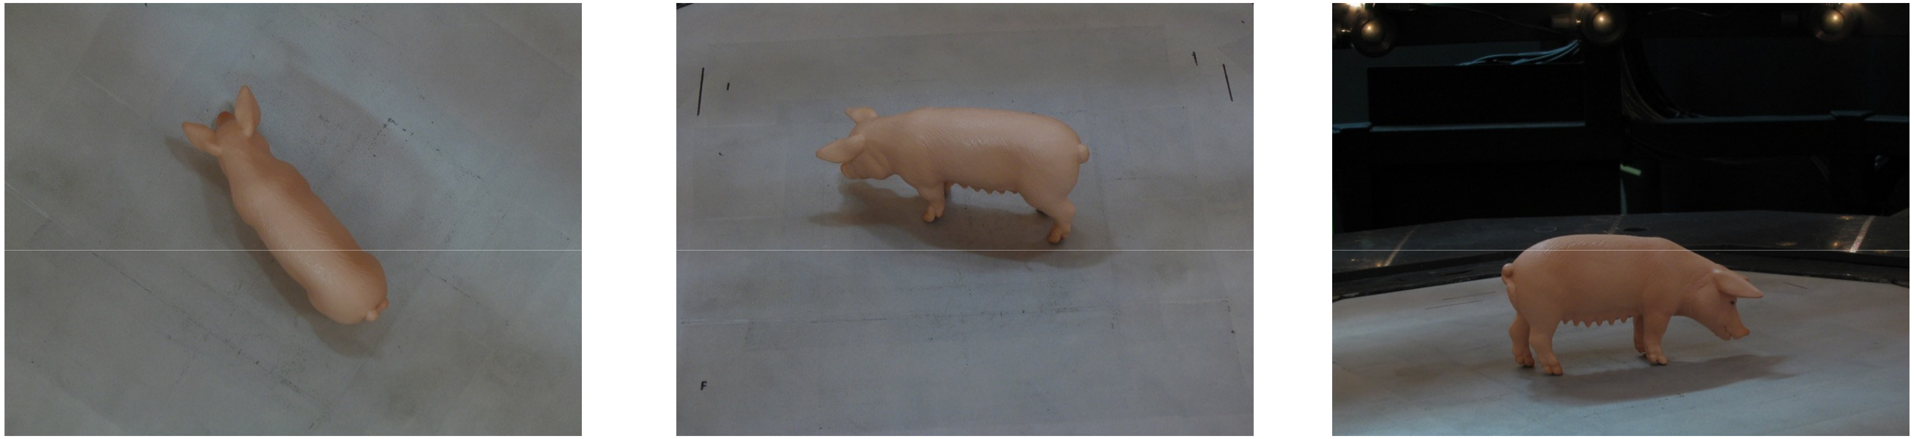

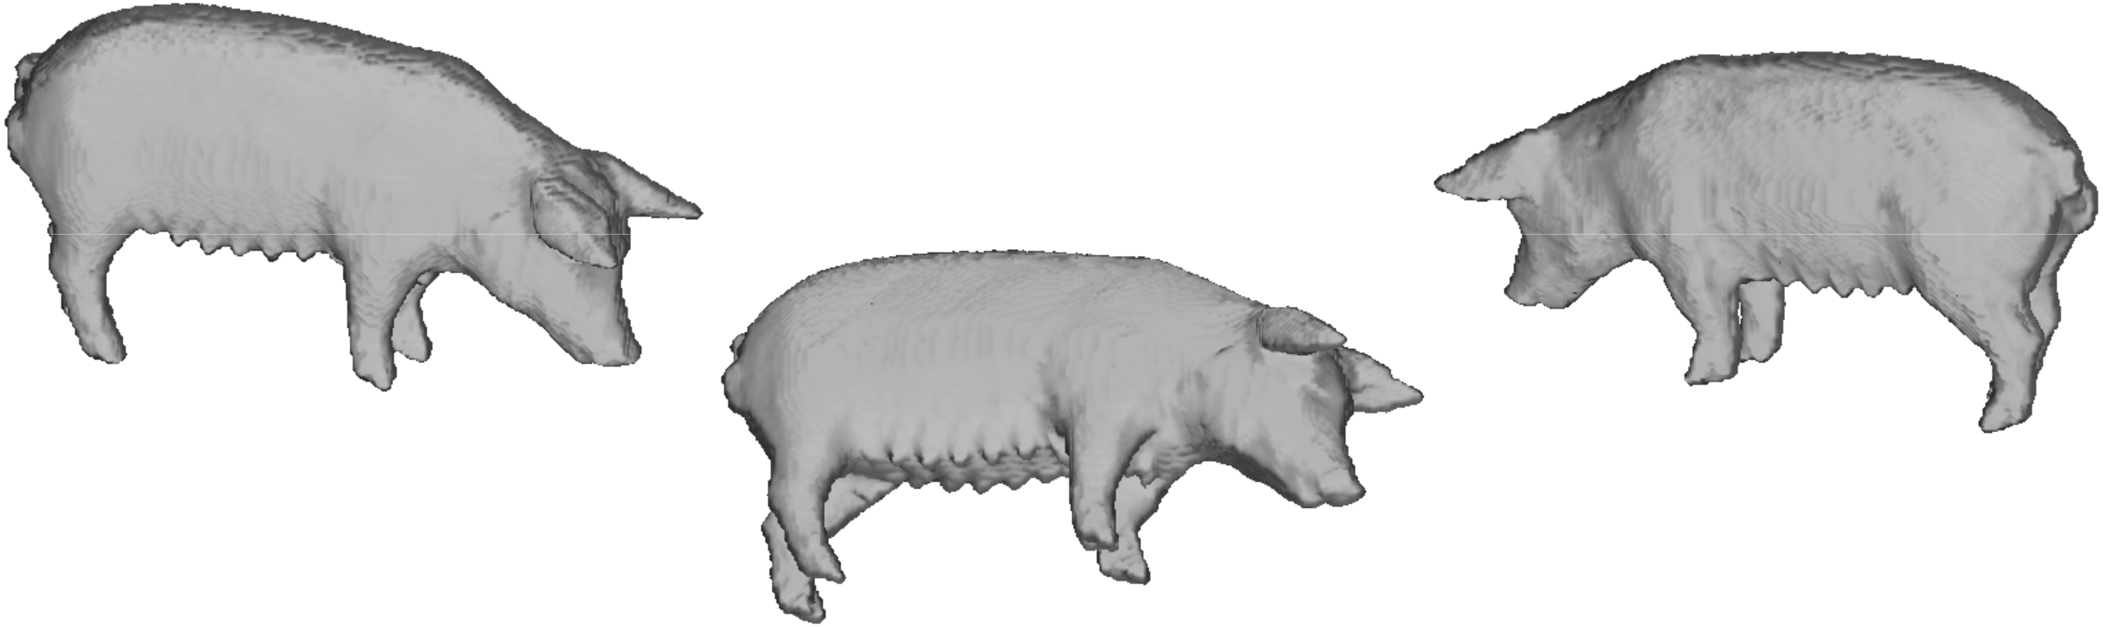

Reconstruction of fine-scale structures

Image data courtesy of Yasutaka Furukawa.

Reconstructing the Niobids statues (450 B.C.)

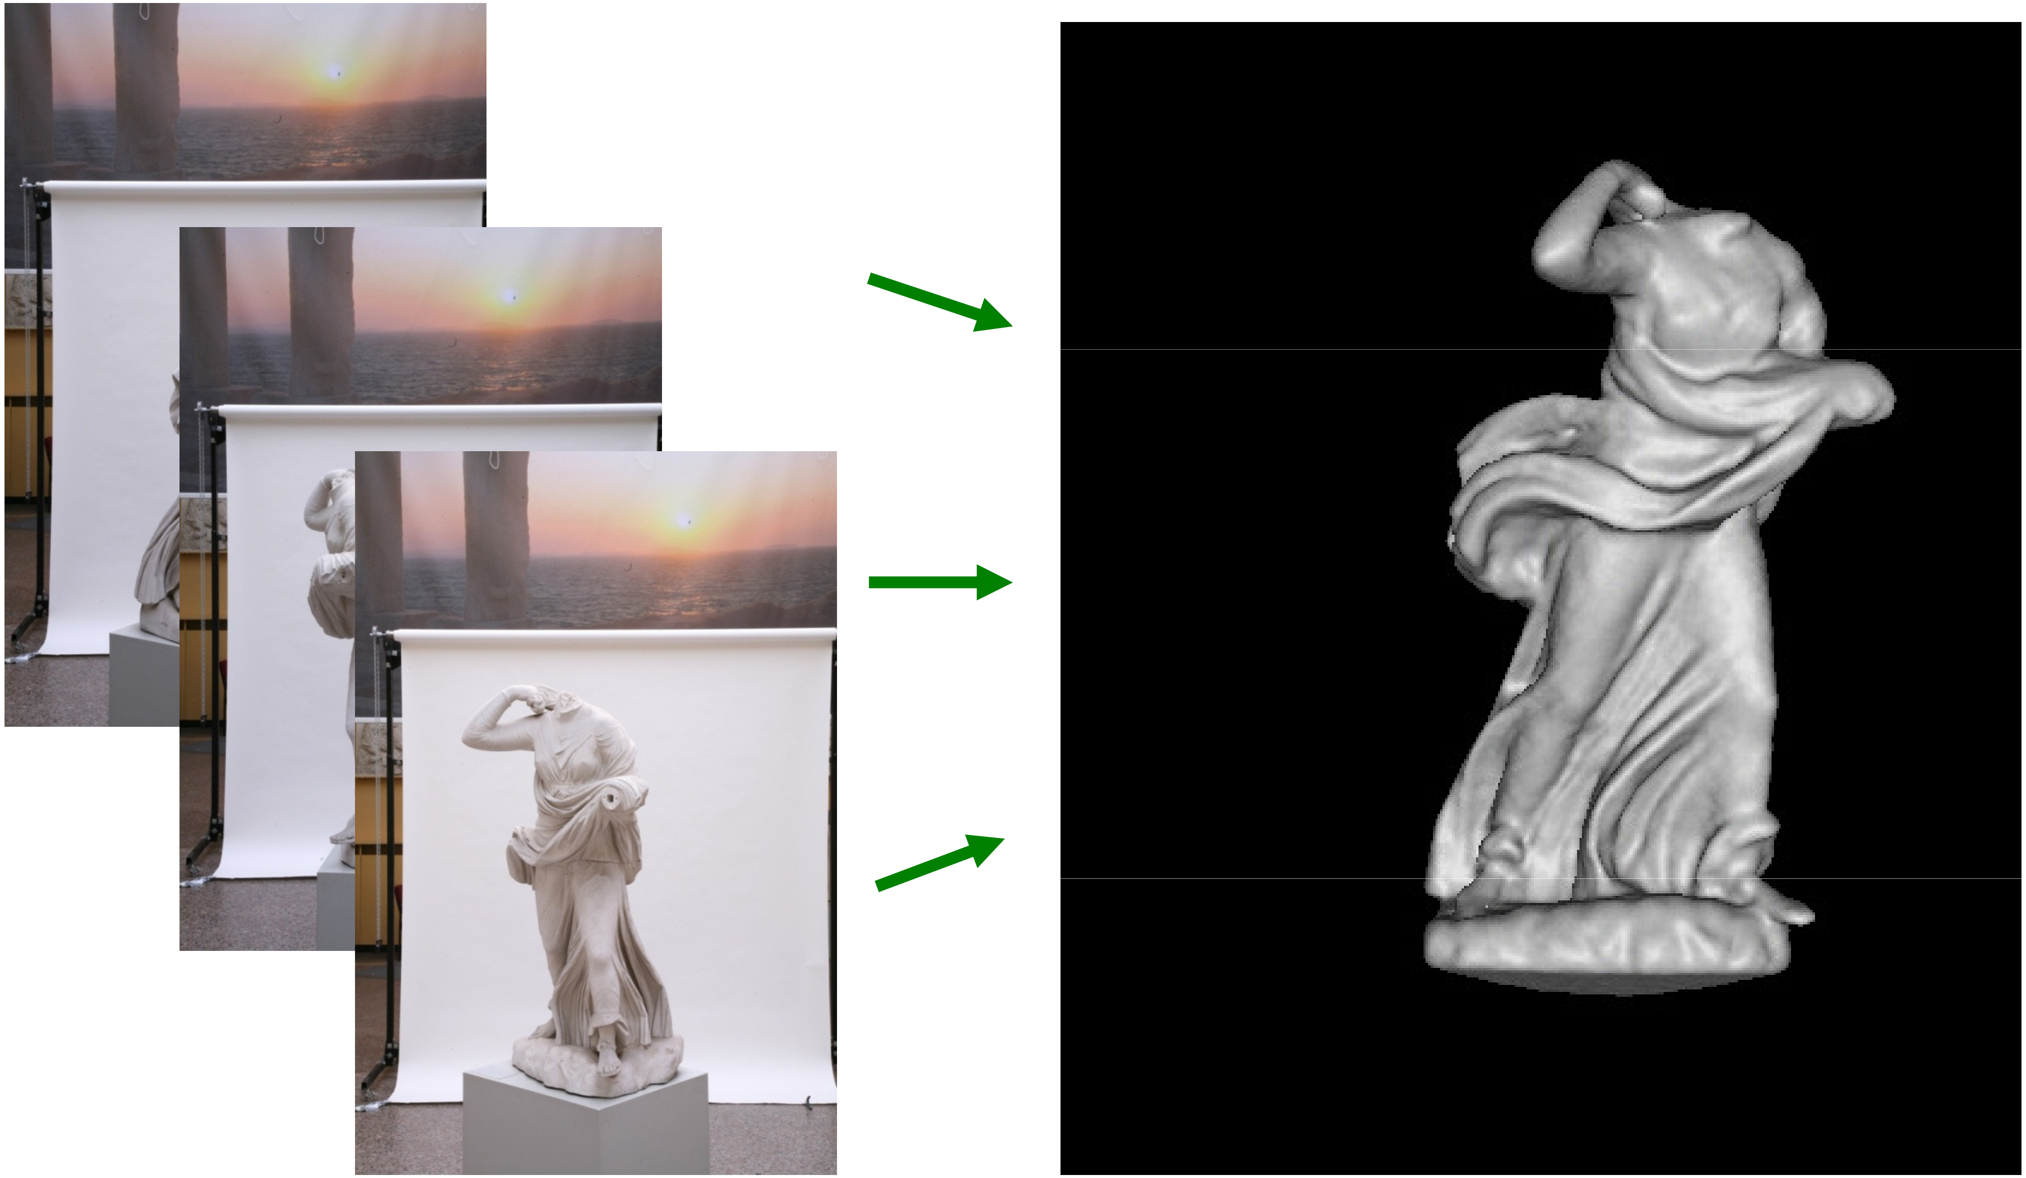

Kolev, Cremers, ECCV ’08, PAMI 2010

Reconstructing the Niobids statues (450 B.C.)

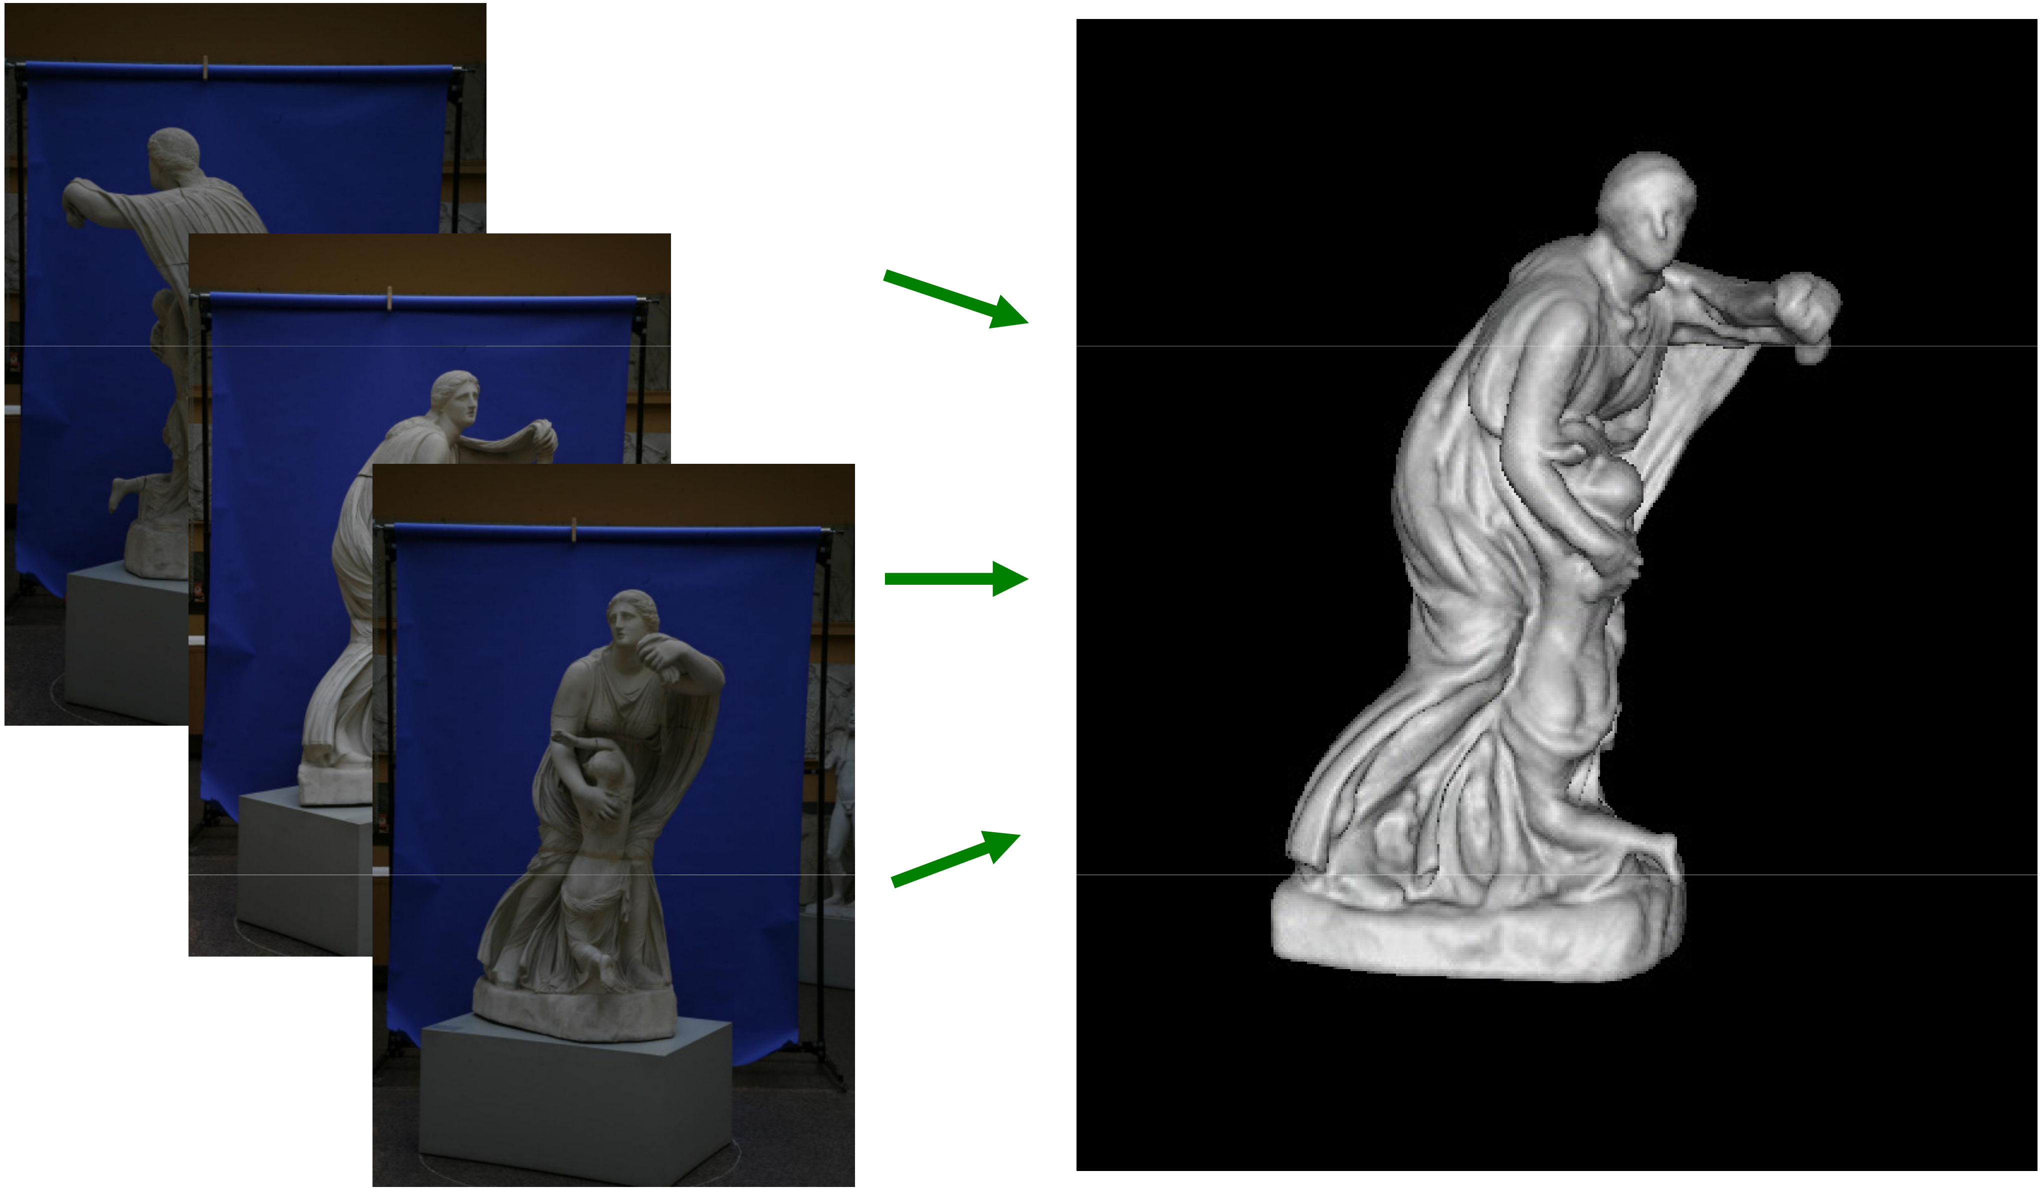

Kolev, Cremers, ECCV ’08, PAMI 2010

Reconstructing the Niobids statues (450 B.C.)

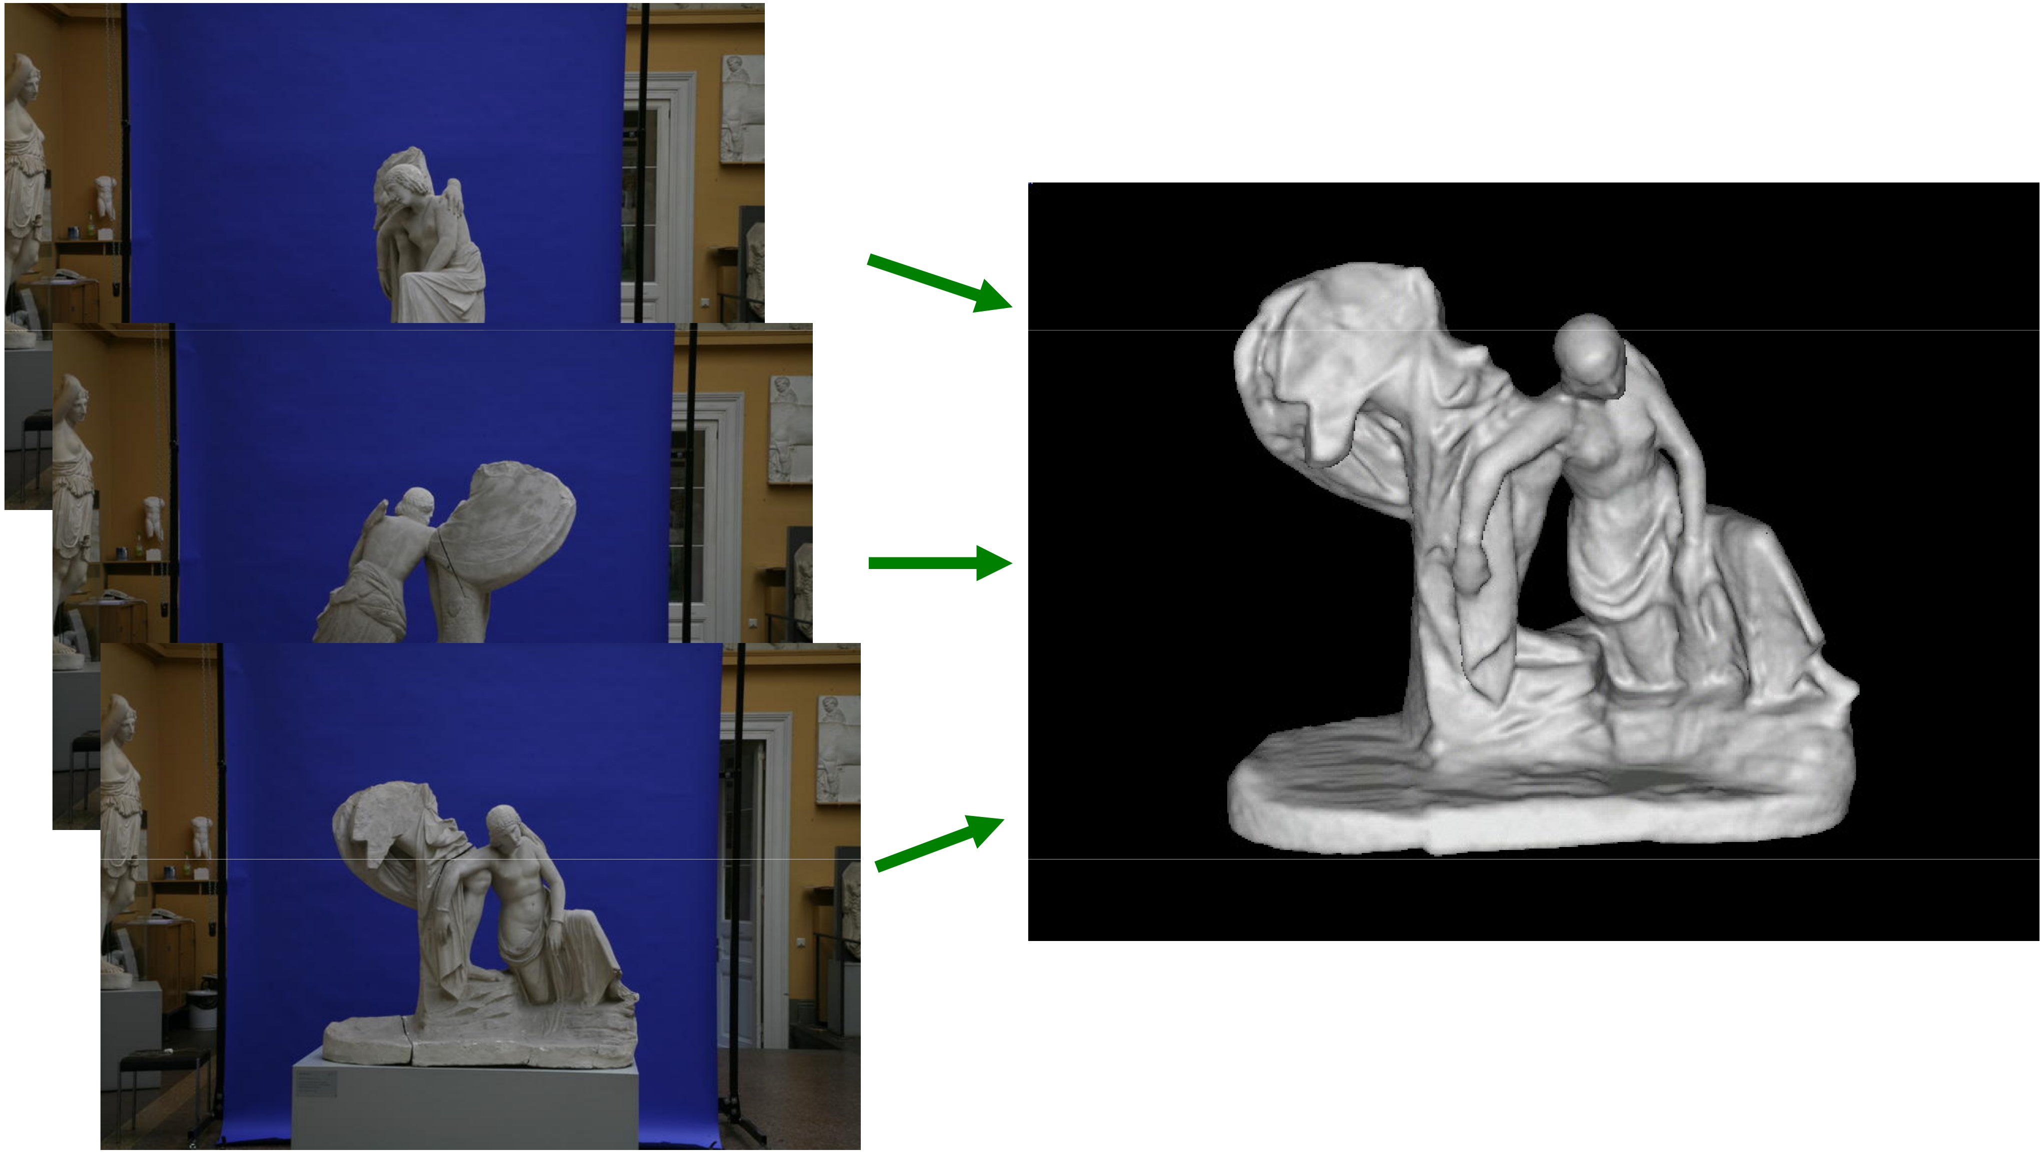

Kolev, Cremers, ECCV ’08, PAMI 2010

Overview

Multiview reconstruction

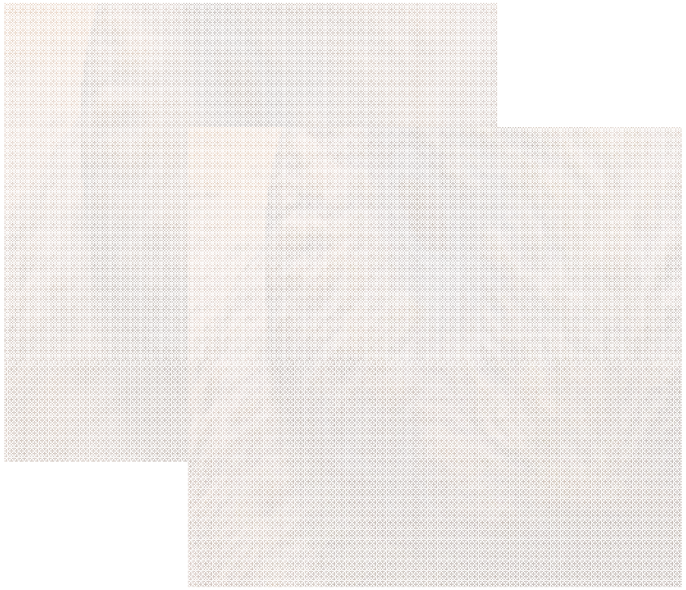

Super-resolution textures

Stereo & silhouettes

Single view reconstruction

Variational single view reconstruction

Can we recover geometry from a single image?

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

Problem: Most approaches do not work well for real-world images.

Variational single view reconstruction

Silhouette-based approaches:

Horry et al. Siggraph '97, Criminisi et al. IJCV '00,

Hoiem et al. Siggraph '05, Prasad et al. CVPR '06,…

Goal: Simple variational approach with minimal user interaction.

Solution: Fixed-volume silhouette-consistent minimal surface.

$$\min_u \int_{\mathbb{R}^3} |\nabla u| \, dx, \quad \text{s.t.} \quad \int_{\mathbb{R}^3} u \, dx = V_0$$

where $u \in BV(\mathbb{R}^3, [0,1])$ denotes the (relaxed) volume occupancy.

Proposition: The relaxed problem is convex.

Toeppe, Oswald, Rother, Cremers, ACCV 2010

In collaboration with Microsoft Research Cambridge

Variational single view reconstruction

Toeppe, Oswald, Rother, Cremers, ACCV 2010

Speedups in GPU computation: GFlops comparison

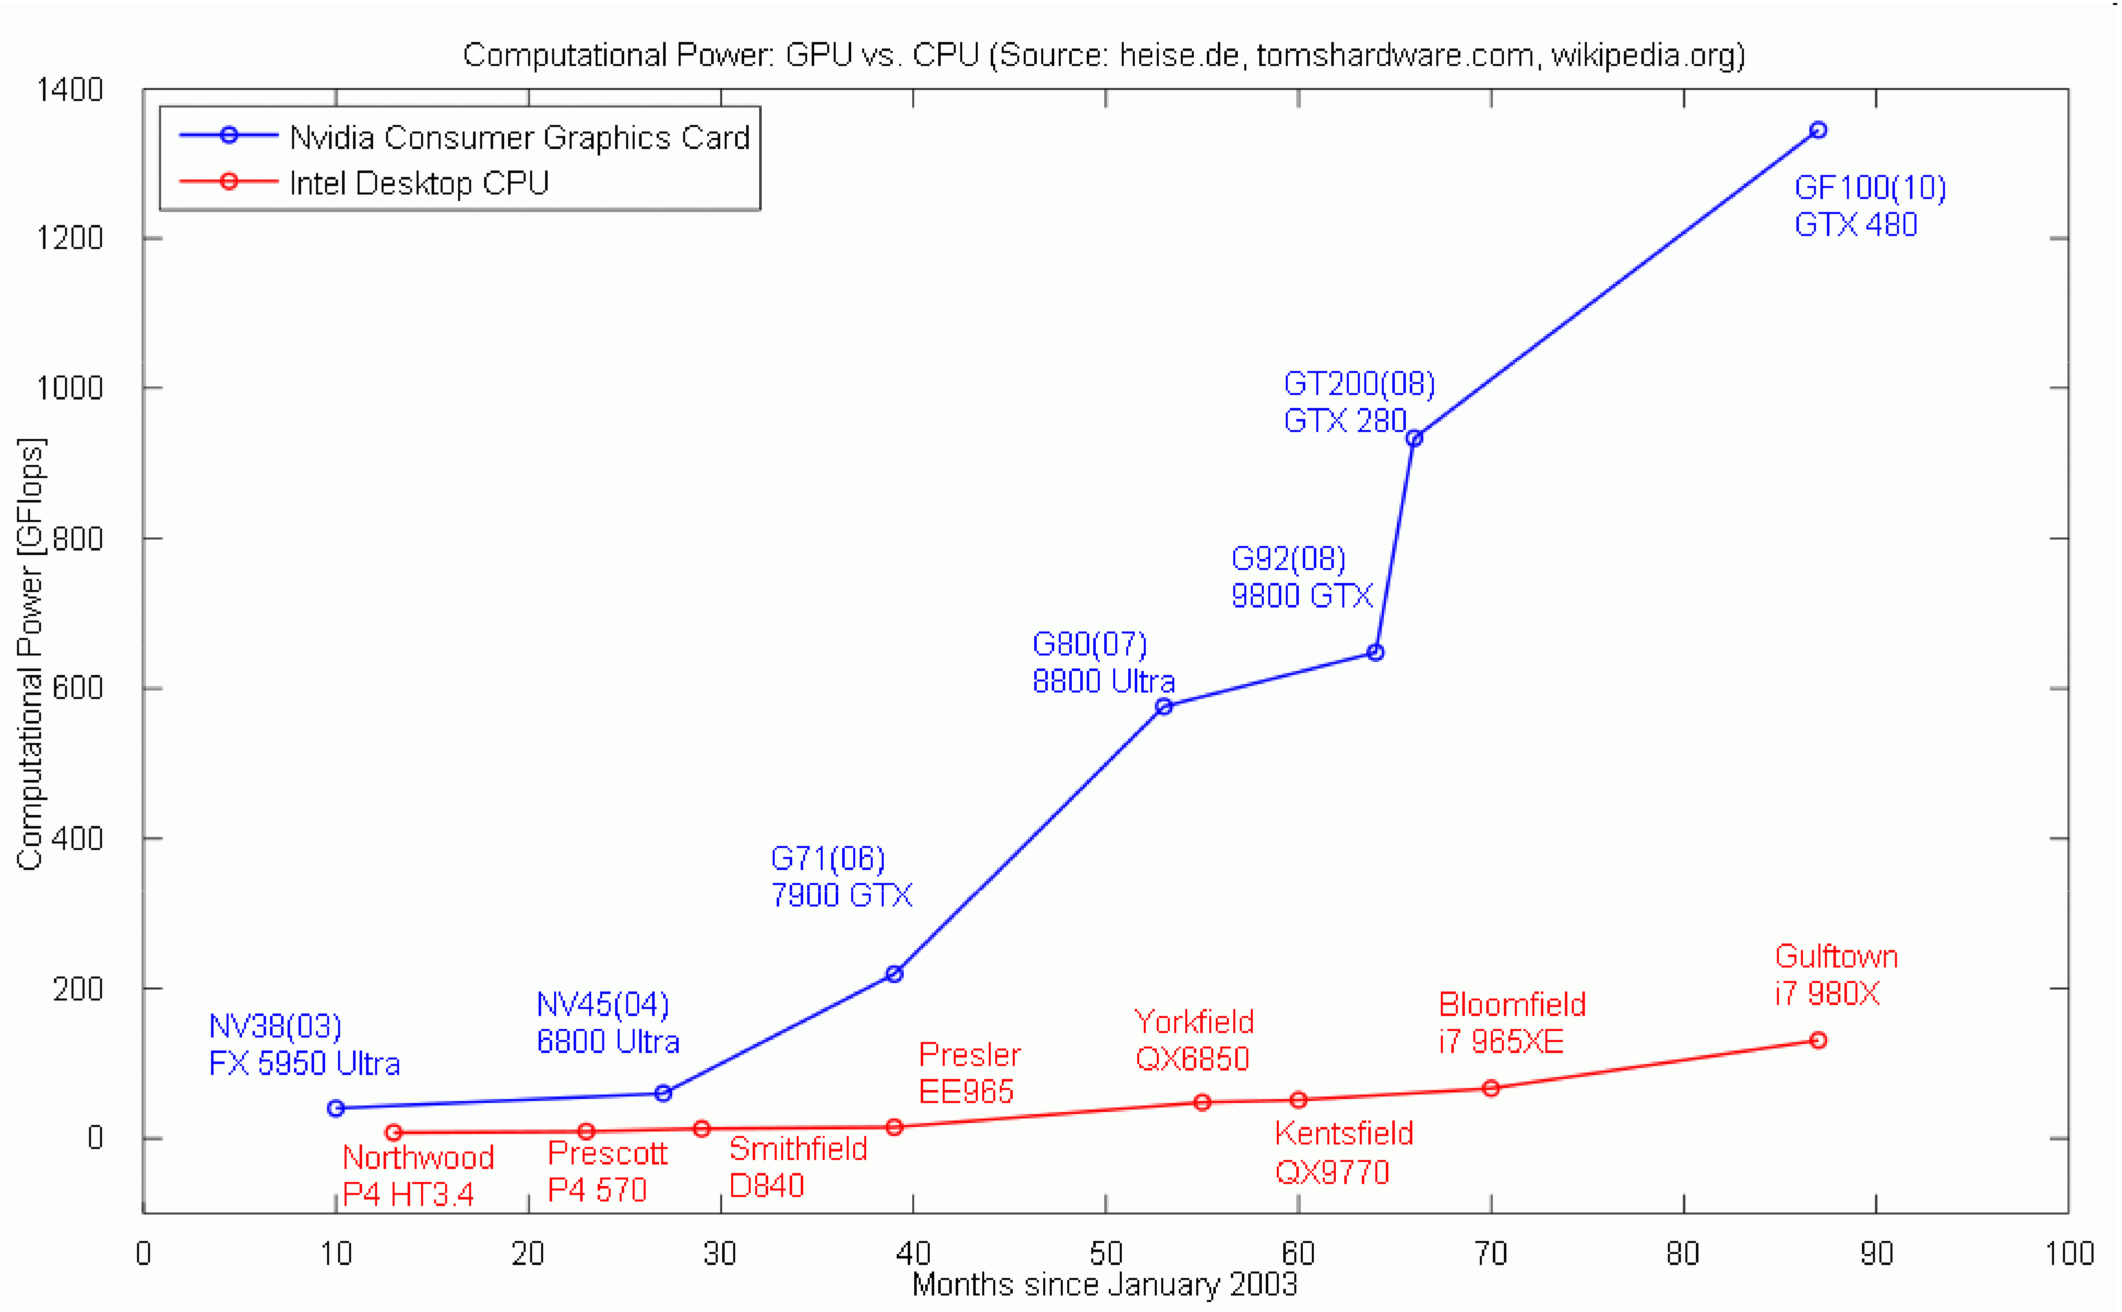

Variational single view reconstruction

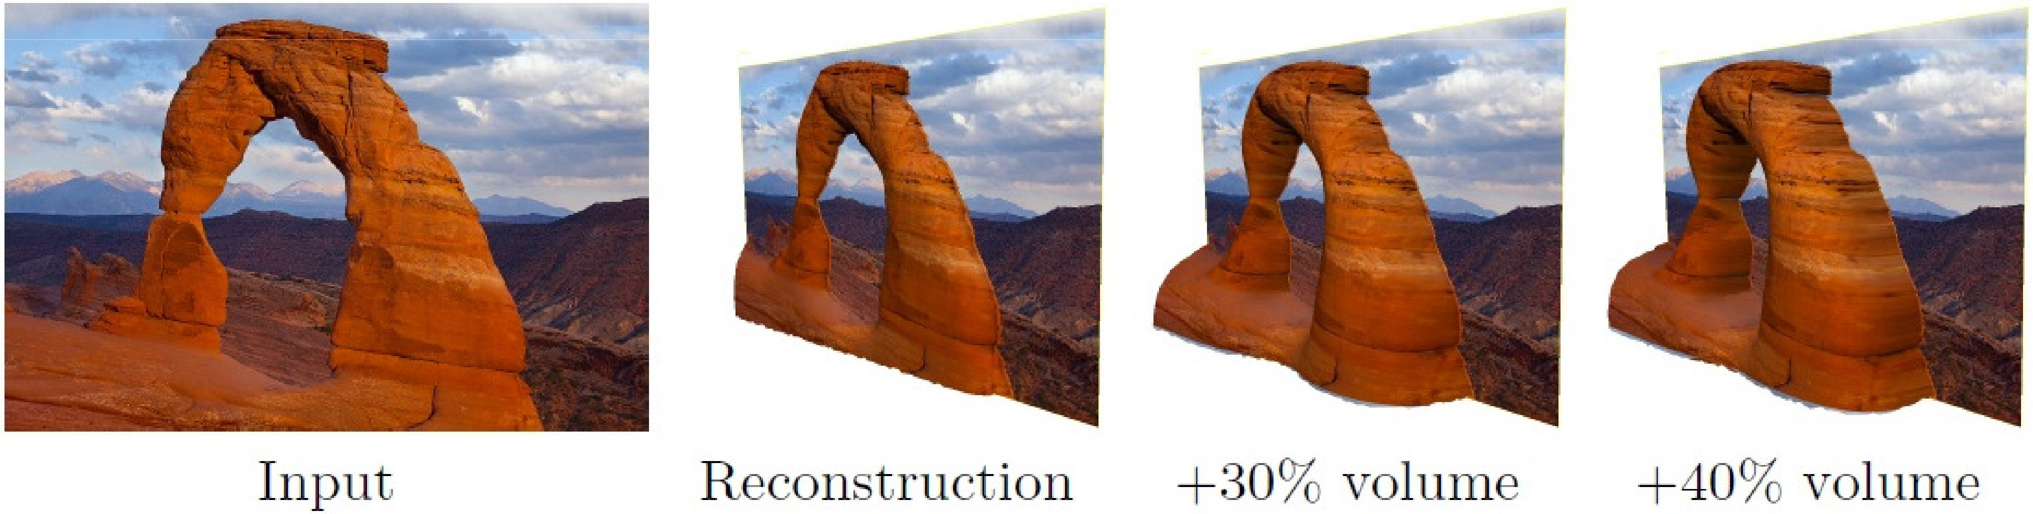

Input Reconstruction +30% volume +40% volume

Computation time approximately 1 second (on GPU).

Toeppe, Oswald, Rother, Cremers, ACCV 2010

Variational single view reconstruction

Toeppe, Oswald, Rother, Cremers, ACCV 2010

Variational single view reconstruction

Toeppe, Oswald, Rother, Cremers, ACCV 2010

Variational single view reconstruction

Toeppe, Oswald, Rother, Cremers, ACCV 2010

Variational single view reconstruction

Toeppe, Oswald, Rother, Cremers, ACCV 2010

Summary

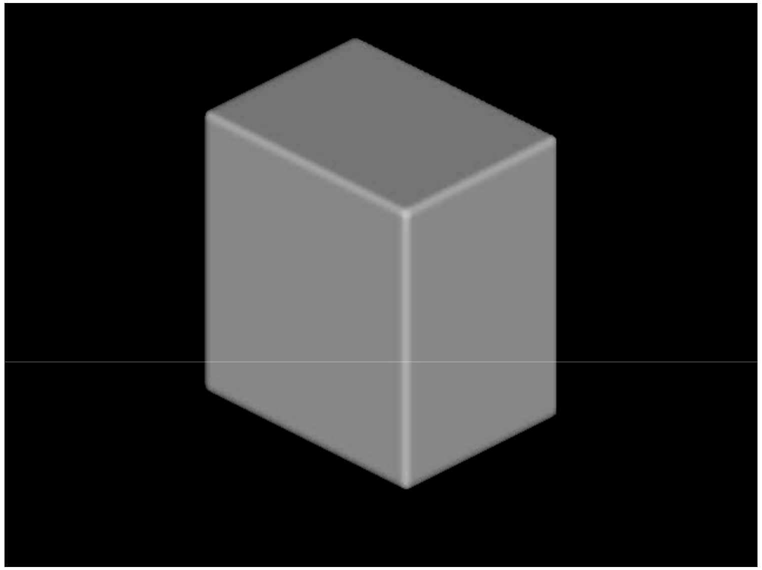

Multiview reconstruction
via convex relaxation

Stereo & silhouettes via convex
optimization & convex constraints

Superresolution texture

Single view reconstruction