A Globally Optimal Data-driven Approach for Image Distortion Estimation

Yuandong Tian and *Srinivasa G. Narasimhan*, Carnegie Mellon University

Problem Statement

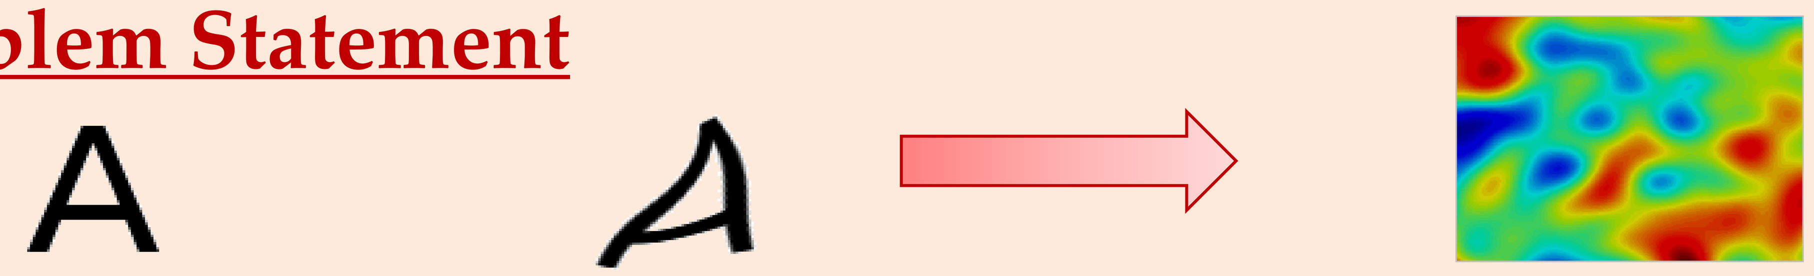

Template $T = I_0$ Distorted image I_p Dense Warping Field $W(x; p)$

Distortion Model

$$W(x; p) = x + B(x)p$$

Bases $B(x)$ → •Linear & Rigid •Spatial nonlinearity •Measurement-based

Parameters p → Fewer dimensions

Previous Works	Generative	Discriminative
Method	Optimization	Direct Mapping
Disadvantages	Local minima	Exponential #samples

Contribution

1. Guaranteed worst-case convergence to **the global optimum** using only $O(C^d \log(1/\varepsilon))$ samples.
2. Decoupling accuracy $1/\varepsilon$ from the dimension d of the parameter space.

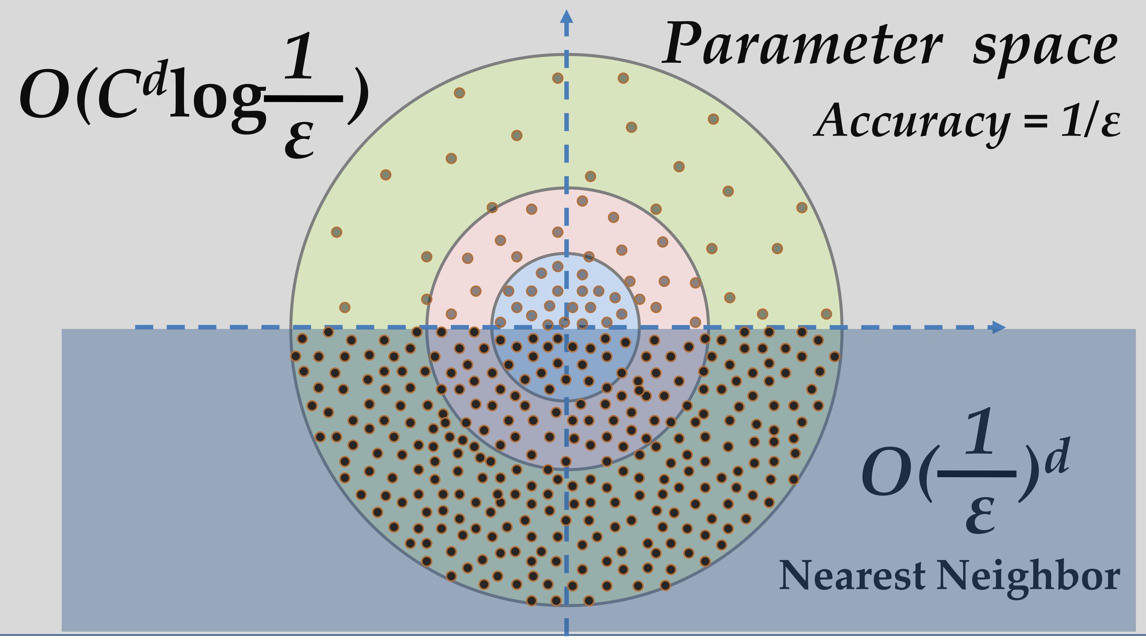

Experimental Result

More results on http://www.cs.cmu.edu/~yuandong/

Rectification

Tracking

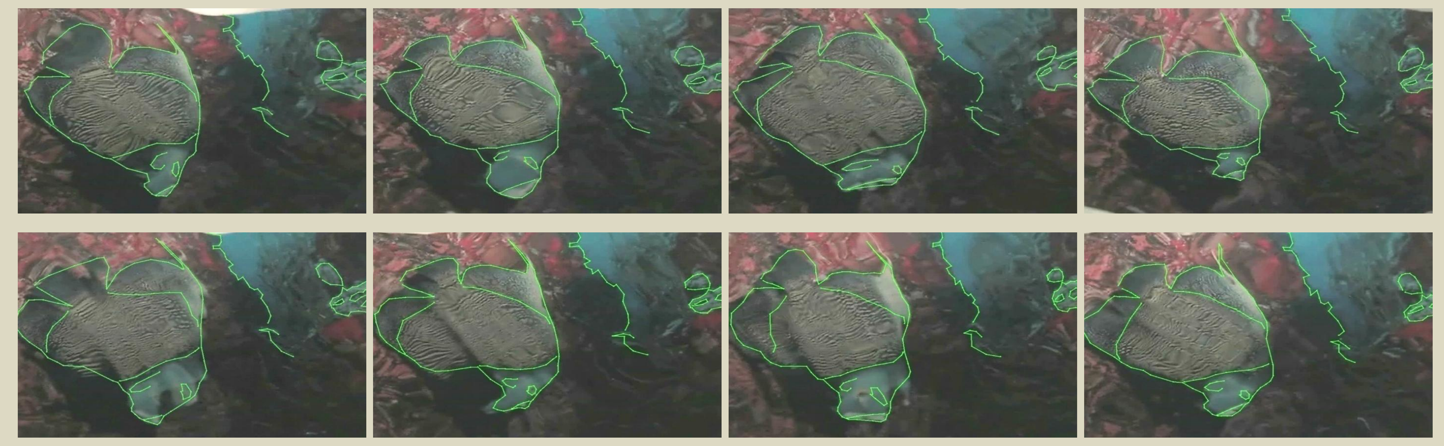

Surface Reconstruction

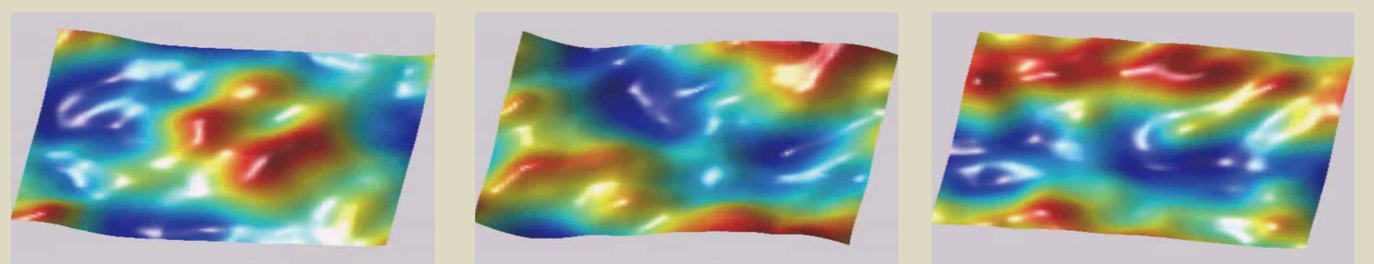

Cloth Deformation

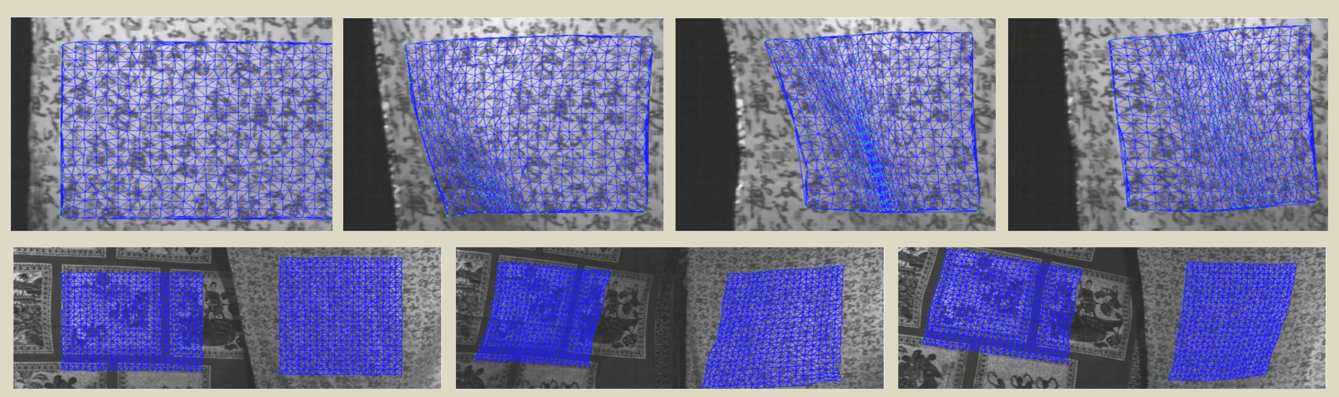

[Data from J.Taylor et.al, CVPR 2010]

Simulation

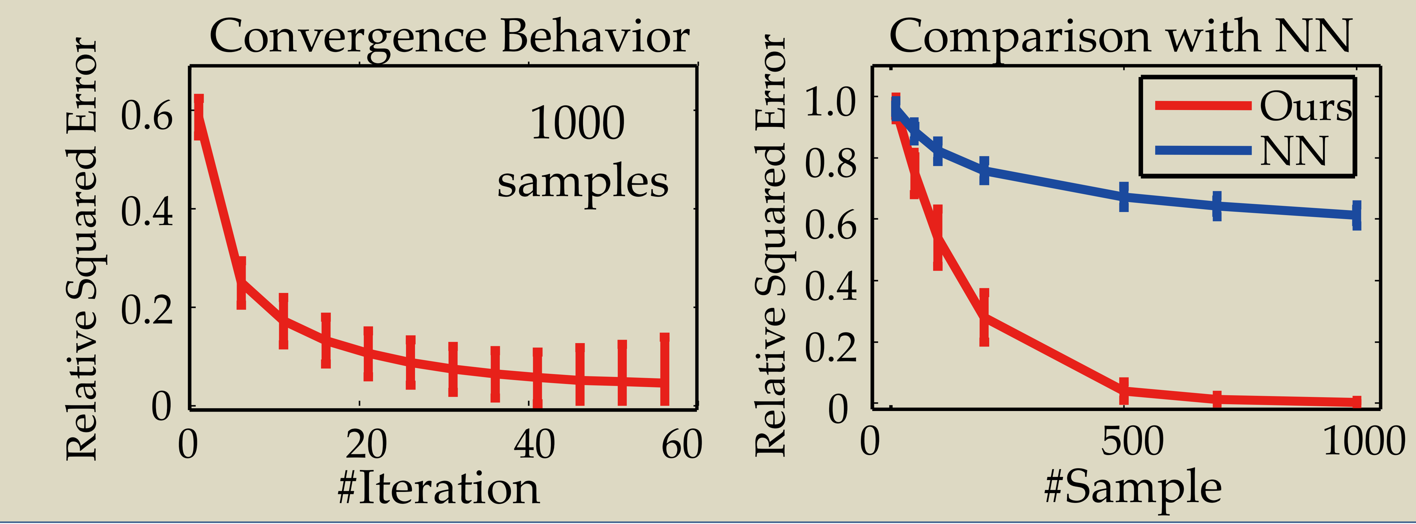

The Algorithm

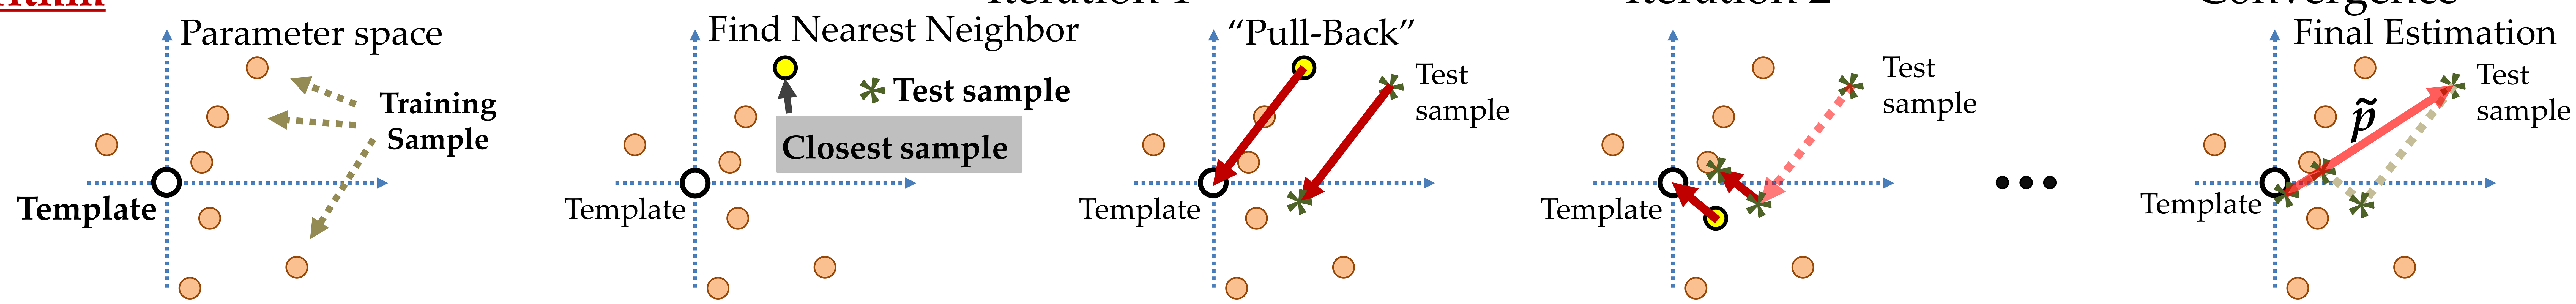

Effectiveness of Nearest Neighbor

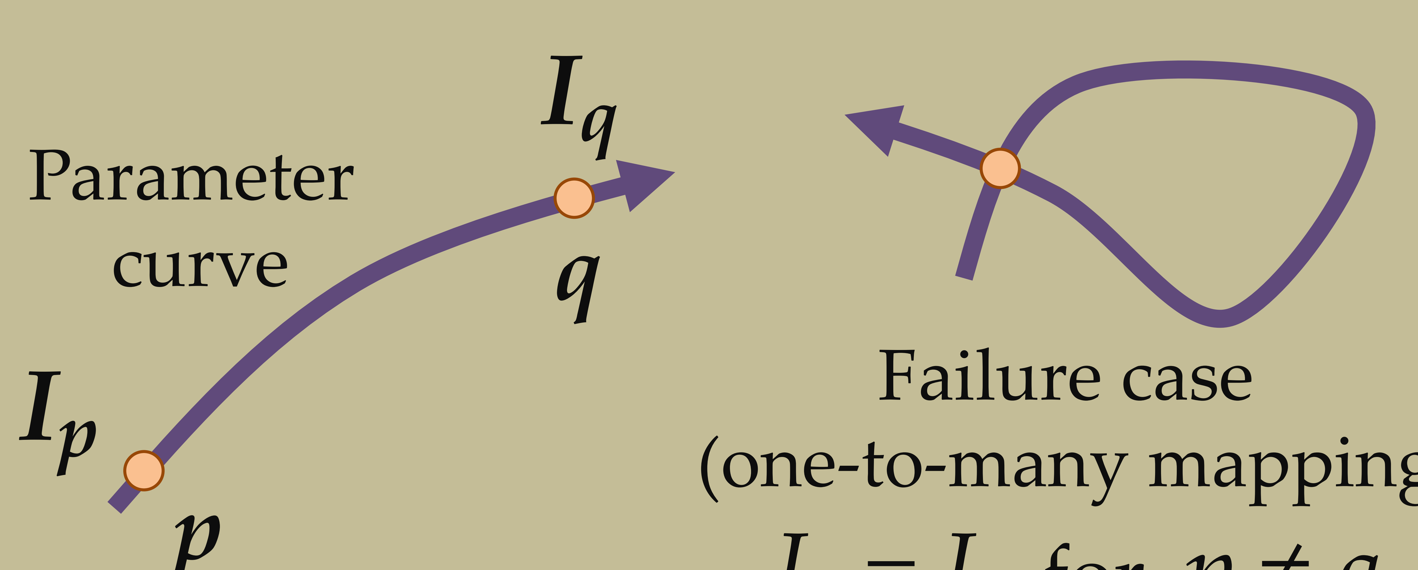

$I_p = I_q$ for $p \neq q$

Lipchitz Continuity:

$$L_1 \| I_p - I_q \| \leq \| p - q \| \leq L_2 \| I_p - I_q \|$$

The Pull-back Operation

For invertible distortion (e.g. affine):

$$H(I_p, q) = \mathrm{Inv}(I_p, q) = I_{p-q}$$

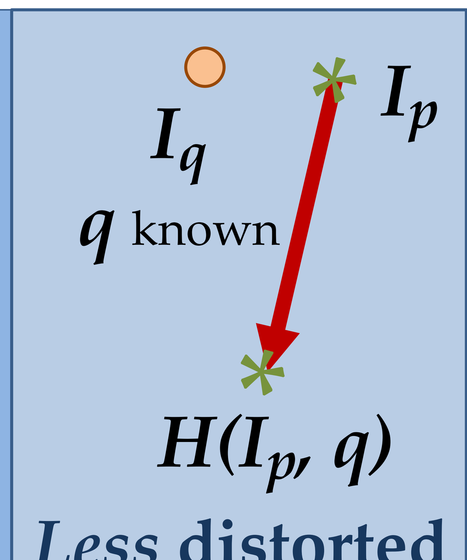

For non-invertible distortion:

In the case of $W(x; p) = x + B(x)p$

We prove $\| I_{p-q} - H(I_p, q) \| \leq R \| p - q \|$

A constant related to $\max(\| \nabla B(x) \|, \| \nabla T \|)$

Failure case:
Resampling artifacts.
E.g. Aliasing

Sample Distribution

For any I_p with $\| p \| \leq r$, there exists I_{tr} : $\| I_{tr} - I_p \| \leq \beta r / L_2$, with $\beta < 1$.

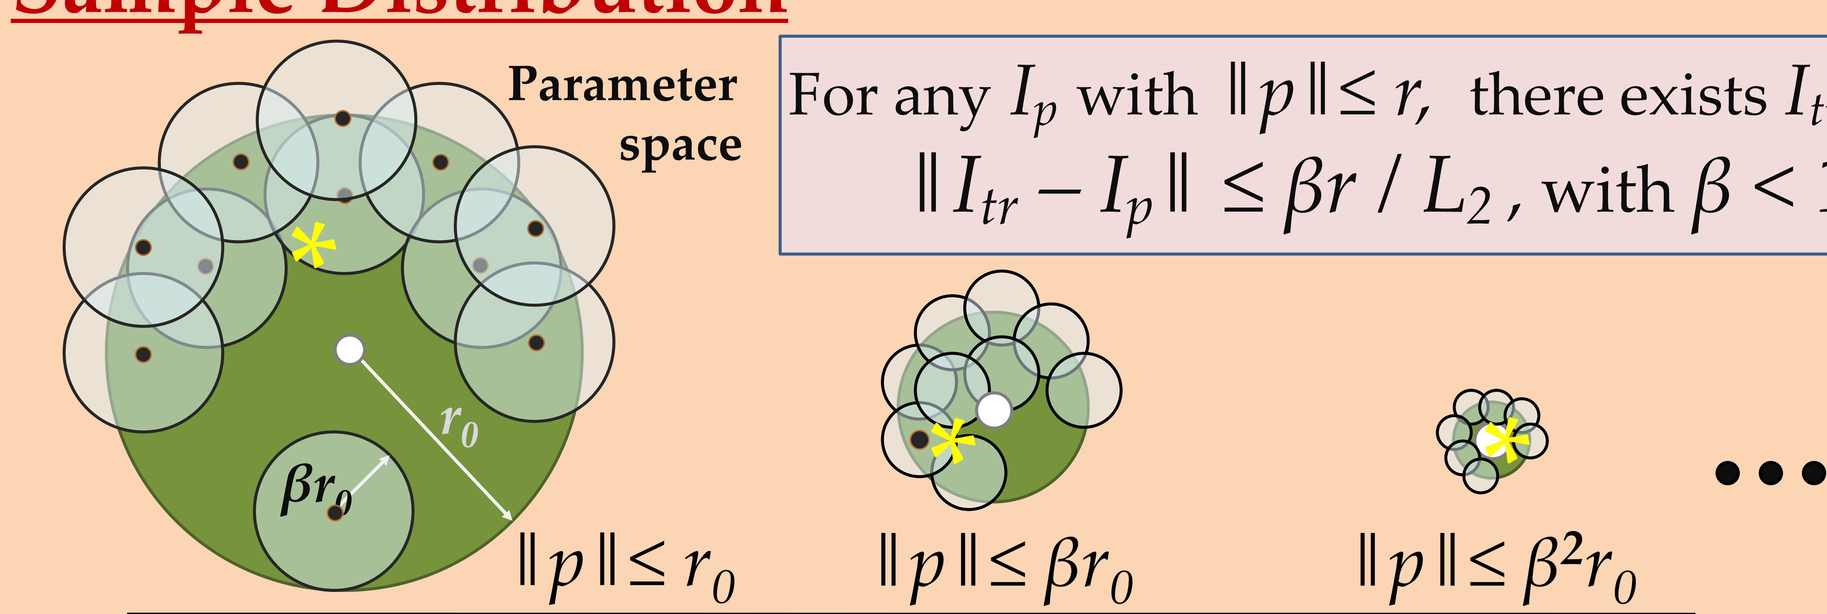

#Sample: $O(L_2/\beta L_1)^d$ $O(L_2/\beta L_1)^d$ $O(L_2/\beta L_1)^d$

Total number of samples for accuracy $1/\varepsilon$:

$$N(\beta, \varepsilon) = O\left[\left(\frac{L_2}{\beta L_1}\right)^d \log_\beta \frac{r_0}{\varepsilon}\right] = O(C^d \log\frac{1}{\varepsilon})$$