Variational Methods for Computer Vision

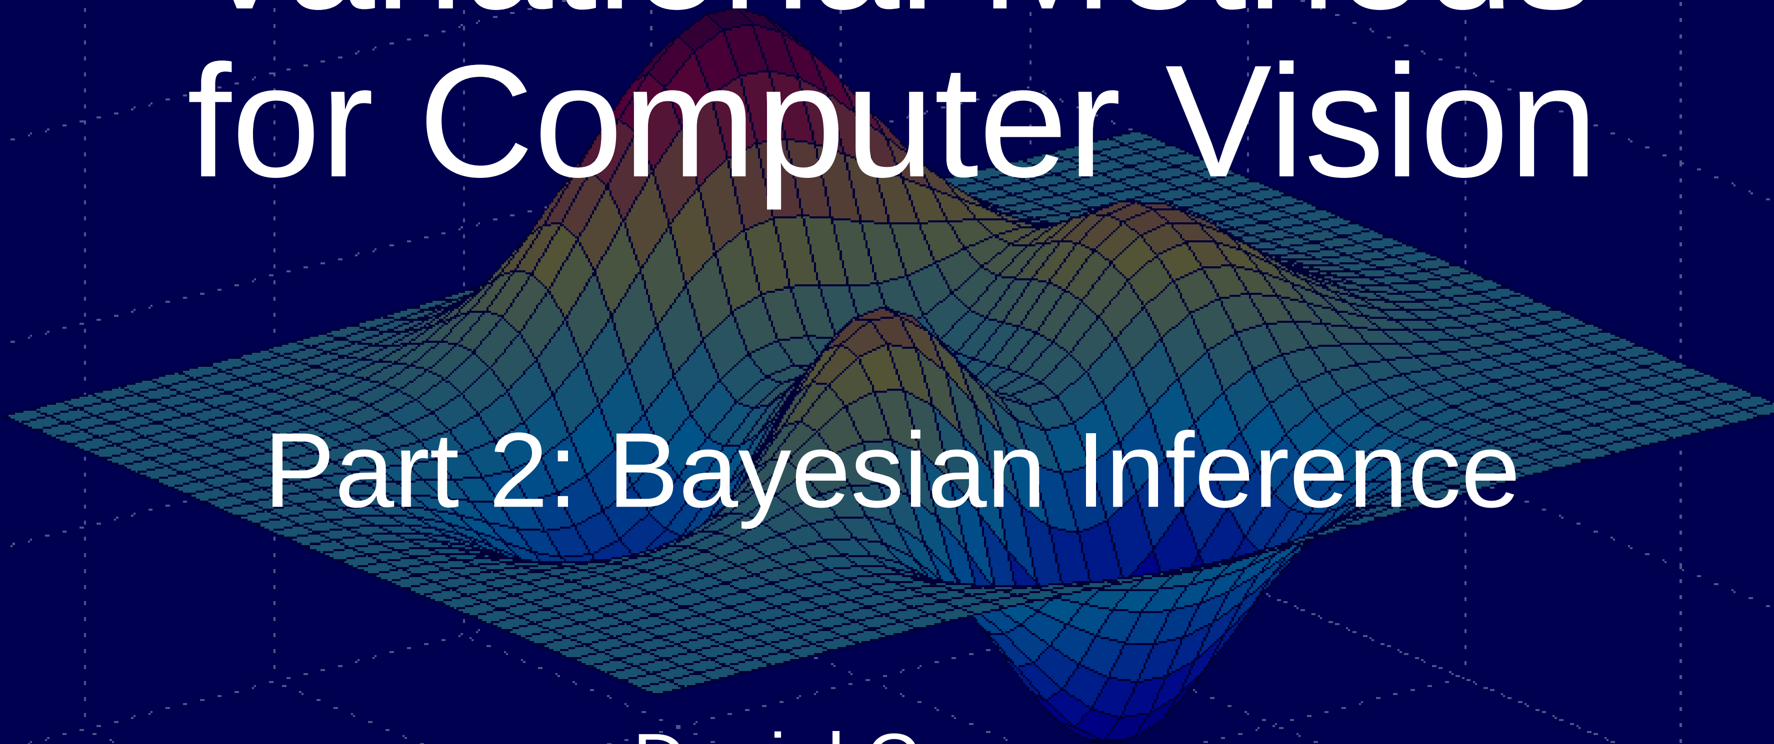

Part 2: Bayesian Inference

Daniel Cremers

Computer Science & Mathematics

TU Munich

Segmentation and Prior Knowledge

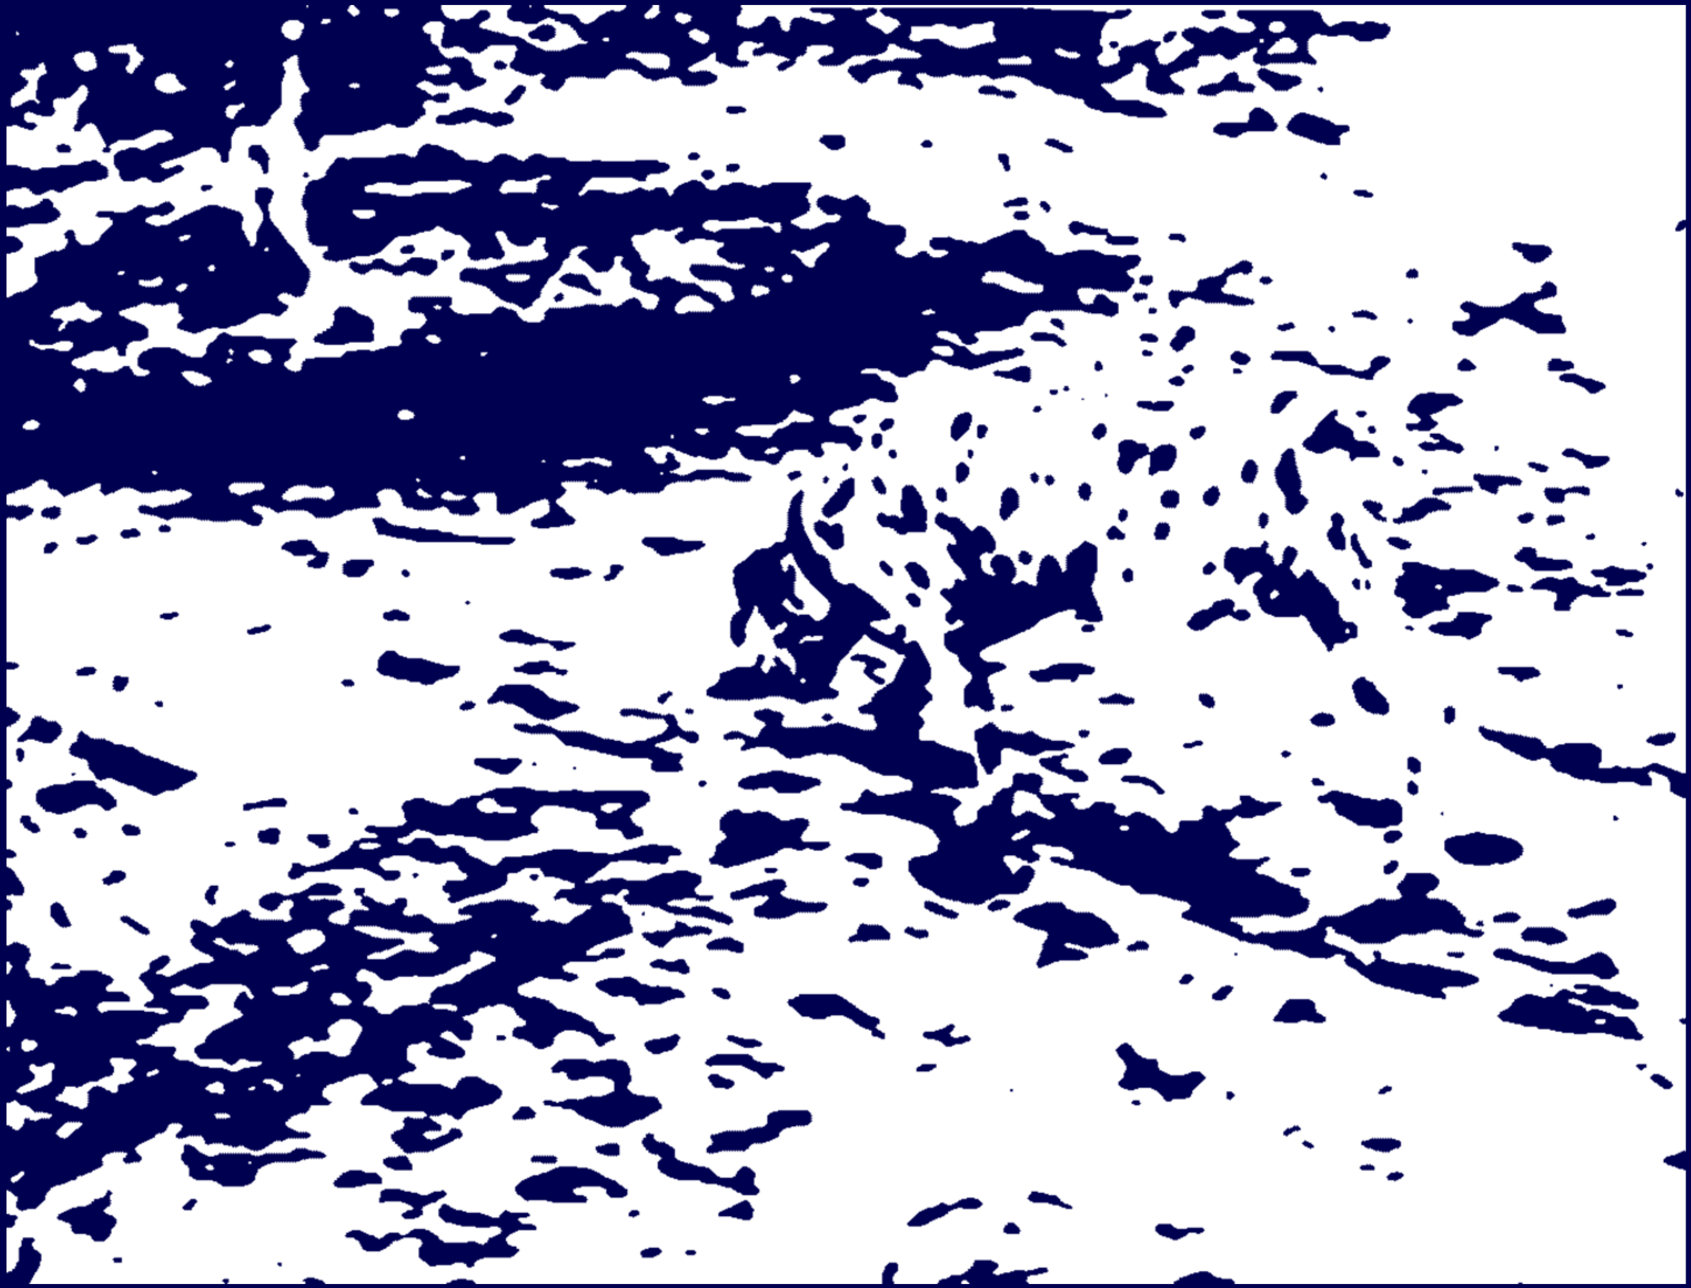

Variational Image Segmentation

Mumford, Shah '89, Chambolle et al. '12:

$$\min_{\Omega_1,\dots,\Omega_n} \sum_i \int_{\Omega_i} f_i \, dx \; + \; \nu \left| \partial \Omega_i \right|$$

$$\text{s.t.} \; \bigcup_i \Omega_i = \Omega, \quad \Omega_i \cap \Omega_j = \emptyset \; \forall i \neq j$$

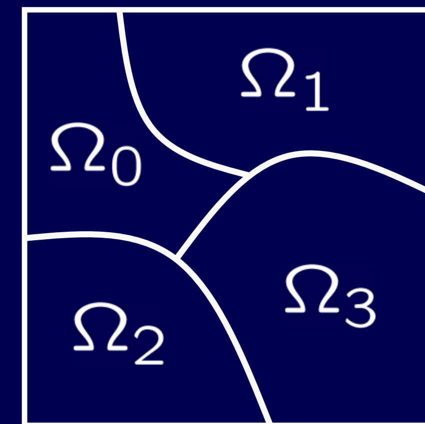

2 label segmentation

10 label segmentation

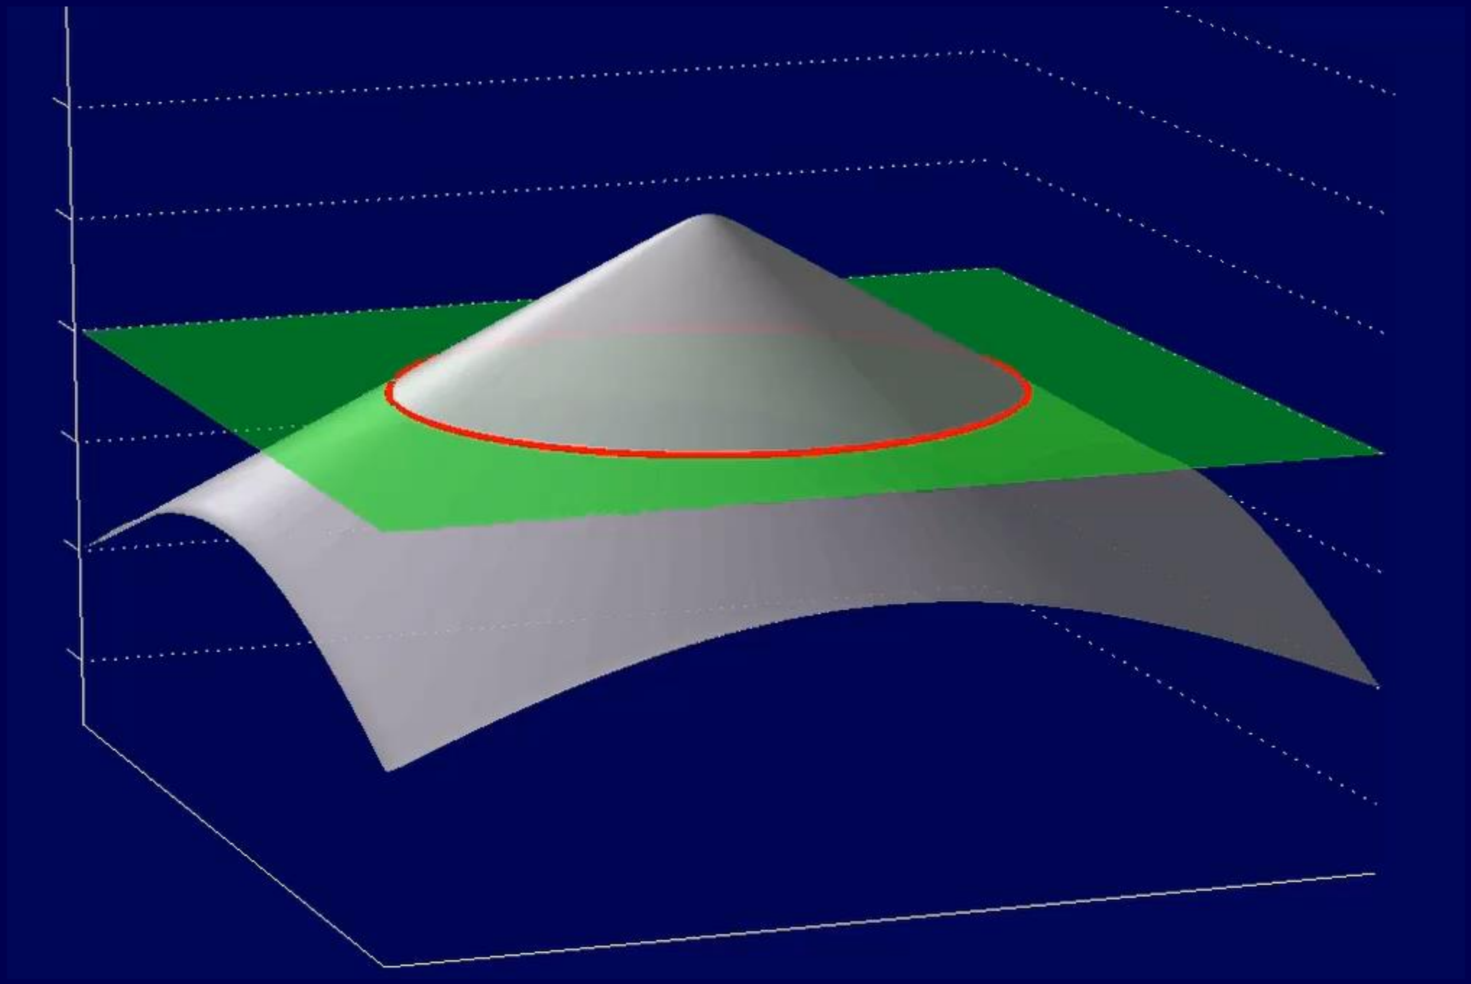

$$C = \{x \in \Omega \,|\, \phi(x) = 0\}, \quad \phi : \Omega \subset \mathbb{R}^2 \to \mathbb{R}$$

Dervieux, Thomasset ’79,’81, Osher, Sethian ’88

Level Set Segmentation

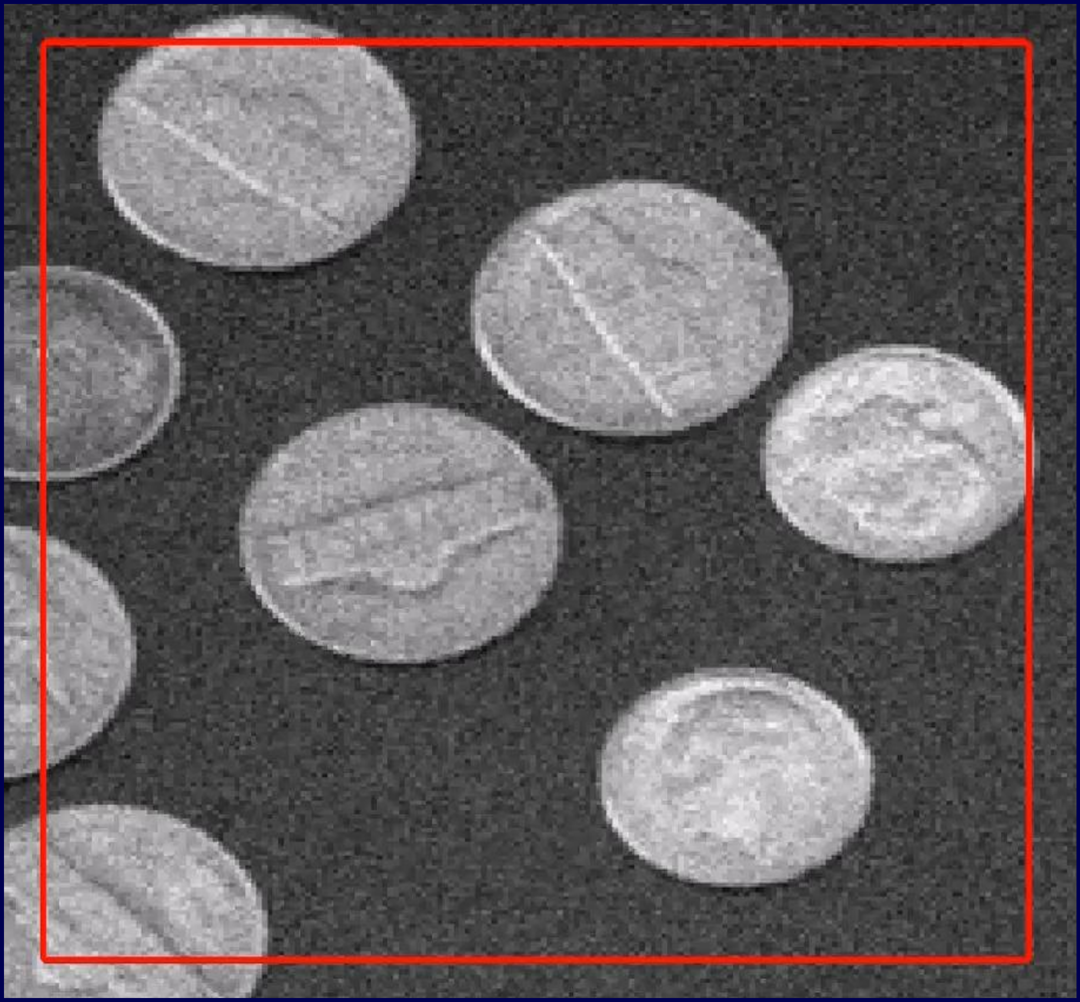

Chan, Vese '01

A Challenging Segmentation Problem

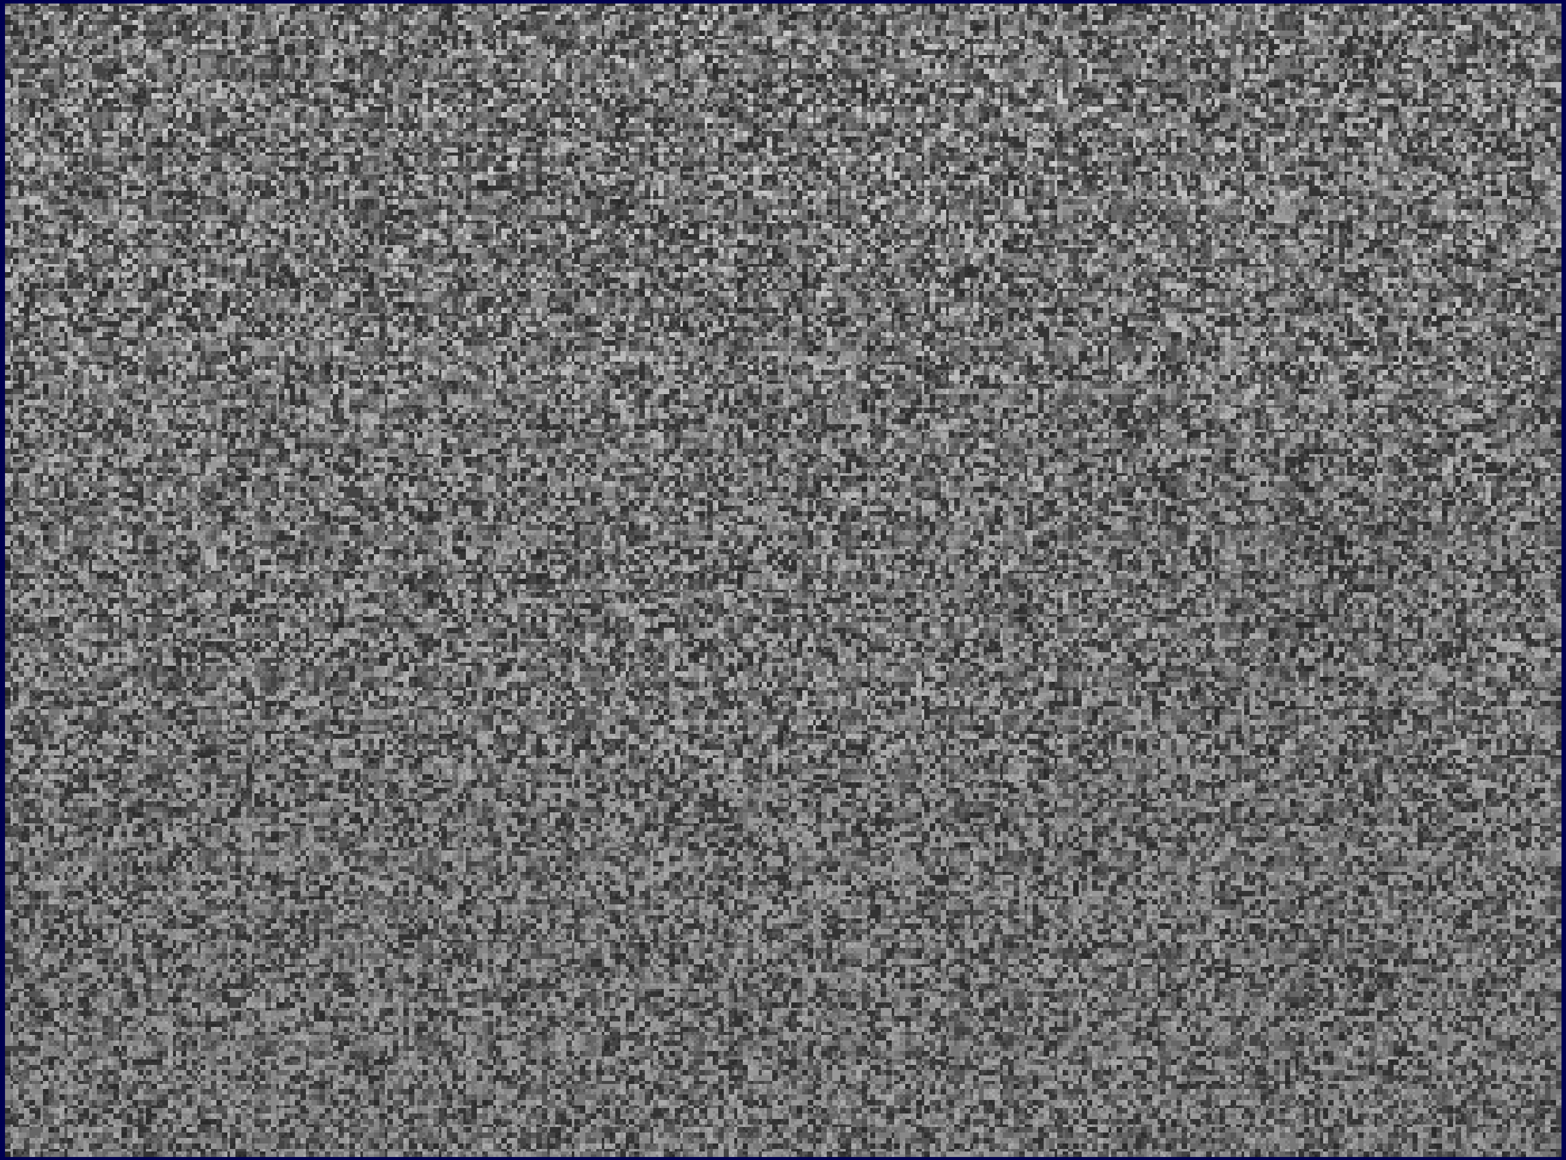

Cremers, IEEE PAMI '06

Statistical Interpretation

Given an image I, find the most likely segmentation by maximizing the conditional probability

$$\mathcal{P}(C \mid I) = \frac{\mathcal{P}(I \mid C)\, \mathcal{P}(C)}{\mathcal{P}(I)}$$

with respect to a curve C.

This is equivalent to minimizing its negative logarithm which (up to a constant) is given by the energy

$$E(C) = E_{data}(C) \; + \; E_{shape}(C)$$

with

$$E_{data} = -\log \mathcal{P}(I \mid C) \quad \text{and} \quad E_{shape} = -\log \mathcal{P}(C)$$

Cremers, Osher, Soatto, IJCV '06

Insufficient Low-Level Information

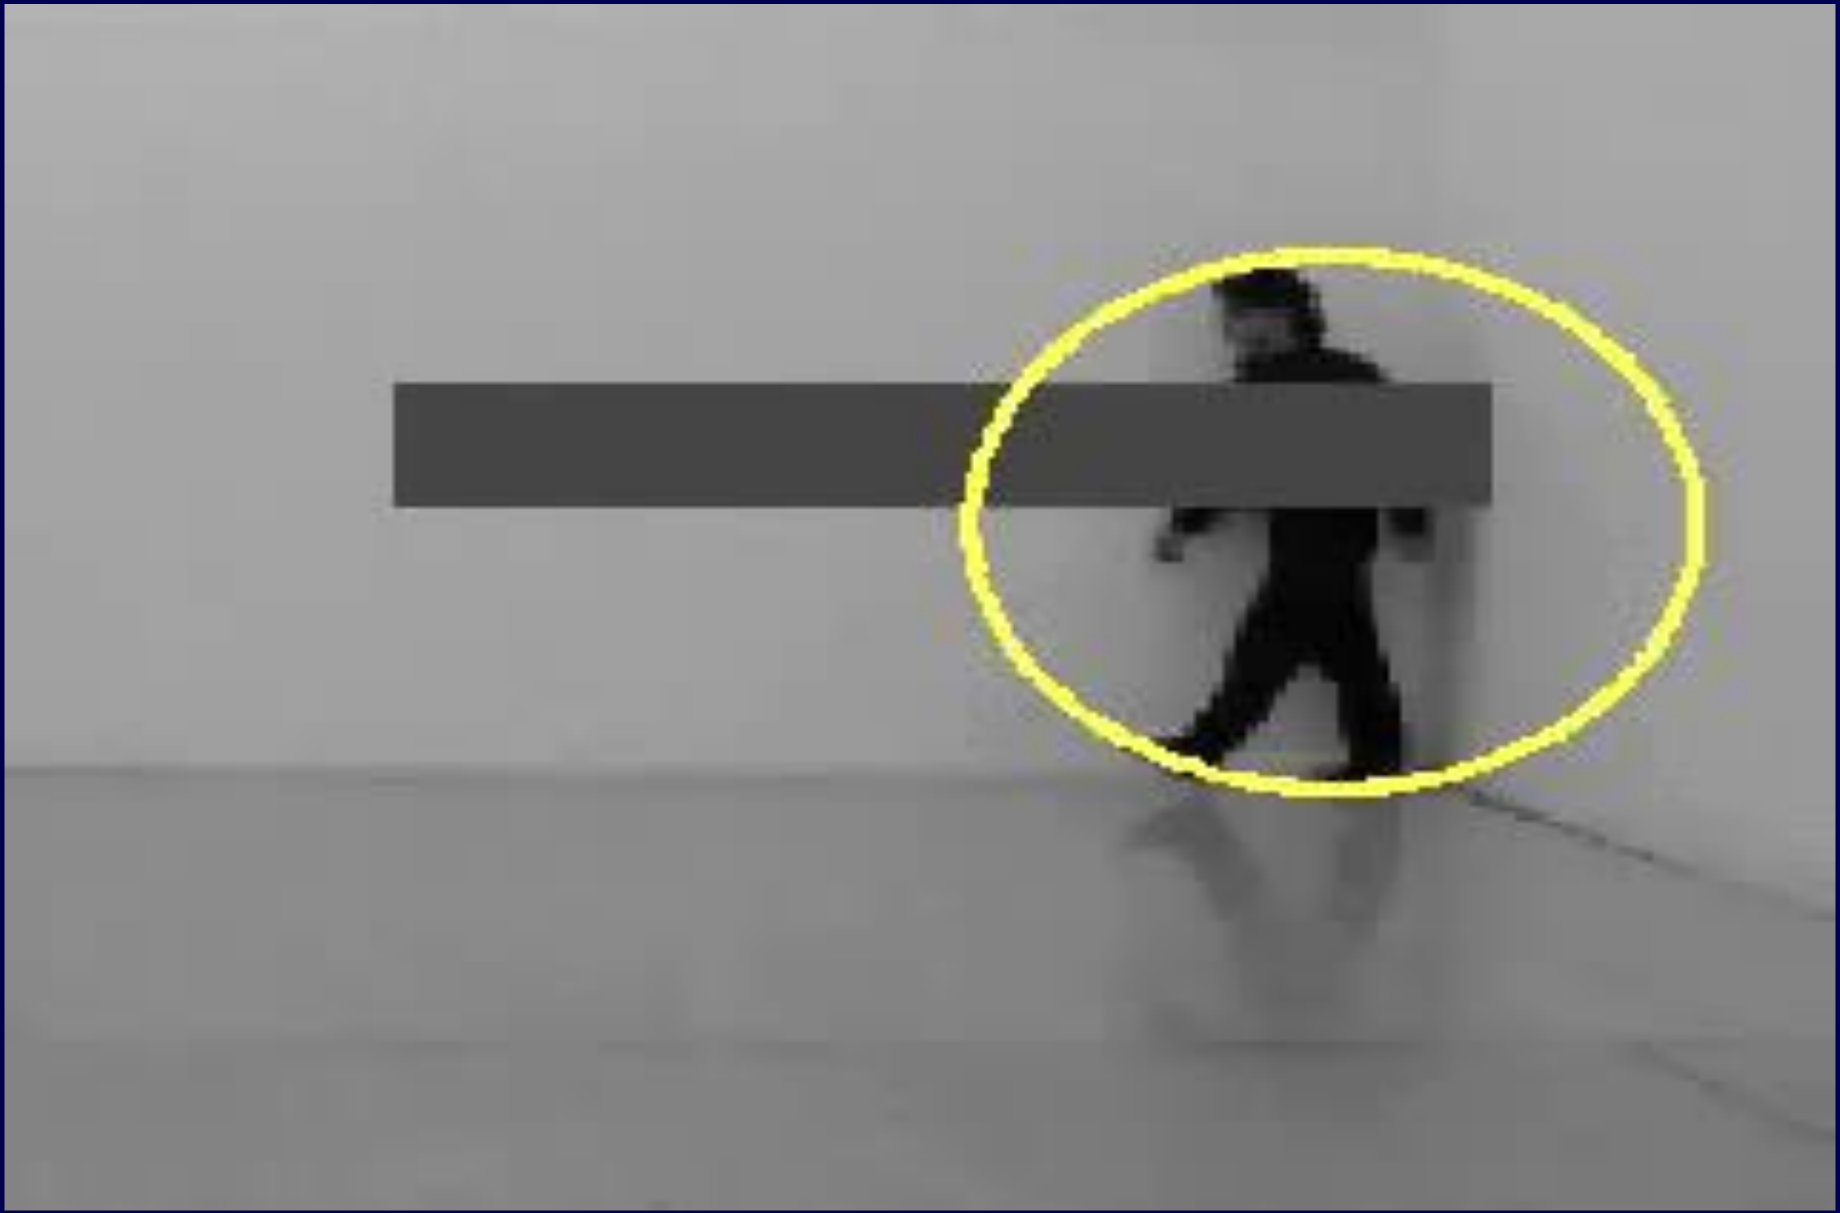

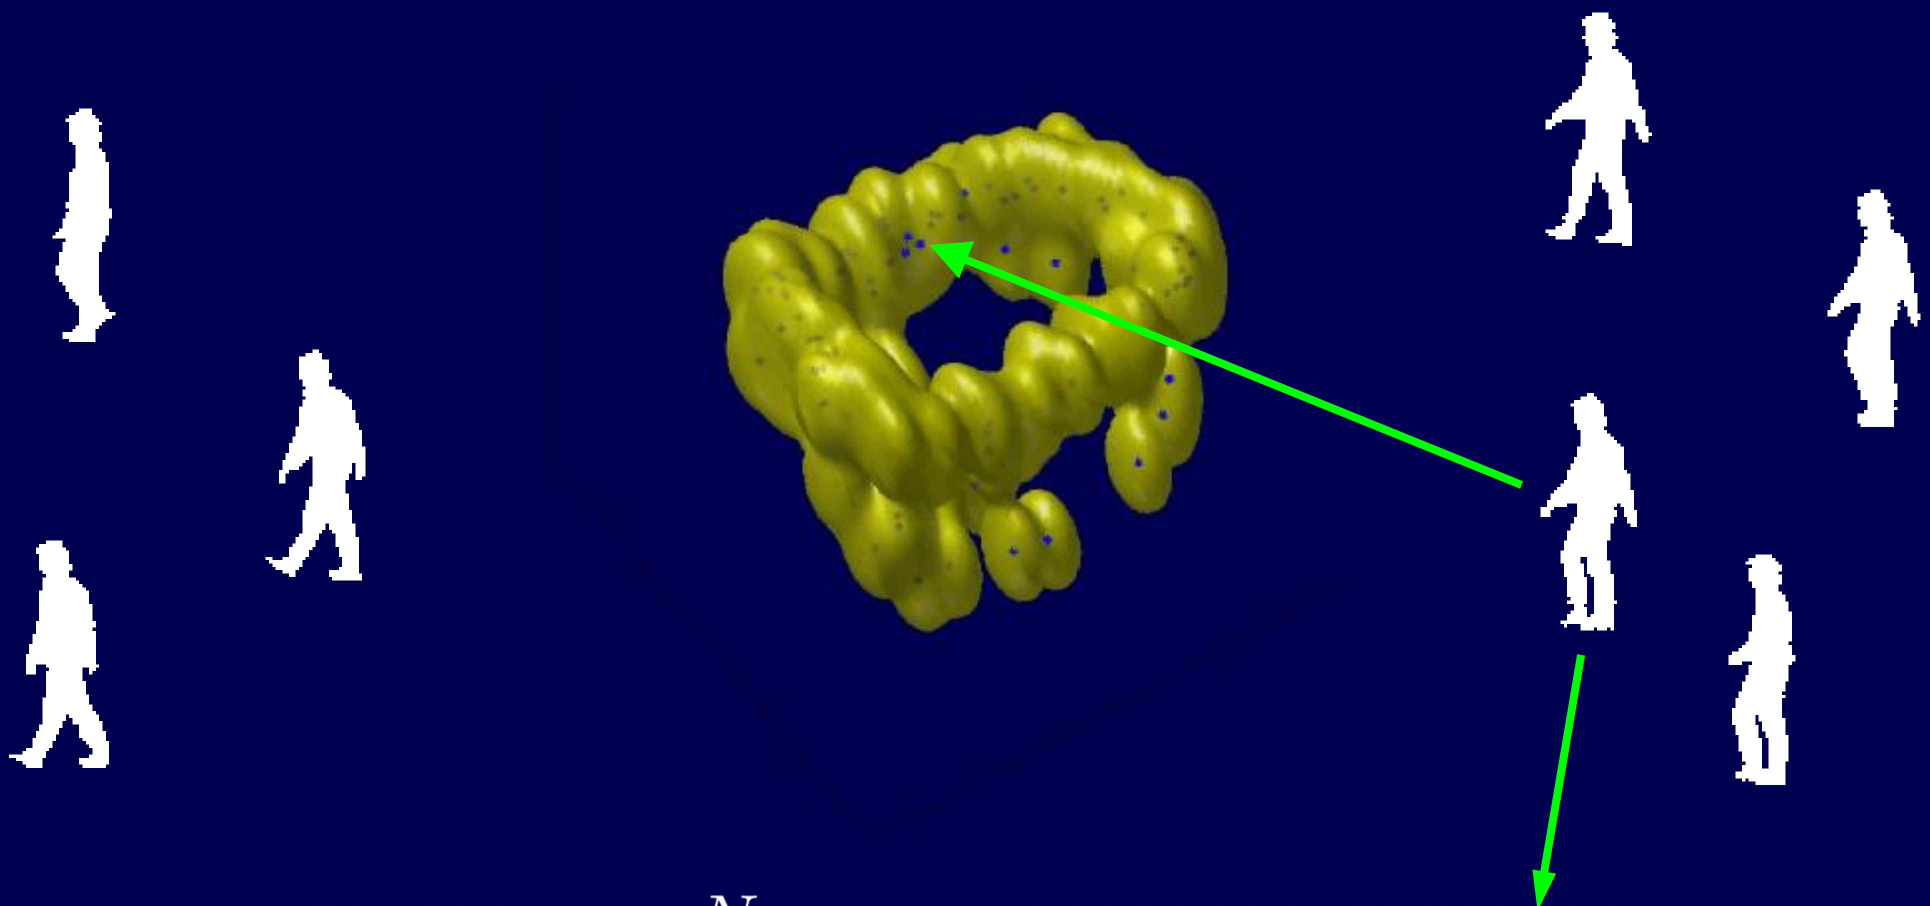

$$\mathcal{P}(\phi) \propto \frac{1}{N} \sum_{i=1}^{N} \exp\left(-\frac{1}{2\sigma^2} d^2(\phi, \phi_i)\right)$$

Cremers, Osher, Soatto, IJCV '06

Rosenblatt '56, Parzen '62

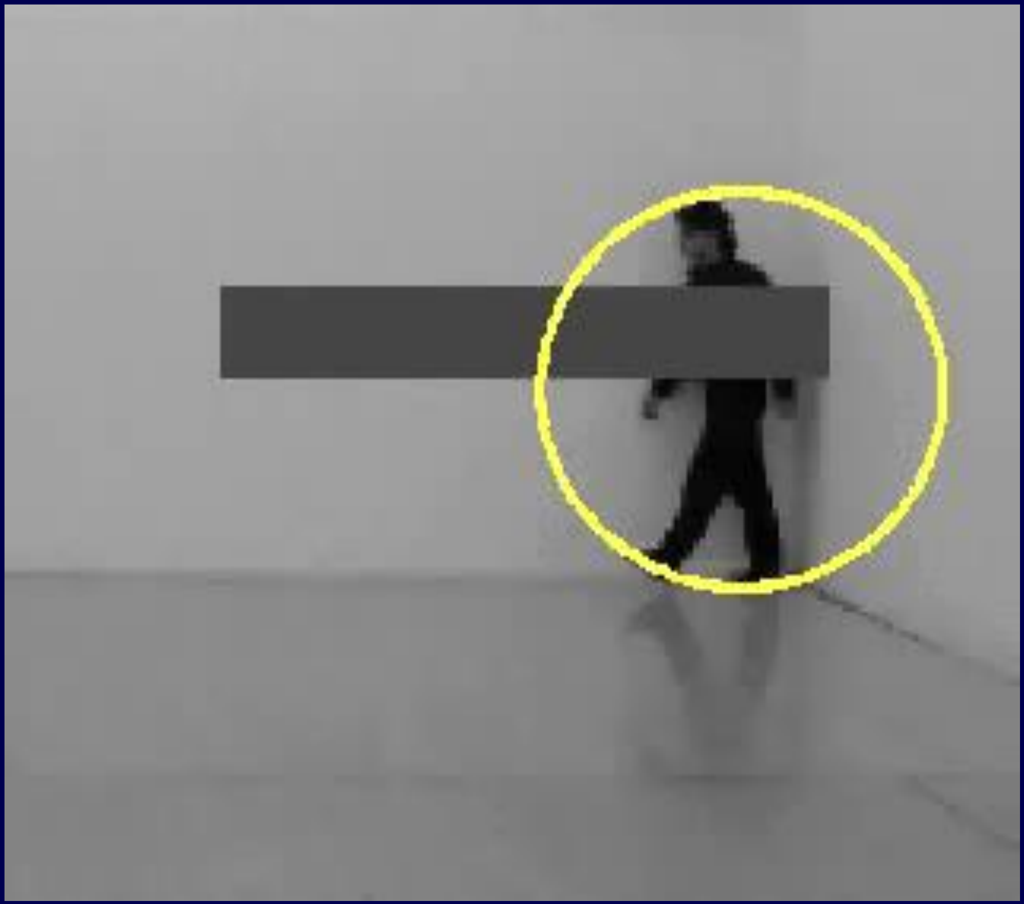

$E_{data} + E_{length} \rightarrow \min$

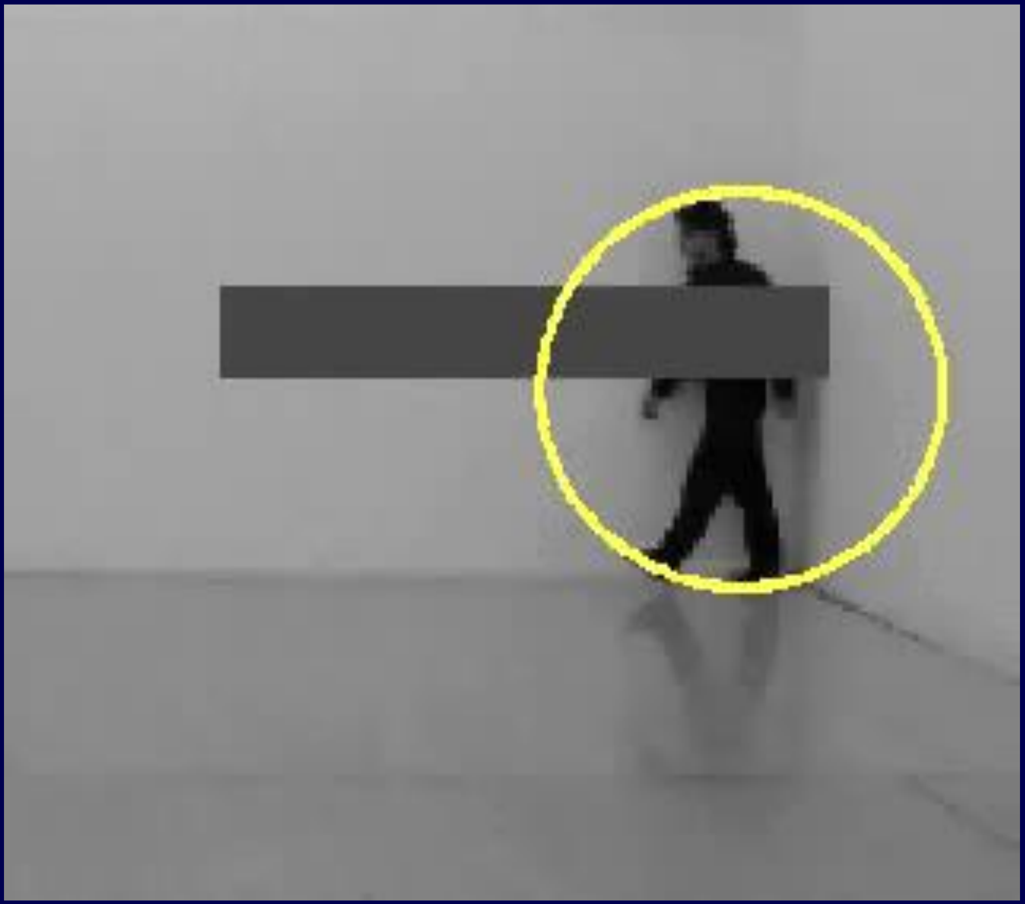

$E_{data} + E_{shape} \rightarrow \min$

Cremers, Osher, Soatto, IJCV '06

Limitations of Static Shape Priors

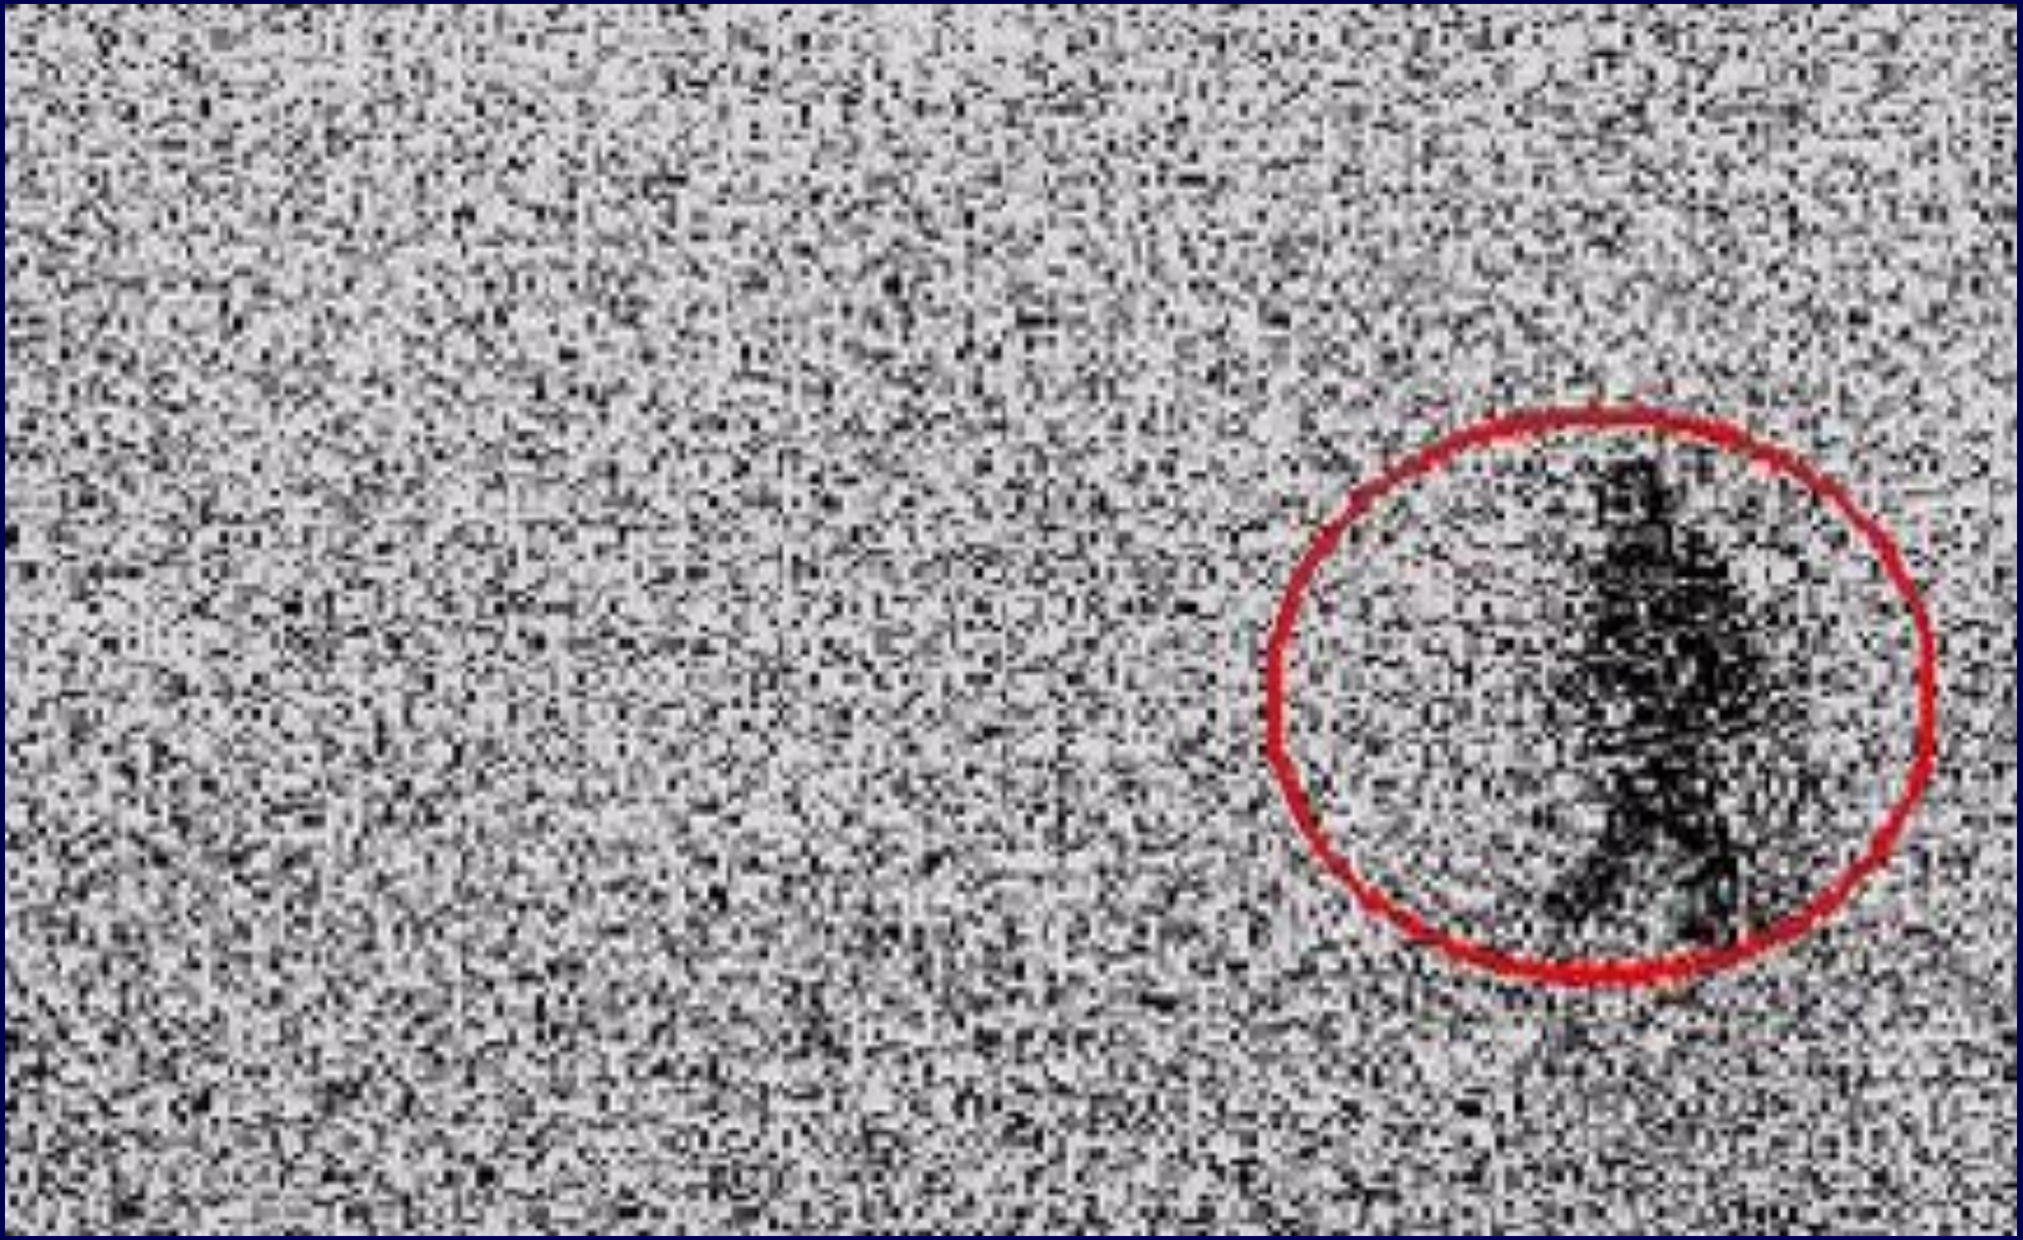

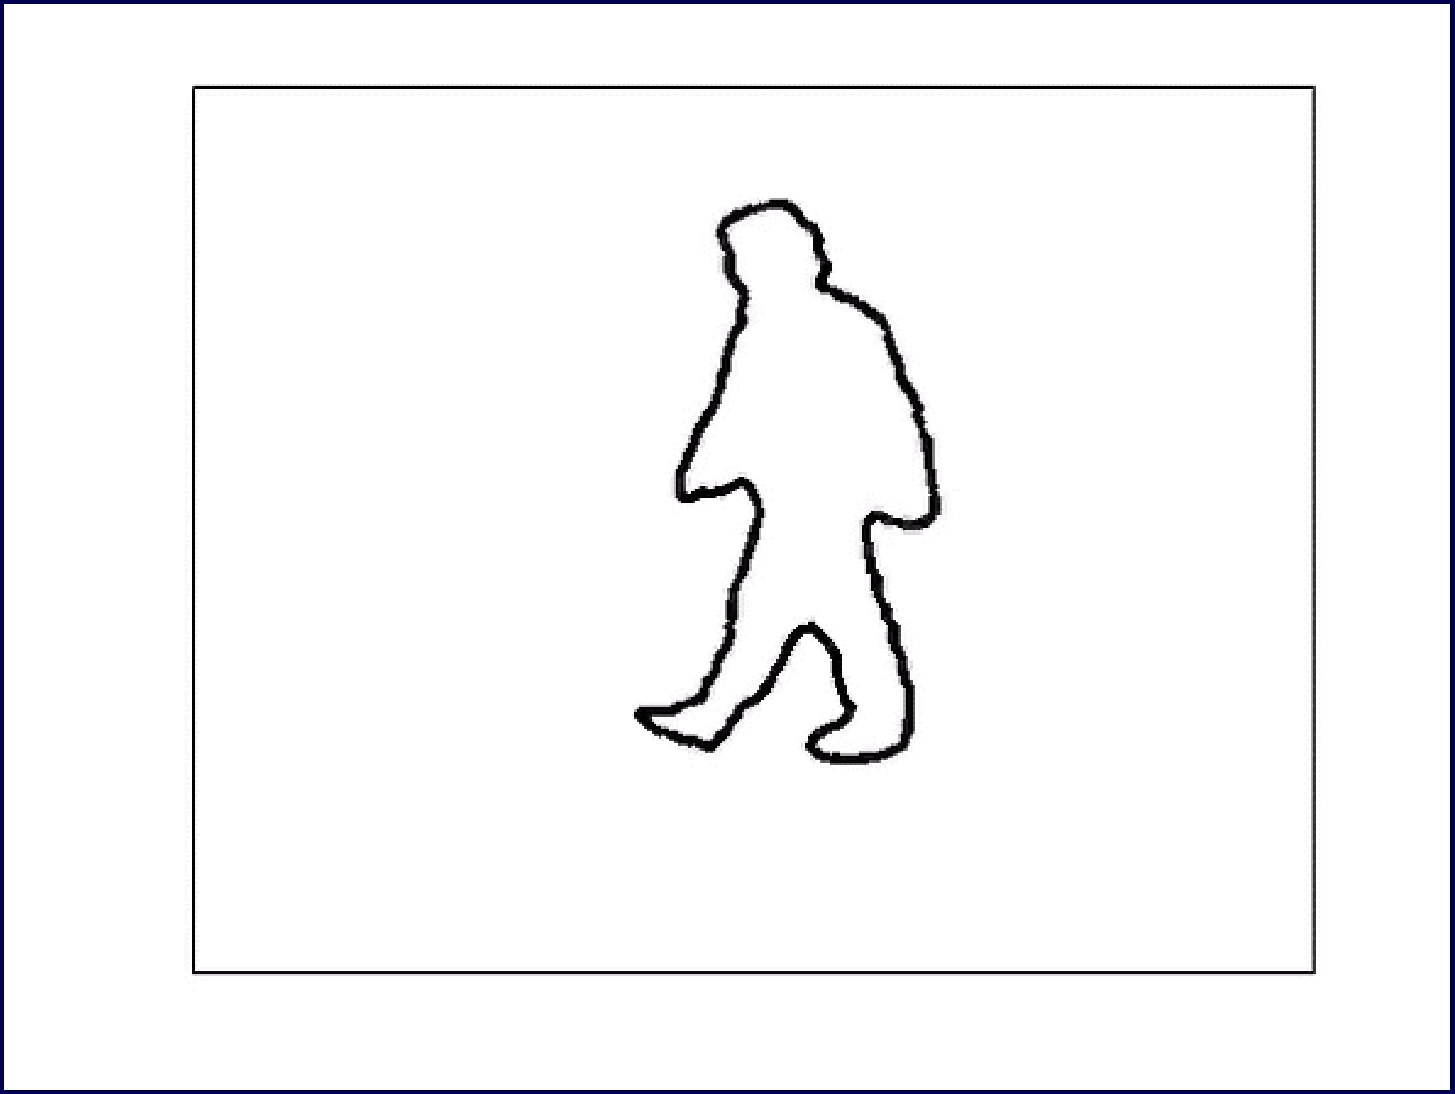

Training sequence

Dynamical Models for Implicit Shapes

1. Low-dim. representation via PCA (*Leventon et al. '00, Tsai et al. '01*):

$$\phi_i(x) \approx \phi_0(x) + \sum_{j=1}^{m} \alpha_{ij} \psi_j(x)$$

mean — $\phi_0(x)$; eigenmodes — $\psi_j(x)$

$$\alpha_{ij} = \int (\phi_i - \phi_0)\, \psi_j \, dx \qquad \boldsymbol{\alpha}_i = (\alpha_{i1}, \dots, \alpha_{im})$$

2. Autoregressive model for the shape coefficients:

$$\boldsymbol{\alpha}_t = \boldsymbol{\mu} + \sum_{i=1}^{k} A_i \, \boldsymbol{\alpha}_{t-i} + \boldsymbol{\eta}$$

mean — $\boldsymbol{\mu}$; transition matrices — A_i; Gaussian noise — $\boldsymbol{\eta}$

3. Synthesize shape vectors $\boldsymbol{\alpha}_t$ and embedding surfaces ϕ_t:

$$\phi_t(x) = \phi_0(x) + \boldsymbol{\alpha}_t^\top \boldsymbol{\psi}(x)$$

Cremers, IEEE PAMI 2006

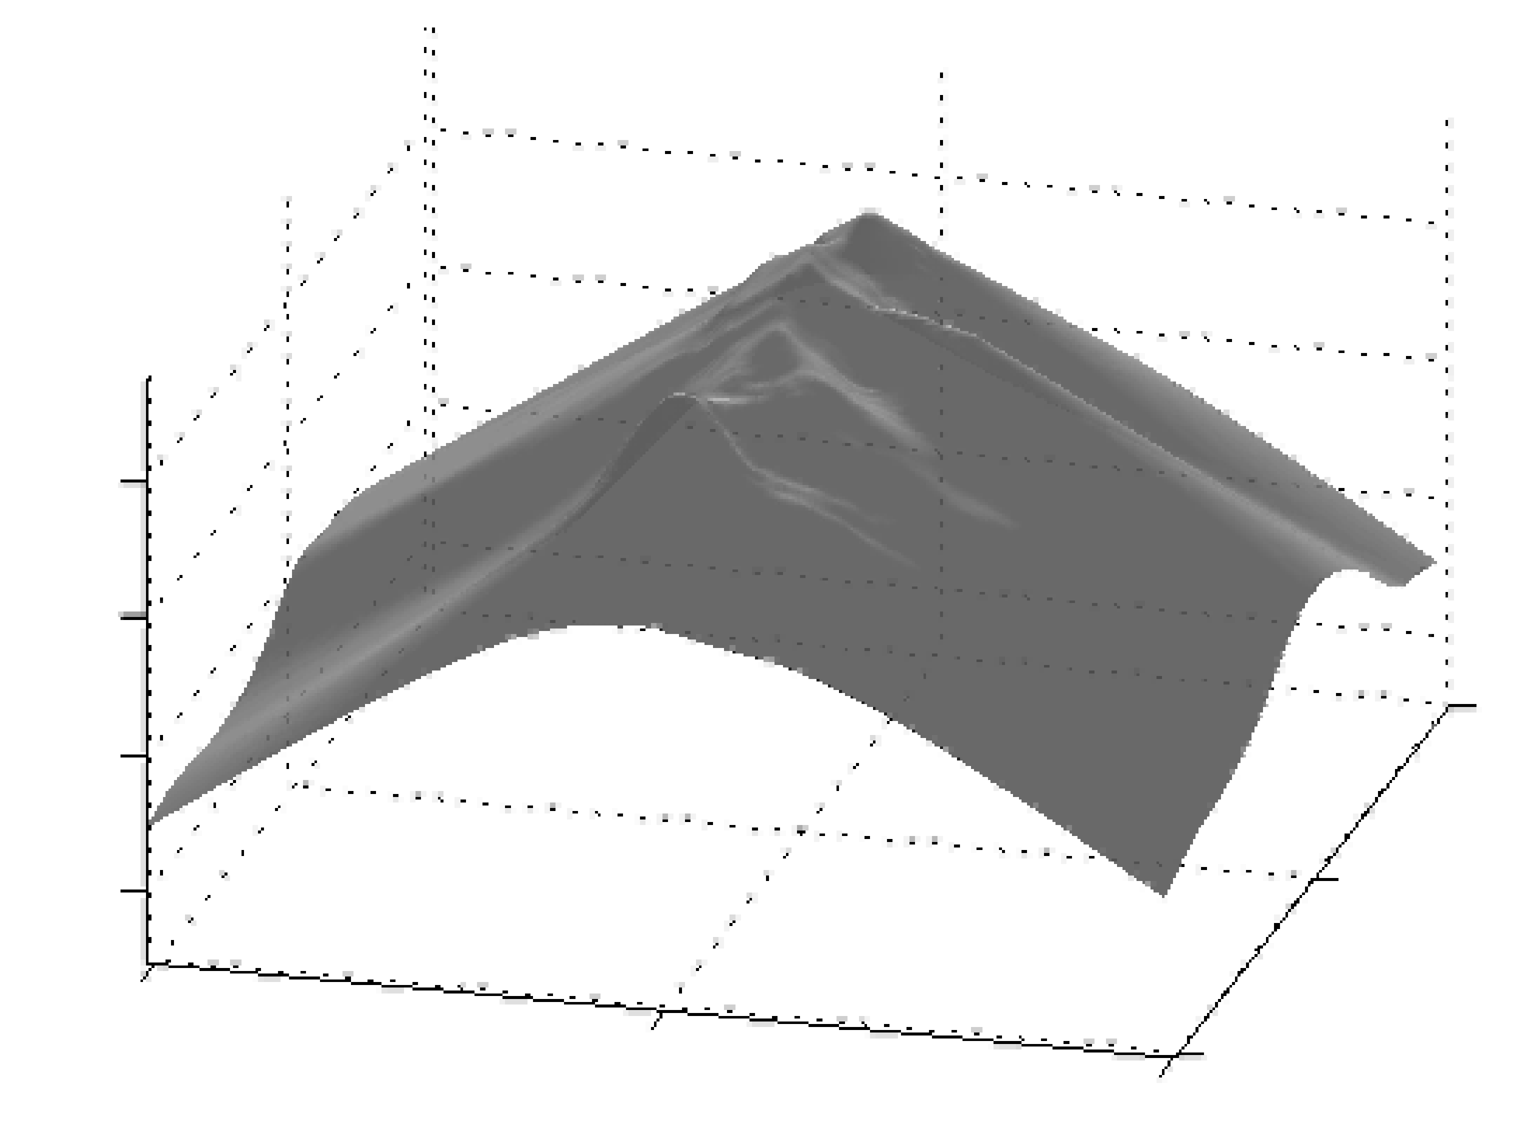

Statistically synthesized embedding functions

Cremers, IEEE PAMI 2006

1. Low-dim. representation via PCA (*Leventon et al. '00, Tsai et al. '01*):

$$\phi_i(x) \approx \phi_0(x) + \sum_{j=1}^{m} \alpha_{ij}\, \psi_j(x)$$

$$\alpha_{ij} = \int (\phi_i - \phi_0)\, \psi_j \, dx \qquad \boldsymbol{\alpha}_i = (\alpha_{i1}, \dots, \alpha_{im})$$

2. Autoregressive model for the shape coefficients:

$$\boldsymbol{\alpha}_t = \boldsymbol{\mu} + \sum_{i=1}^{k} A_i\, \boldsymbol{\alpha}_{t-i} + \boldsymbol{\eta}$$

3. Synthesize shape vectors and embedding surfaces:

$$\phi_t(x) = \phi_0(x) + \boldsymbol{\alpha}_t^\top\, \boldsymbol{\psi}(x)$$

Cremers, IEEE PAMI 2006

1. Low-dim. representation via PCA (*Leventon et al. '00, Tsai et al. '01*):

$$\phi_i(x) \approx \phi_0(x) + \sum_{j=1}^{m} \alpha_{ij}\,\psi_j(x)$$

$$\alpha_{ij} = \int (\phi_i - \phi_0)\,\psi_j\,dx \qquad \boldsymbol{\alpha}_i = (\alpha_{i1}, \dots, \alpha_{im})$$

2. Autoregressive model for the shape coefficients:

$$\boldsymbol{\alpha}_t = \boldsymbol{\mu} + \sum_{i=1}^{k} A_i\,\boldsymbol{\alpha}_{t-i} + \boldsymbol{\eta}$$

3. Synthesize shape vectors and embedding surfaces:

$$\phi_t(x) = \phi_0(T_{\theta_t}\,x) + \boldsymbol{\alpha}_t^\top\,\boldsymbol{\psi}(T_{\theta_t}\,x)$$

Cremers, IEEE PAMI 2006

Bayesian Aposteriori Maximization:

$$\hat{\boldsymbol{\alpha}}_t, \hat{\theta}_t = \arg\max_{\boldsymbol{\alpha}_t, \theta_t} \mathcal{P}\left(\boldsymbol{\alpha}_t, \theta_t \,\middle|\, I_t, \hat{\boldsymbol{\alpha}}_{1:t-1}, \hat{\theta}_{1:t-1}\right)$$

$$= \arg\max_{\boldsymbol{\alpha}_t, \theta_t} \mathcal{P}\left(I_t \,\middle|\, \boldsymbol{\alpha}_t, \theta_t\right) \mathcal{P}\left(\boldsymbol{\alpha}_t \,\middle|\, \hat{\boldsymbol{\alpha}}_{1:t-1}\right)$$

$$E(\boldsymbol{\alpha}_t, \theta_t) = -\log \mathcal{P} = E_{dat}(\boldsymbol{\alpha}_t, \theta_t, I_t) + E_{dyn}(\boldsymbol{\alpha}_t, \hat{\boldsymbol{\alpha}}_{1:t-1})$$

$$E_{dyn} = \frac{1}{2}(\boldsymbol{\alpha}_t - \boldsymbol{v})^\top C^{-1}(\boldsymbol{\alpha}_t - \boldsymbol{v}), \qquad \boldsymbol{v} \equiv \mu + \sum_{i=1}^{k} A_i \hat{\boldsymbol{\alpha}}_{t-i}$$

nodel

Optimization by gradient descent: $\frac{d\boldsymbol{\alpha}_t}{d\tau} = -\frac{\partial E}{\partial \boldsymbol{\alpha}_t}, \quad \frac{d\theta_t}{d\tau} = -\frac{\partial E}{\partial \theta_t}$

Dynamical Shape Priors

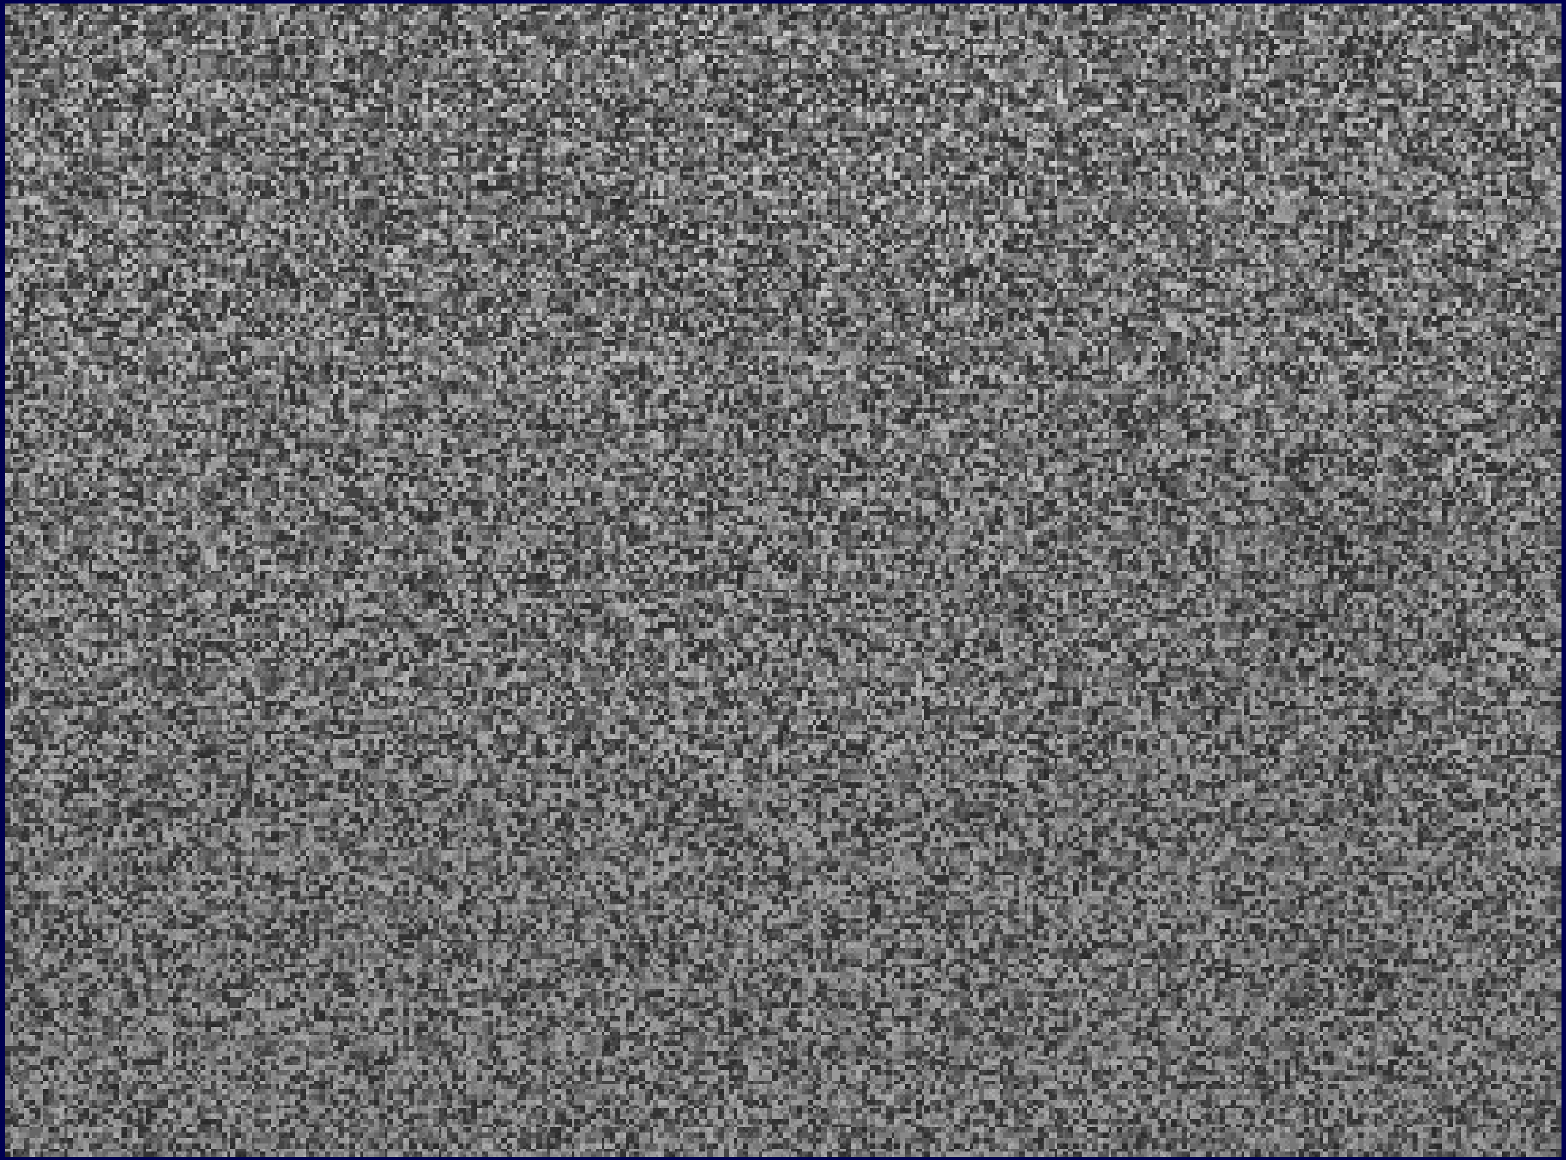

Cremers, IEEE PAMI '06

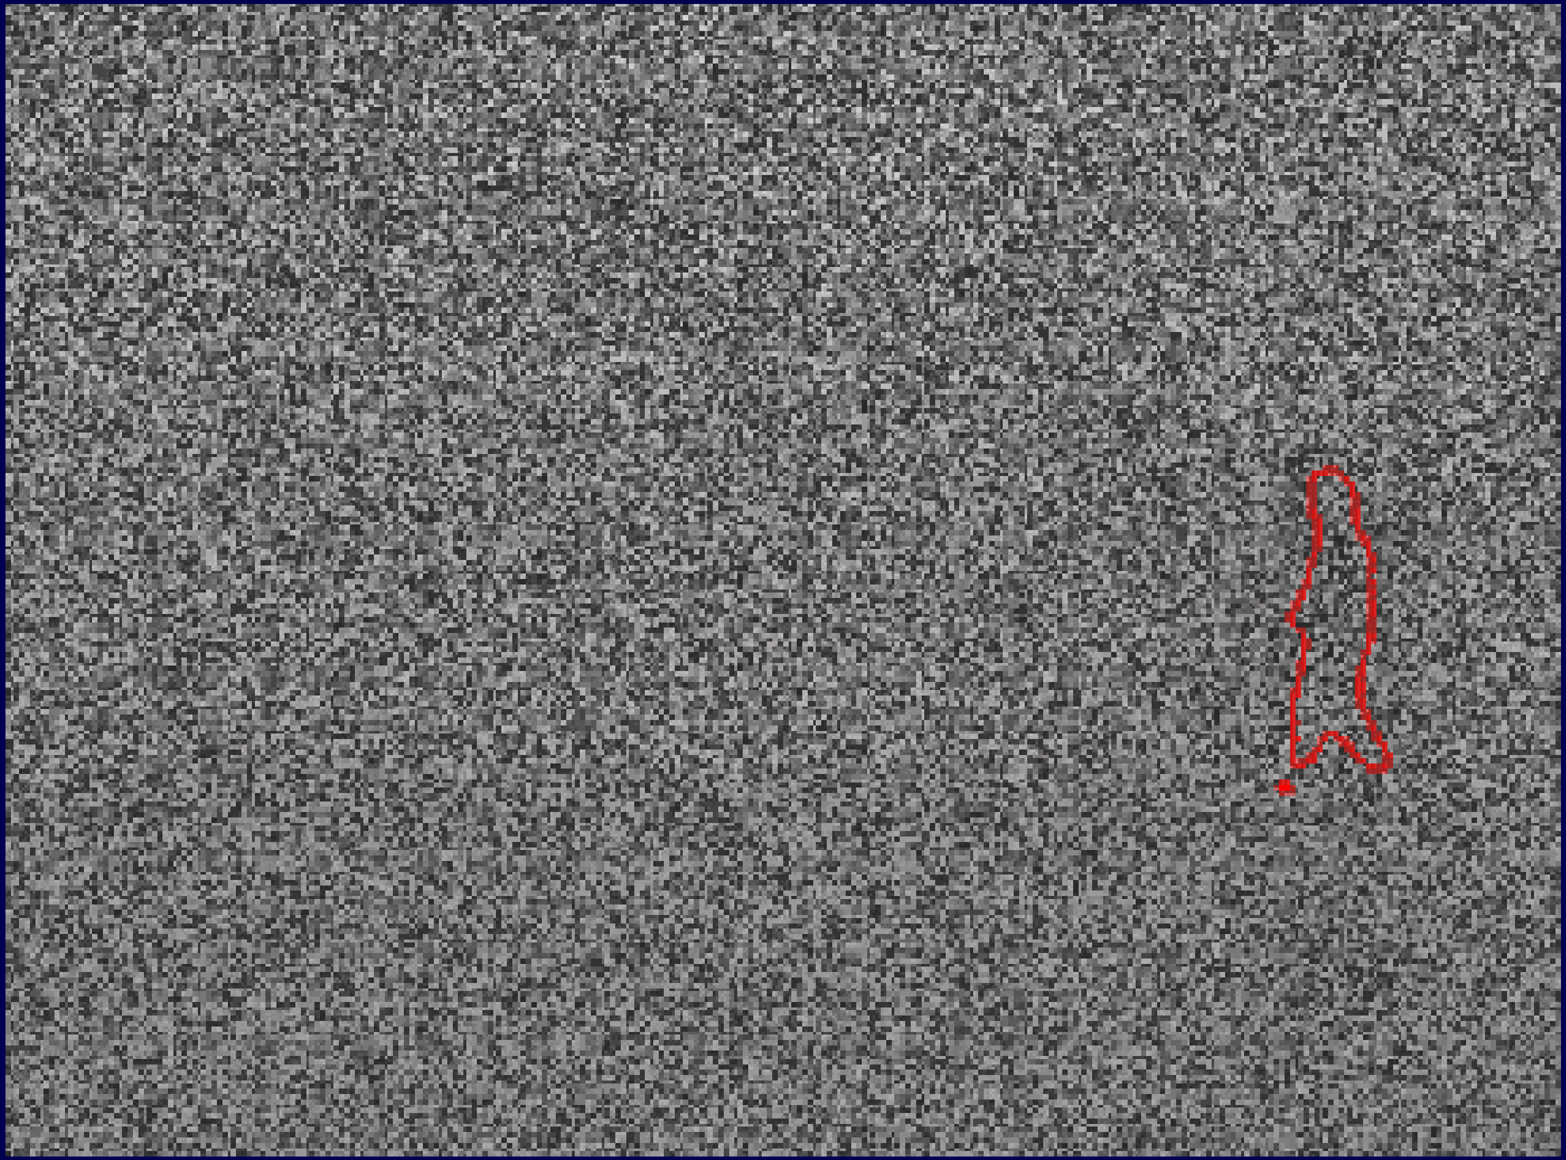

Cremers, IEEE PAMI ’06

Robustness to Speed Variation

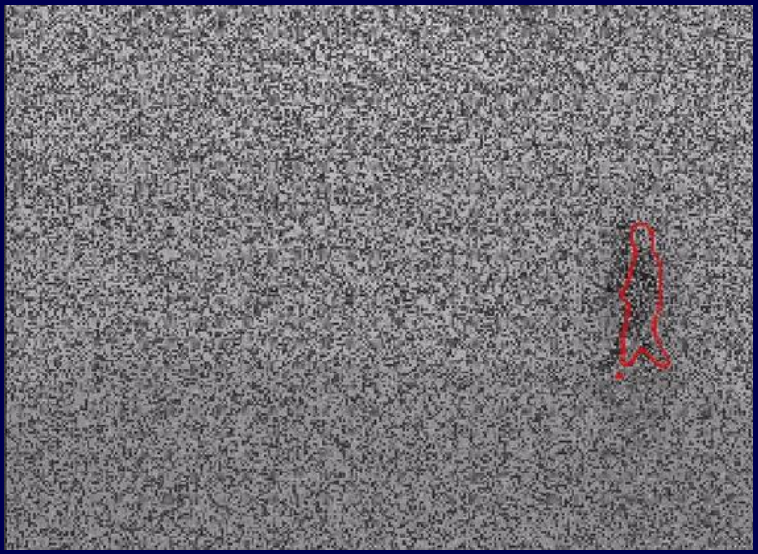

2 x faster

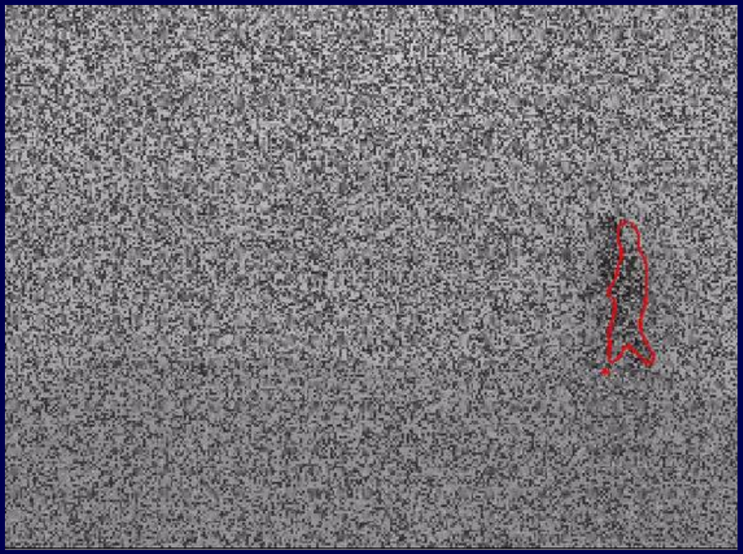

3 x faster

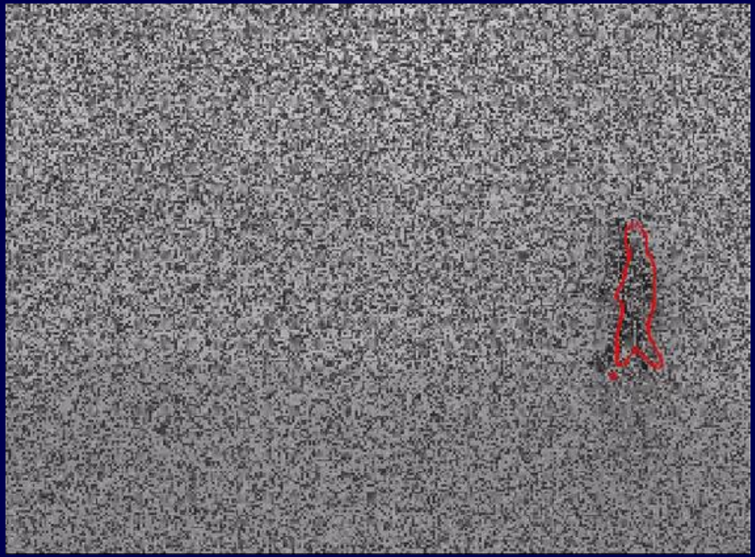

4 x faster

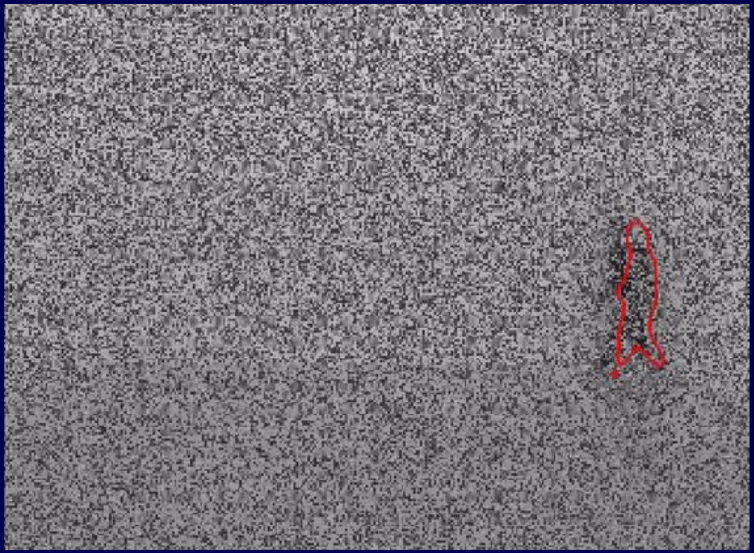

5 x faster

3D Heart Tracking with Dynamical Prior

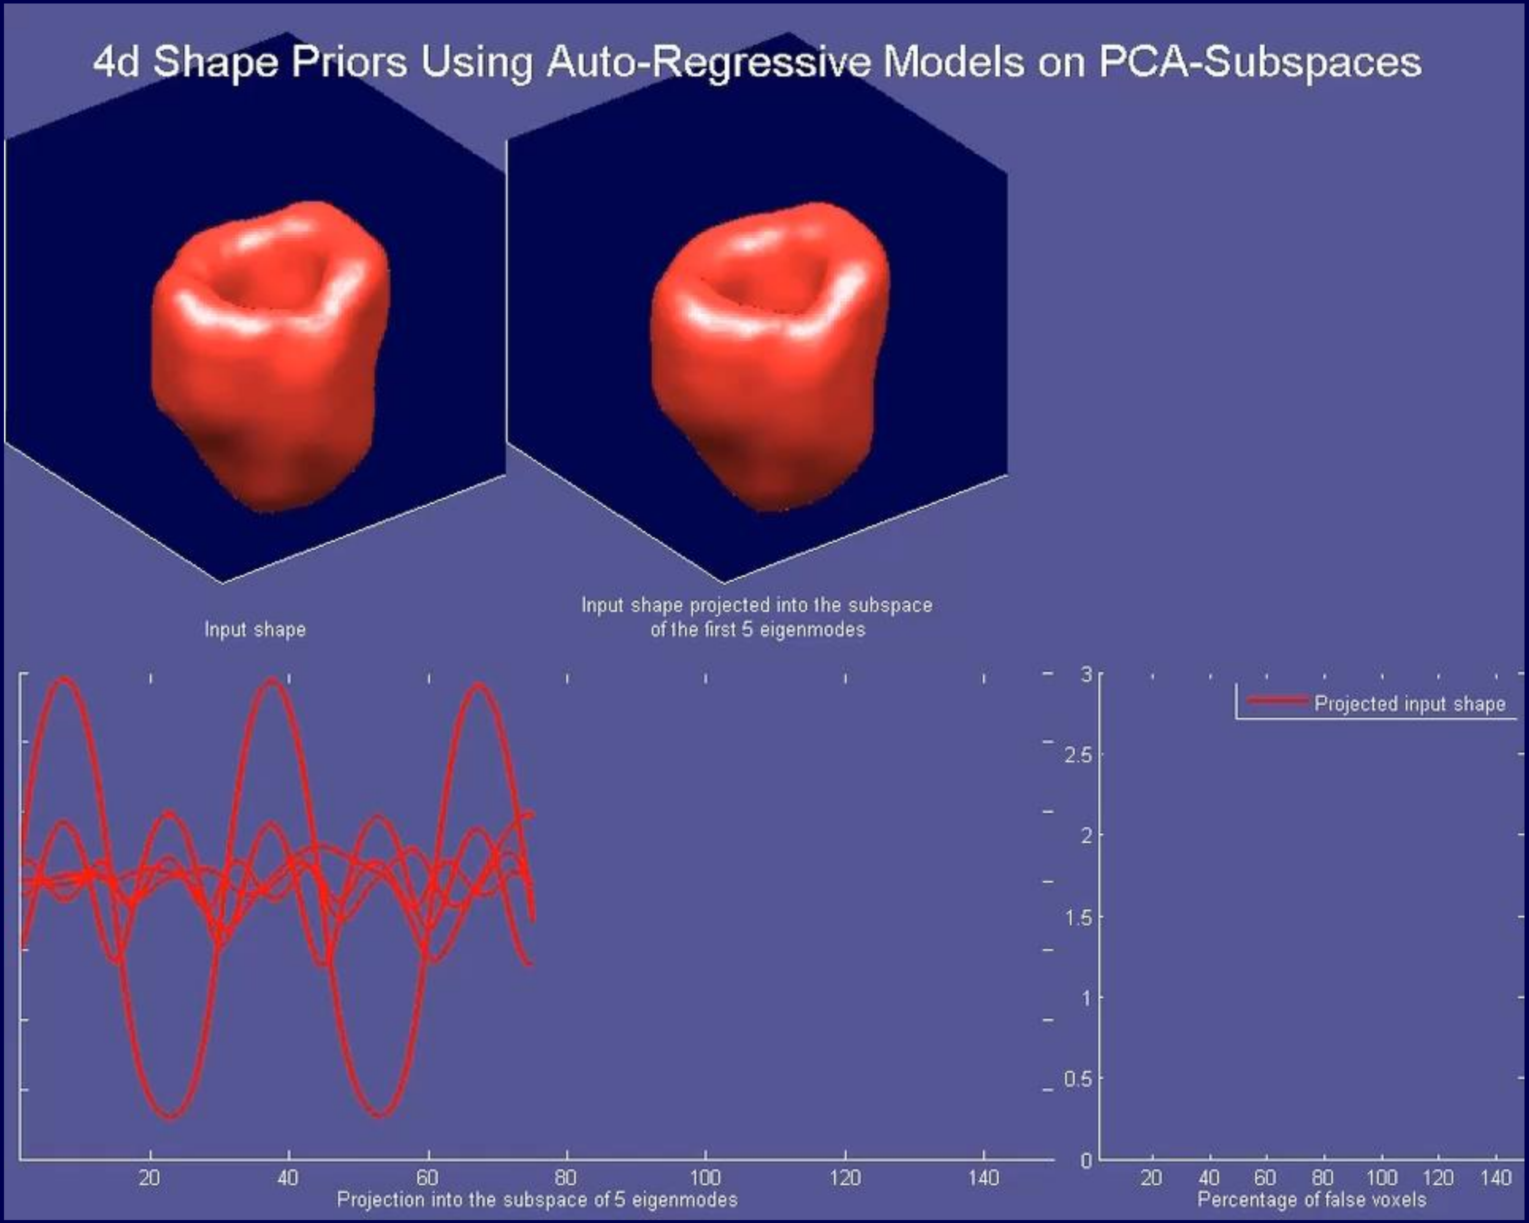

Implementation: Timo Kohlberger

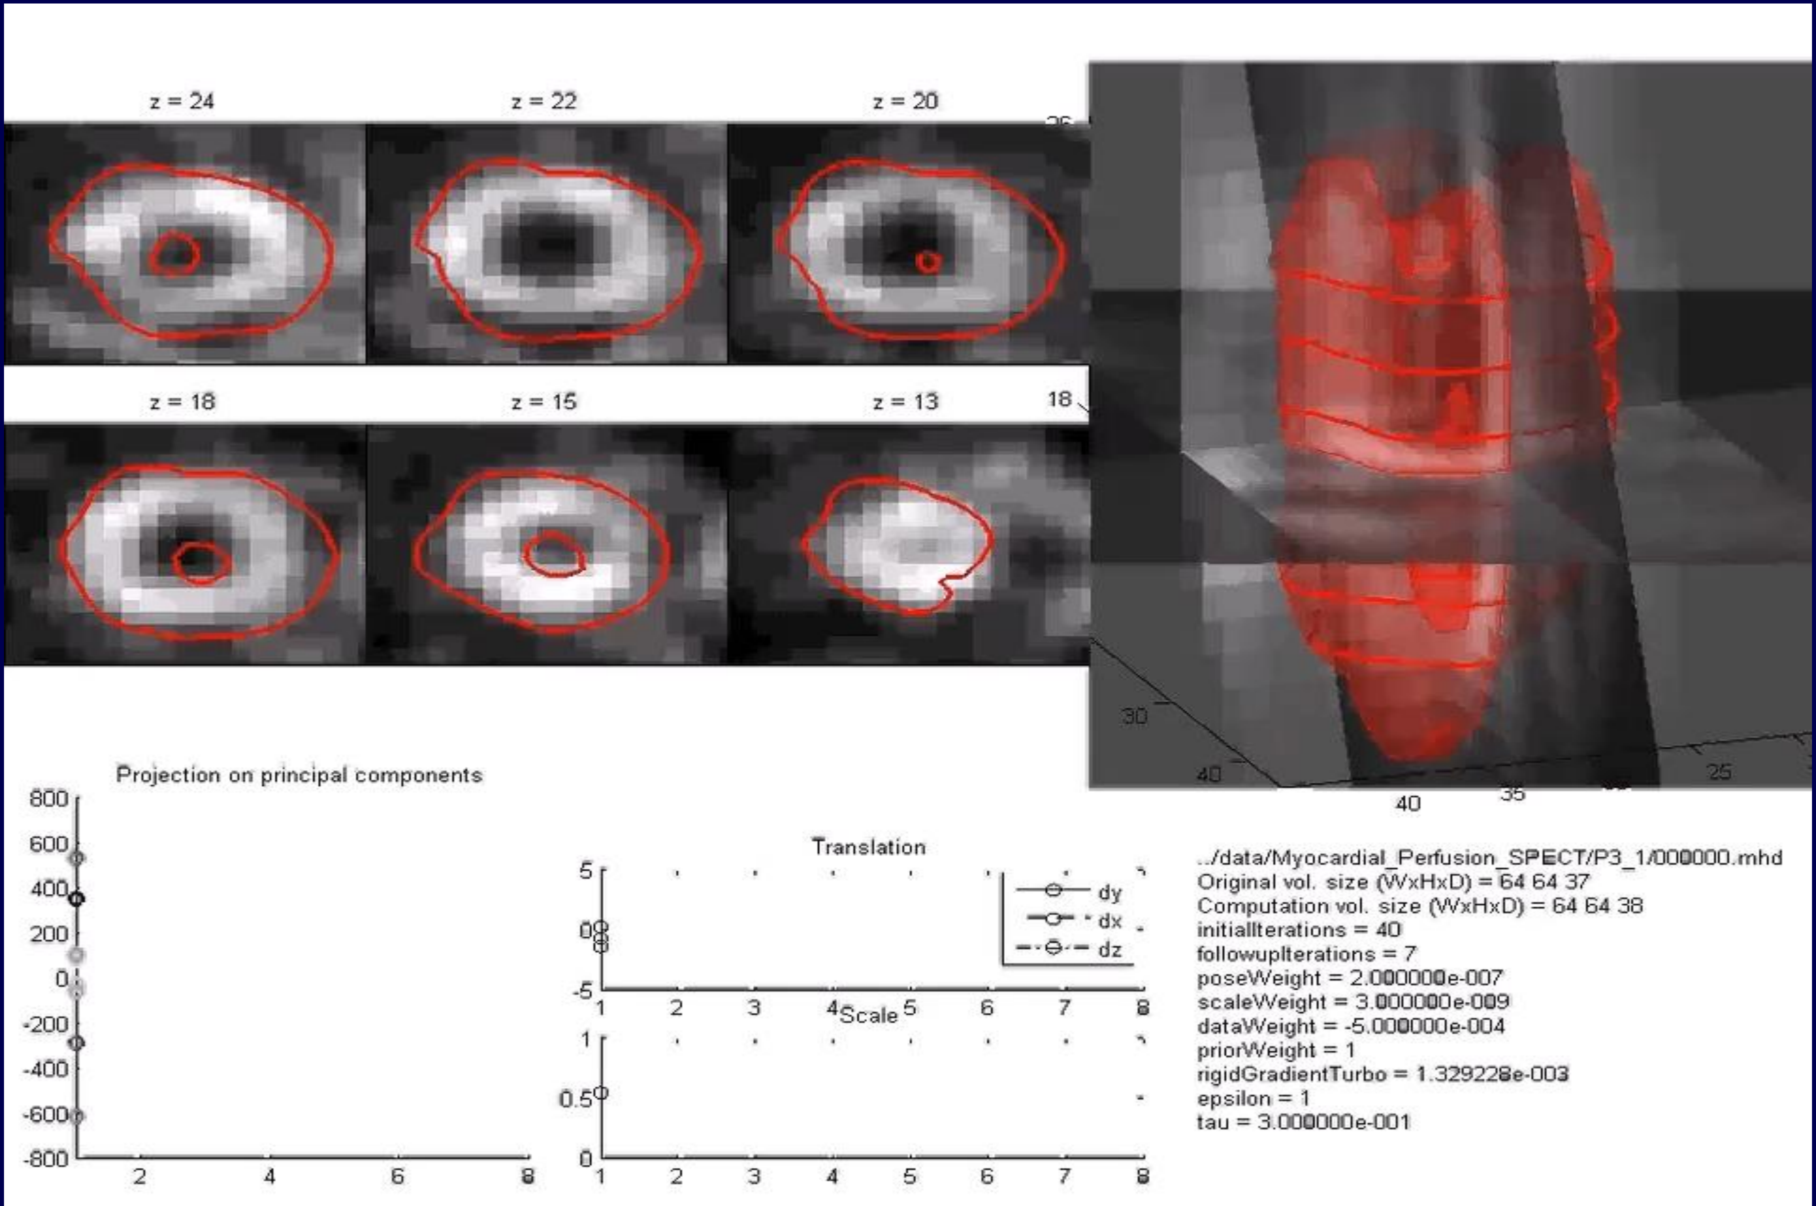

Implementation: Timo Kohlberger

Dynamical Shape Priors

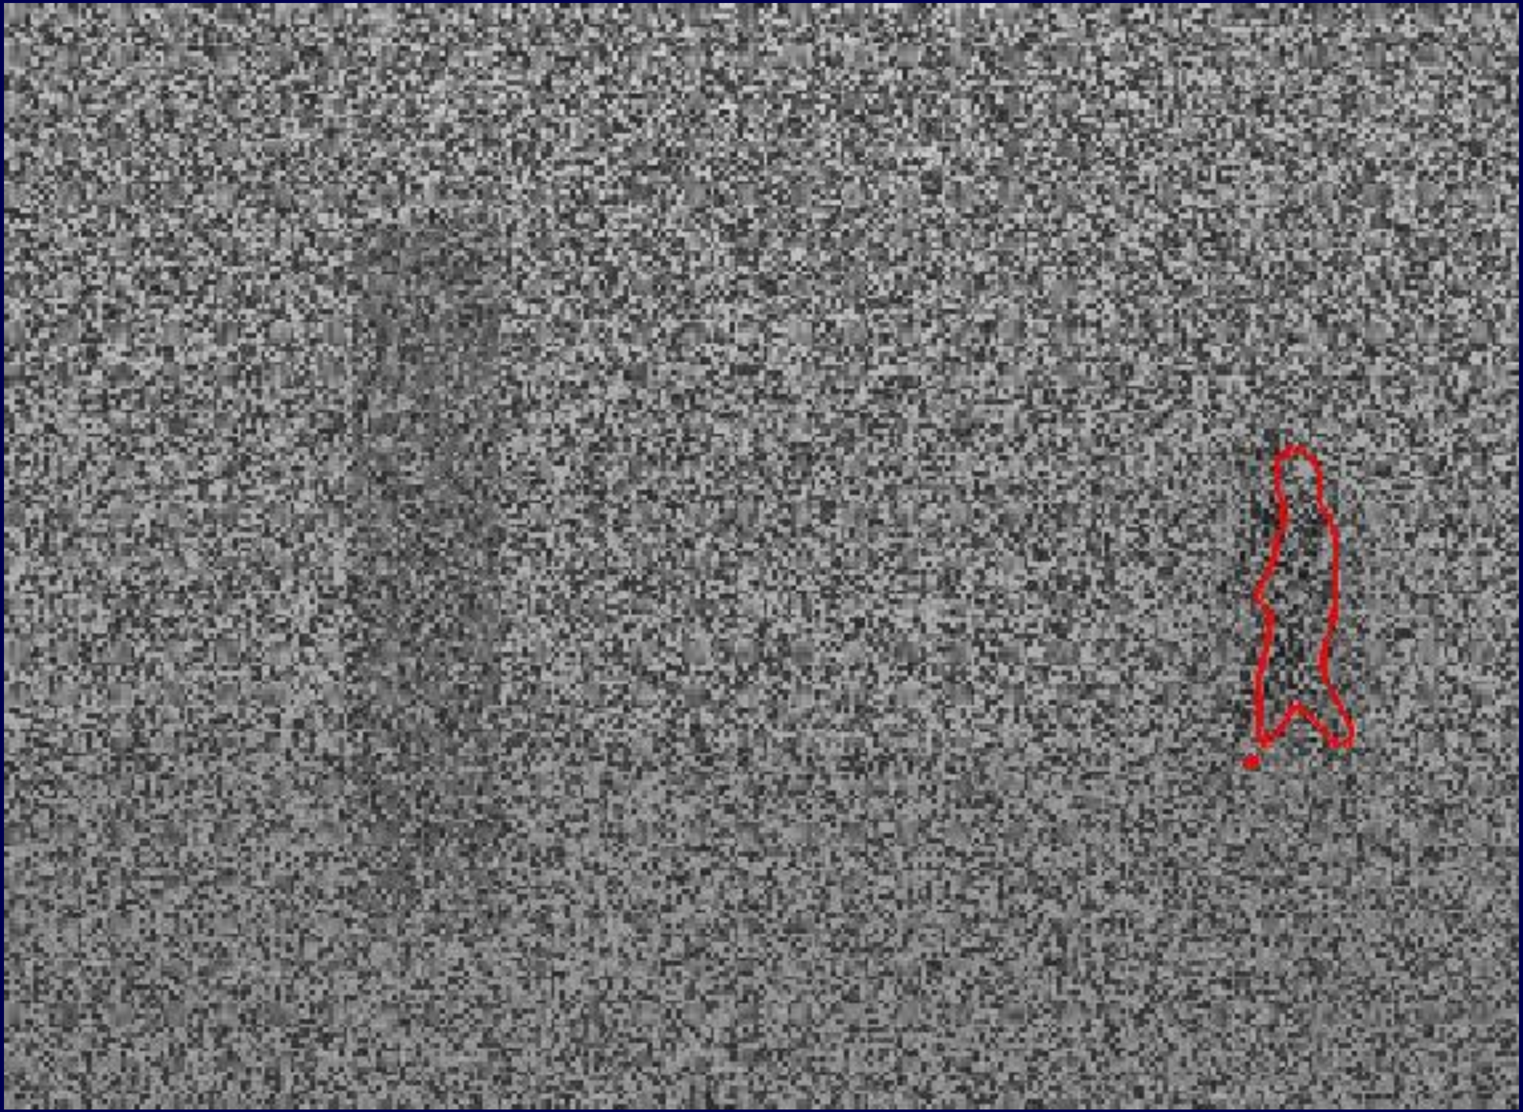

Dynamical shape prior

Cremers, IEEE PAMI '06

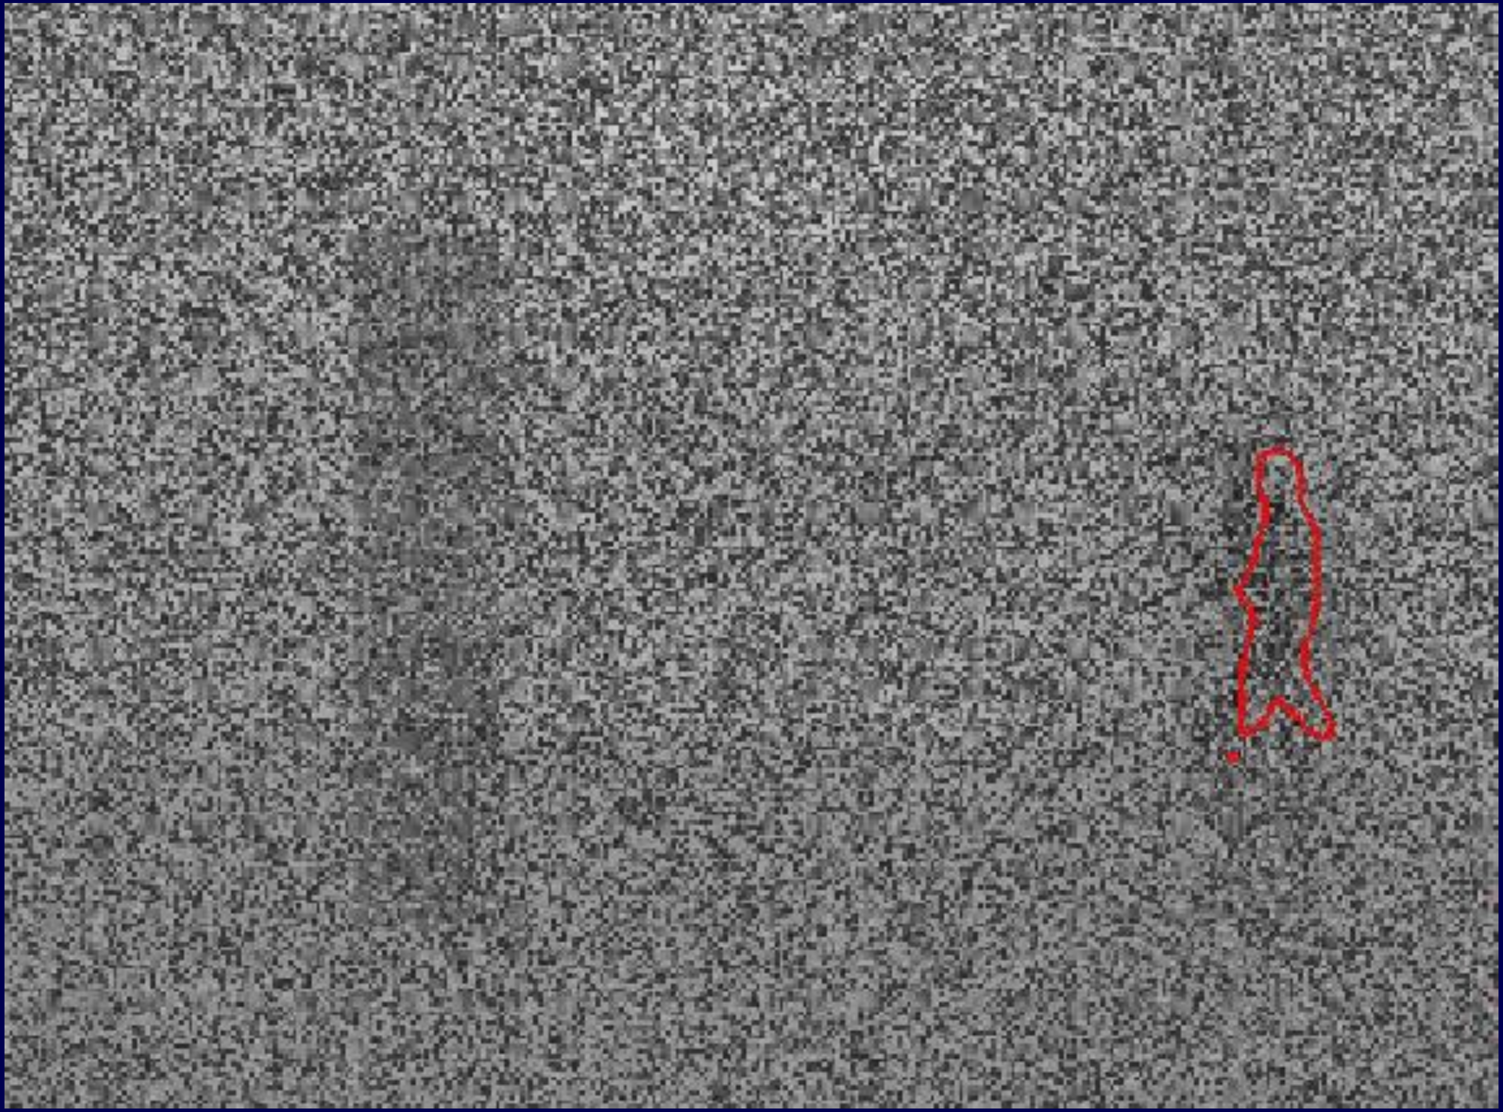

Dynamical prior of shape and translation

Cremers, IEEE PAMI '06

Bayesian inference

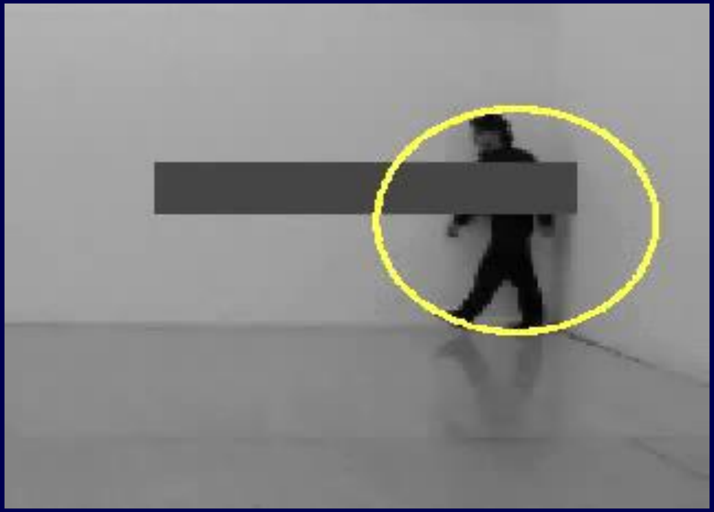

Statistical shape priors

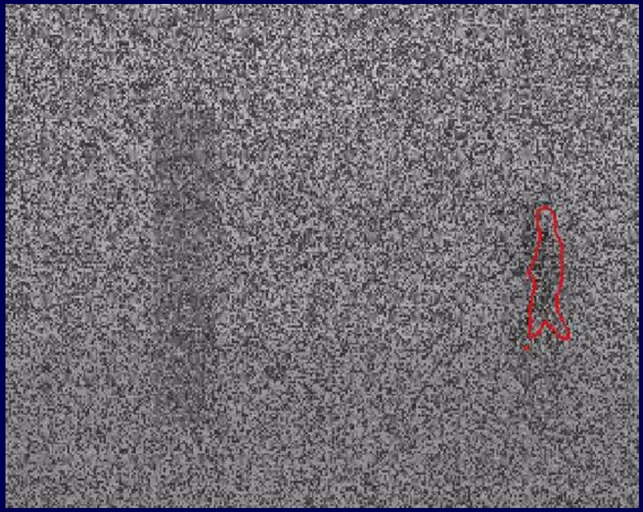

Dynamical shape priors

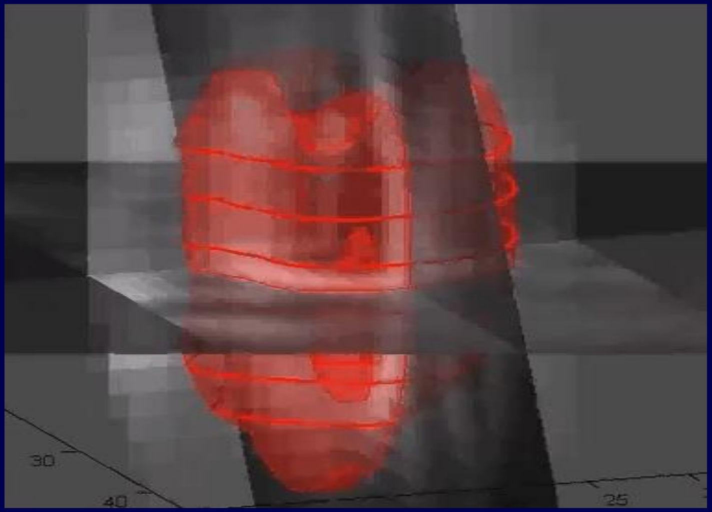

3D heart tracking