The Multiscale Structure of Non-Differentiable Image Manifolds

Michael Wakin

Electrical Engineering
Colorado School of Mines

Joint work with Richard Baraniuk, Hyeokho Choi, David Donoho

Models for Image Structure

- Not all N-pixel images are created equal

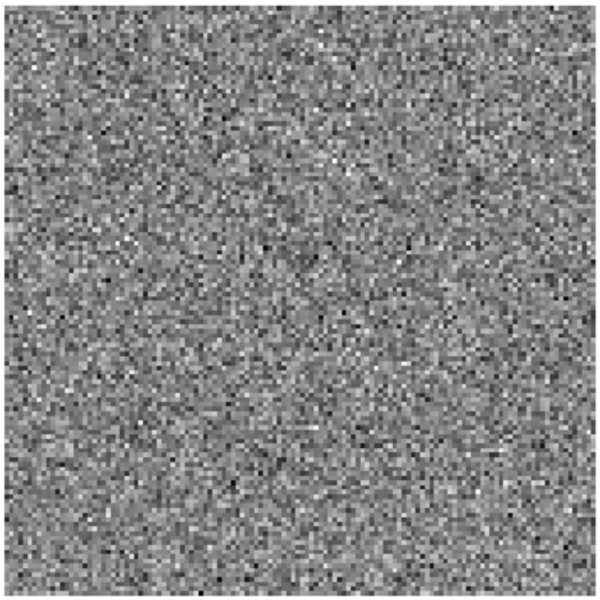

- Models capture ***concise*** structure
  - few degrees of freedom
  - permit effective denoising, compression, registration, detection, classification, segmentation, estimation, …

Geometry: Where are the Images?

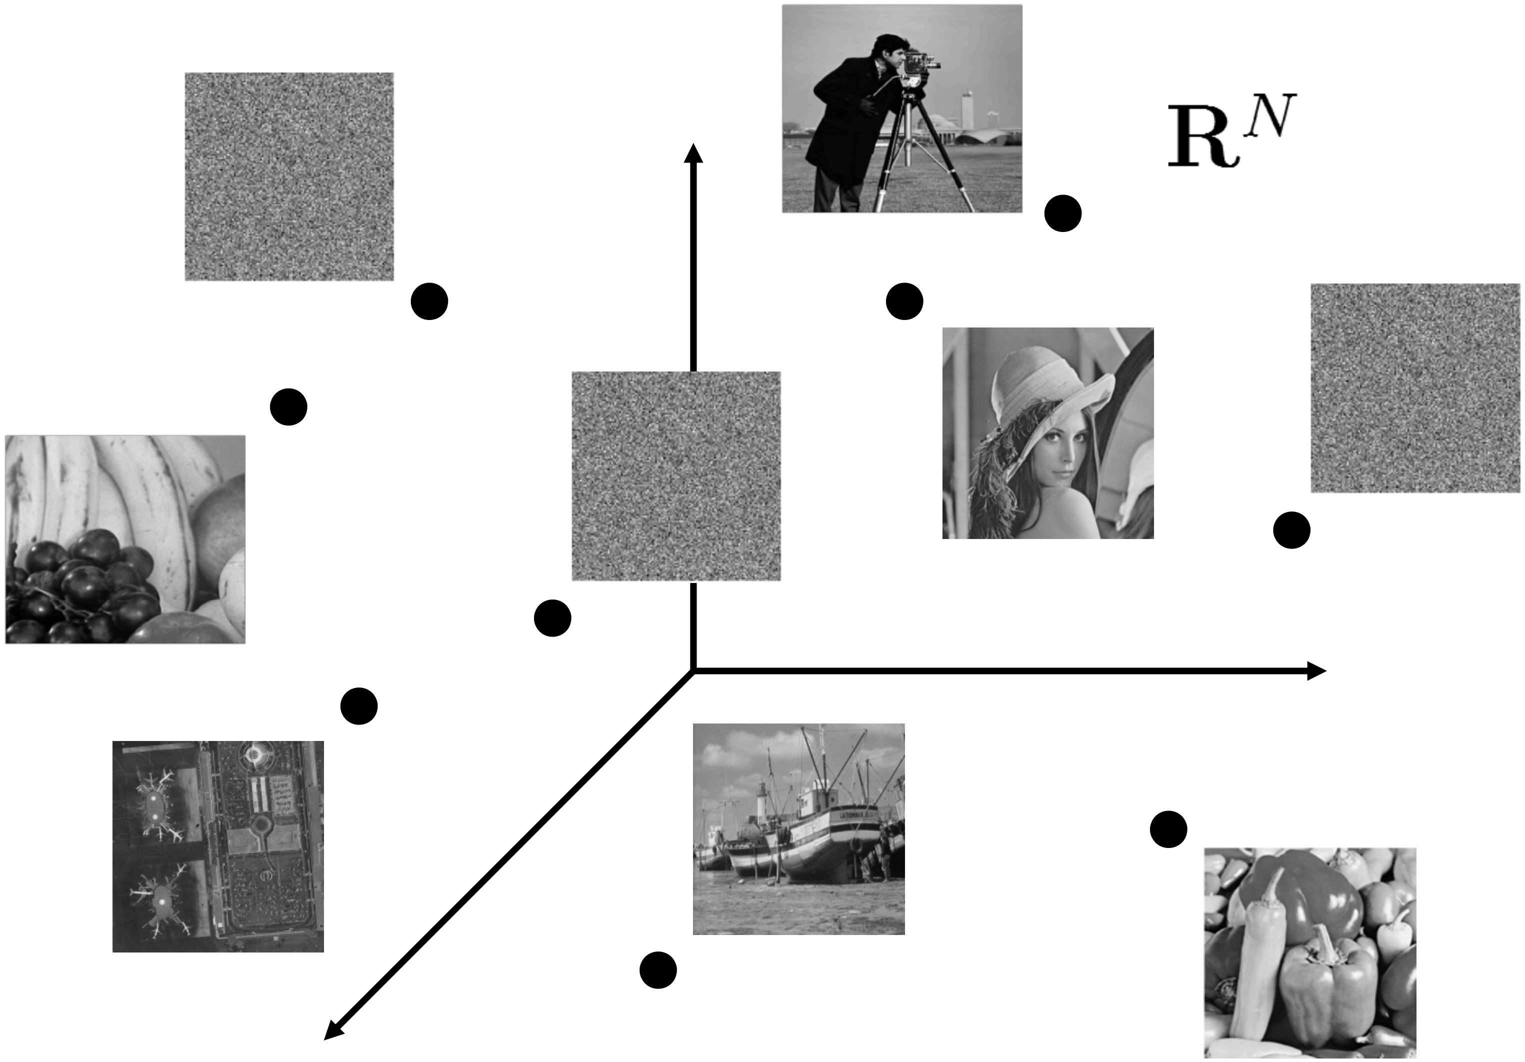

concise **models ⇔ *low-dimensional* geometry**

Linear Subspace Models

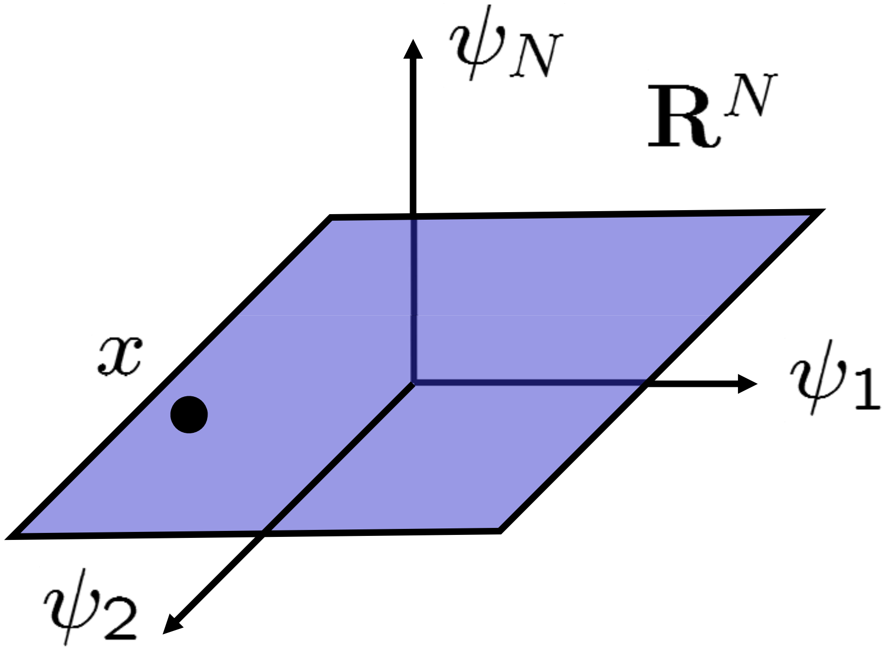

e.g., 2D Fourier basis with bandlimited images

Many Image Families are Highly Nonlinear

Sparse Models: Unions of Subspaces

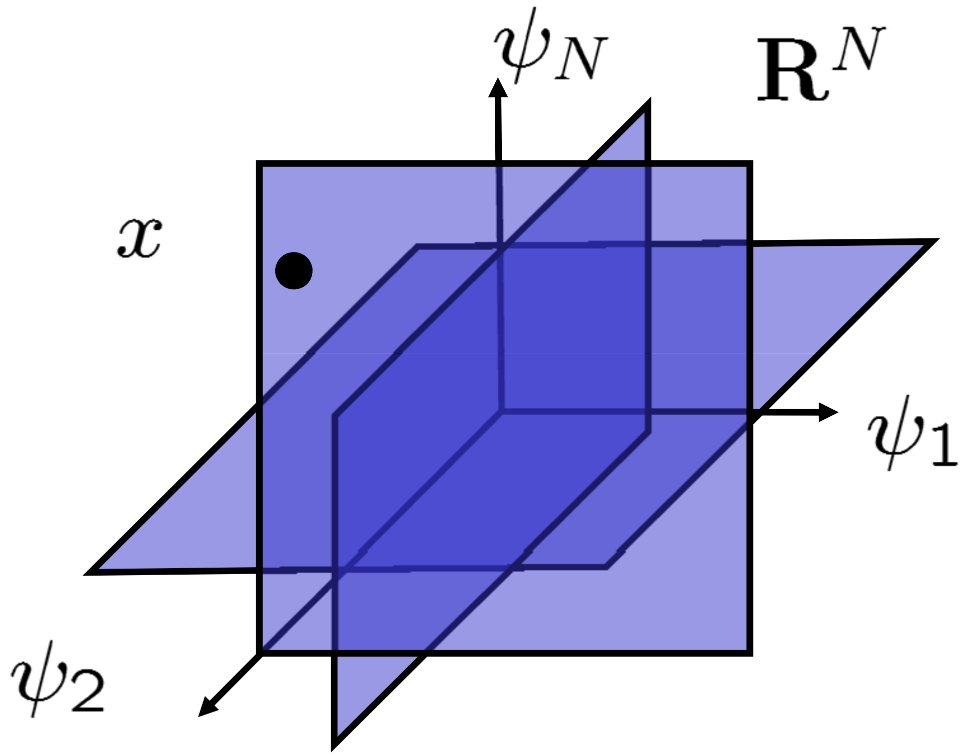

e.g., wavelet bases with piecewise smooth images

What more can we say about nonlinear signal families?

Manifold Models

- *K-dimensional parameter* $\theta \in \Theta$ captures degrees of freedom in signal $f_\theta \in R^N$

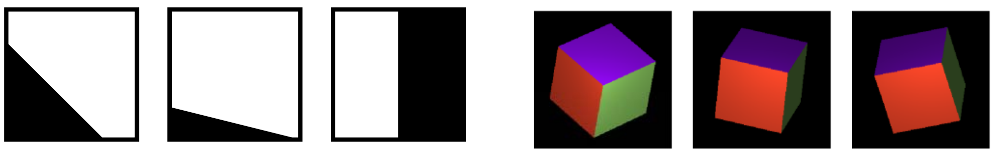

- Signal class $F = \{ f_\theta : \theta \in \Theta \}$ forms a *K-dimensional manifold*
  - also nonparametric collections: faces, handwritten digits, shape spaces, etc.

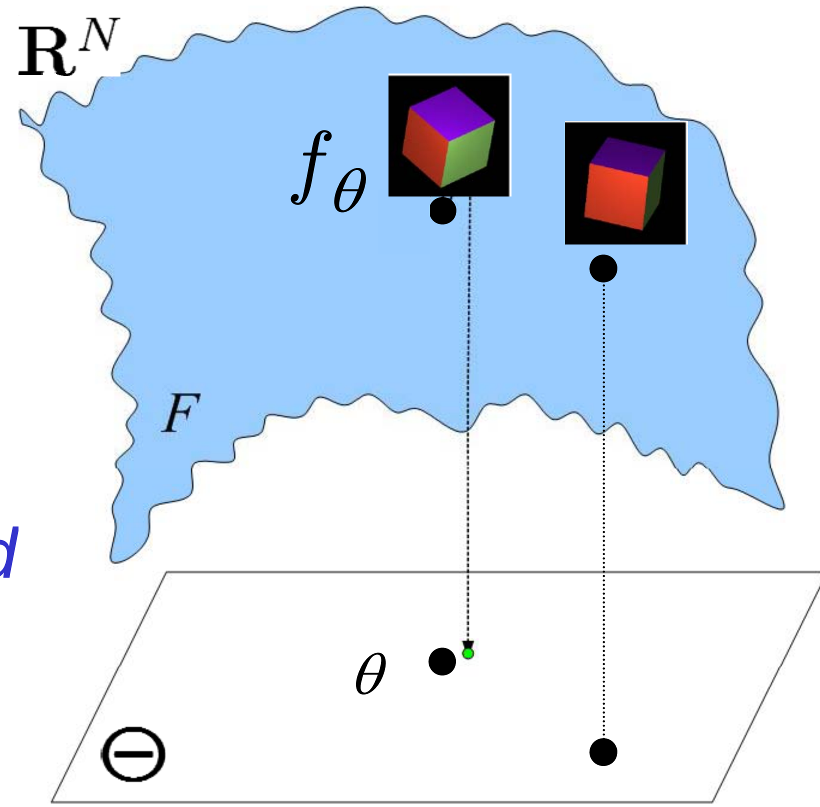

- Generally *nonlinear*

- Surprise: Often *non-differentiable*

Overview

- Motivating application: parameter estimation
- Non-differentiability from edge migration
- Parameter estimation (revisited)
- Non-differentiability from edge occlusion
- Manifolds in Compressive Sensing

Application: Parameter Estimation

- Given an observed image $I = f_\theta$, can we recover the underlying articulation parameters θ
  - efficiently, and
  - with high precision?

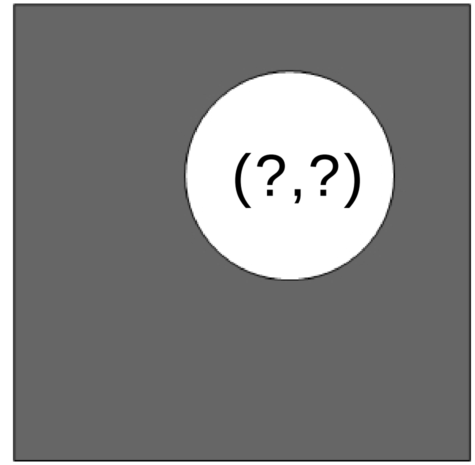

- Given a *noisy* image $I \approx f_\theta$, can we do the same?

- Relevant in pose estimation, image registration, computer vision, edge detection, ...

Newton's Method

- Optimization problem $\min_{\theta' \in \Theta} \|I - f_{\theta'}\|_2$

- For a *differentiable* manifold, project onto tangent planes

$$\theta^{(k+1)} \leftarrow \theta^{(k)} + [H(\theta^{(k)})]^{-1} J(\theta^{(k)})$$

$$J_i = \langle f_\theta - I, \tau_\theta^i \rangle$$

tangent

$$H_{ij} = \langle \tau_\theta^i, \tau_\theta^j \rangle$$

Newton's Method...

$$\min_{\theta' \in \Theta} \| I - f_{\theta'} \|_2$$

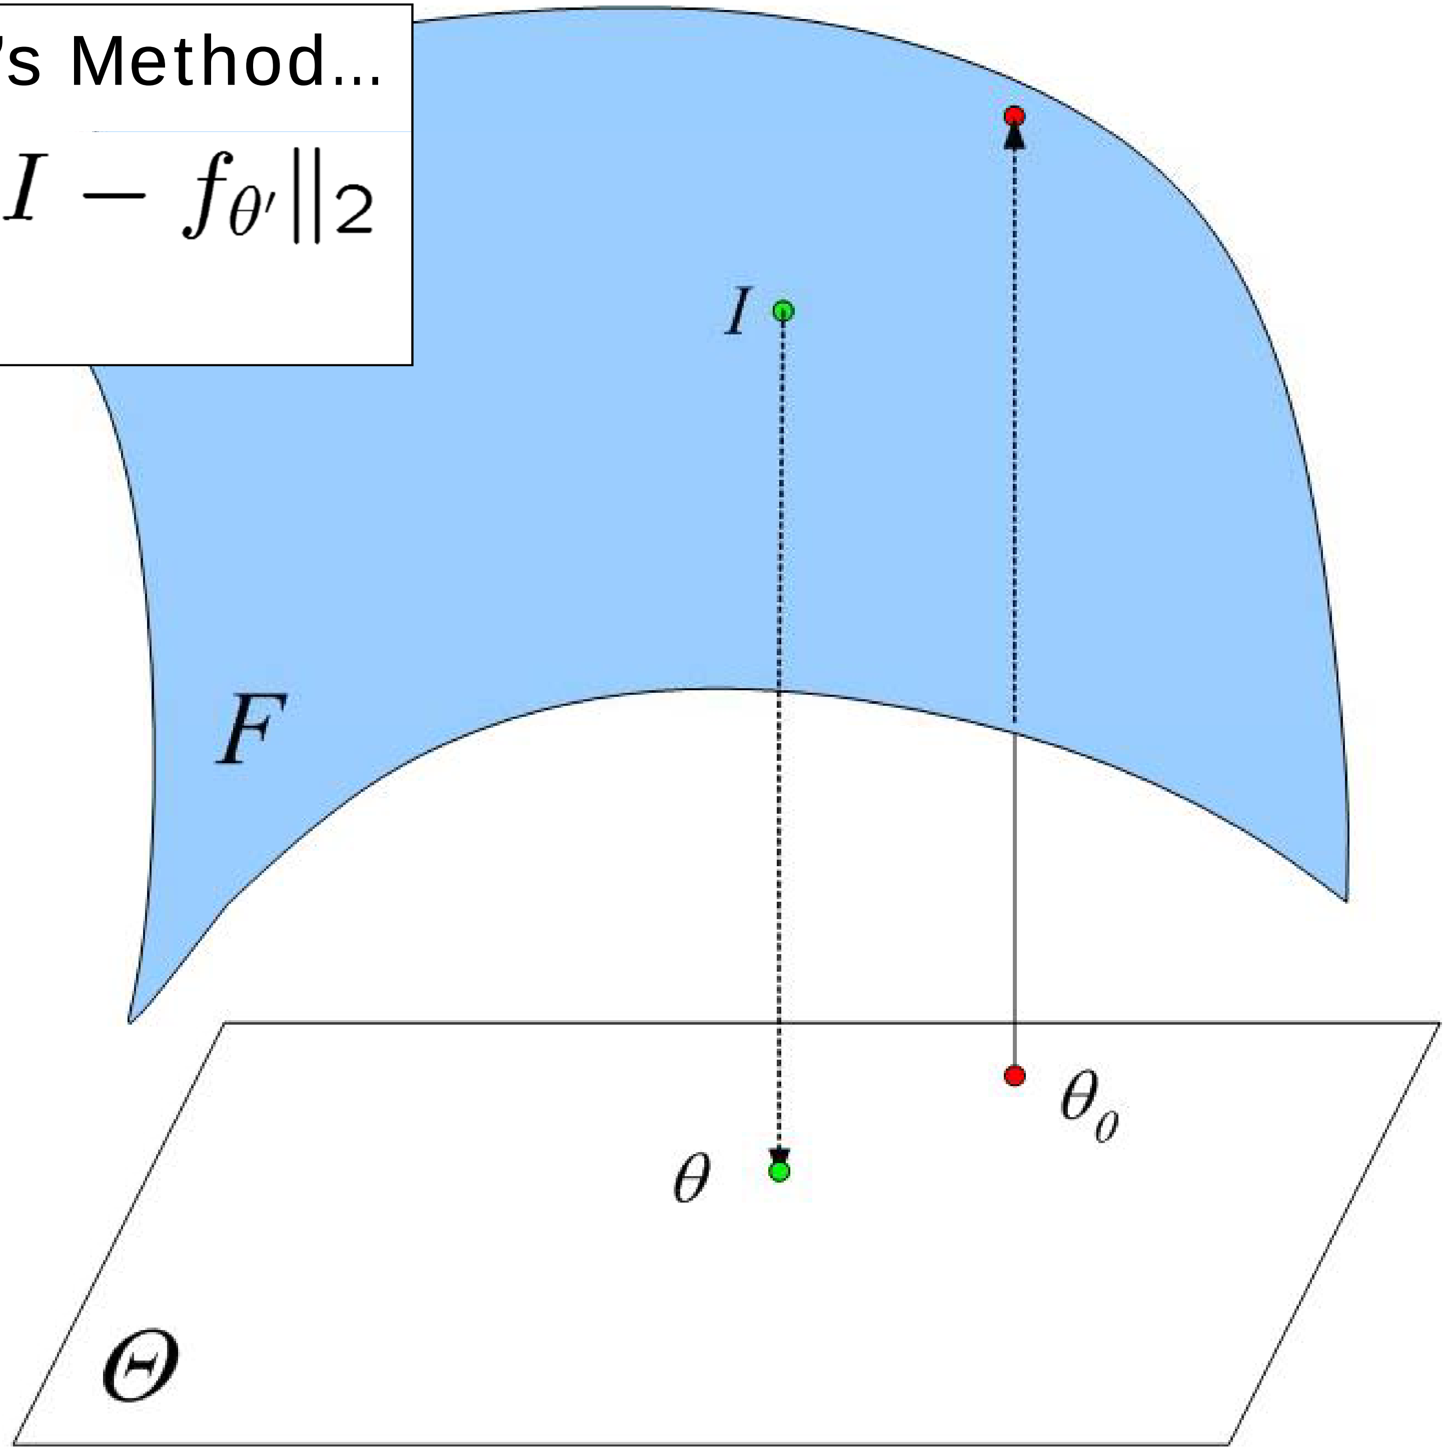

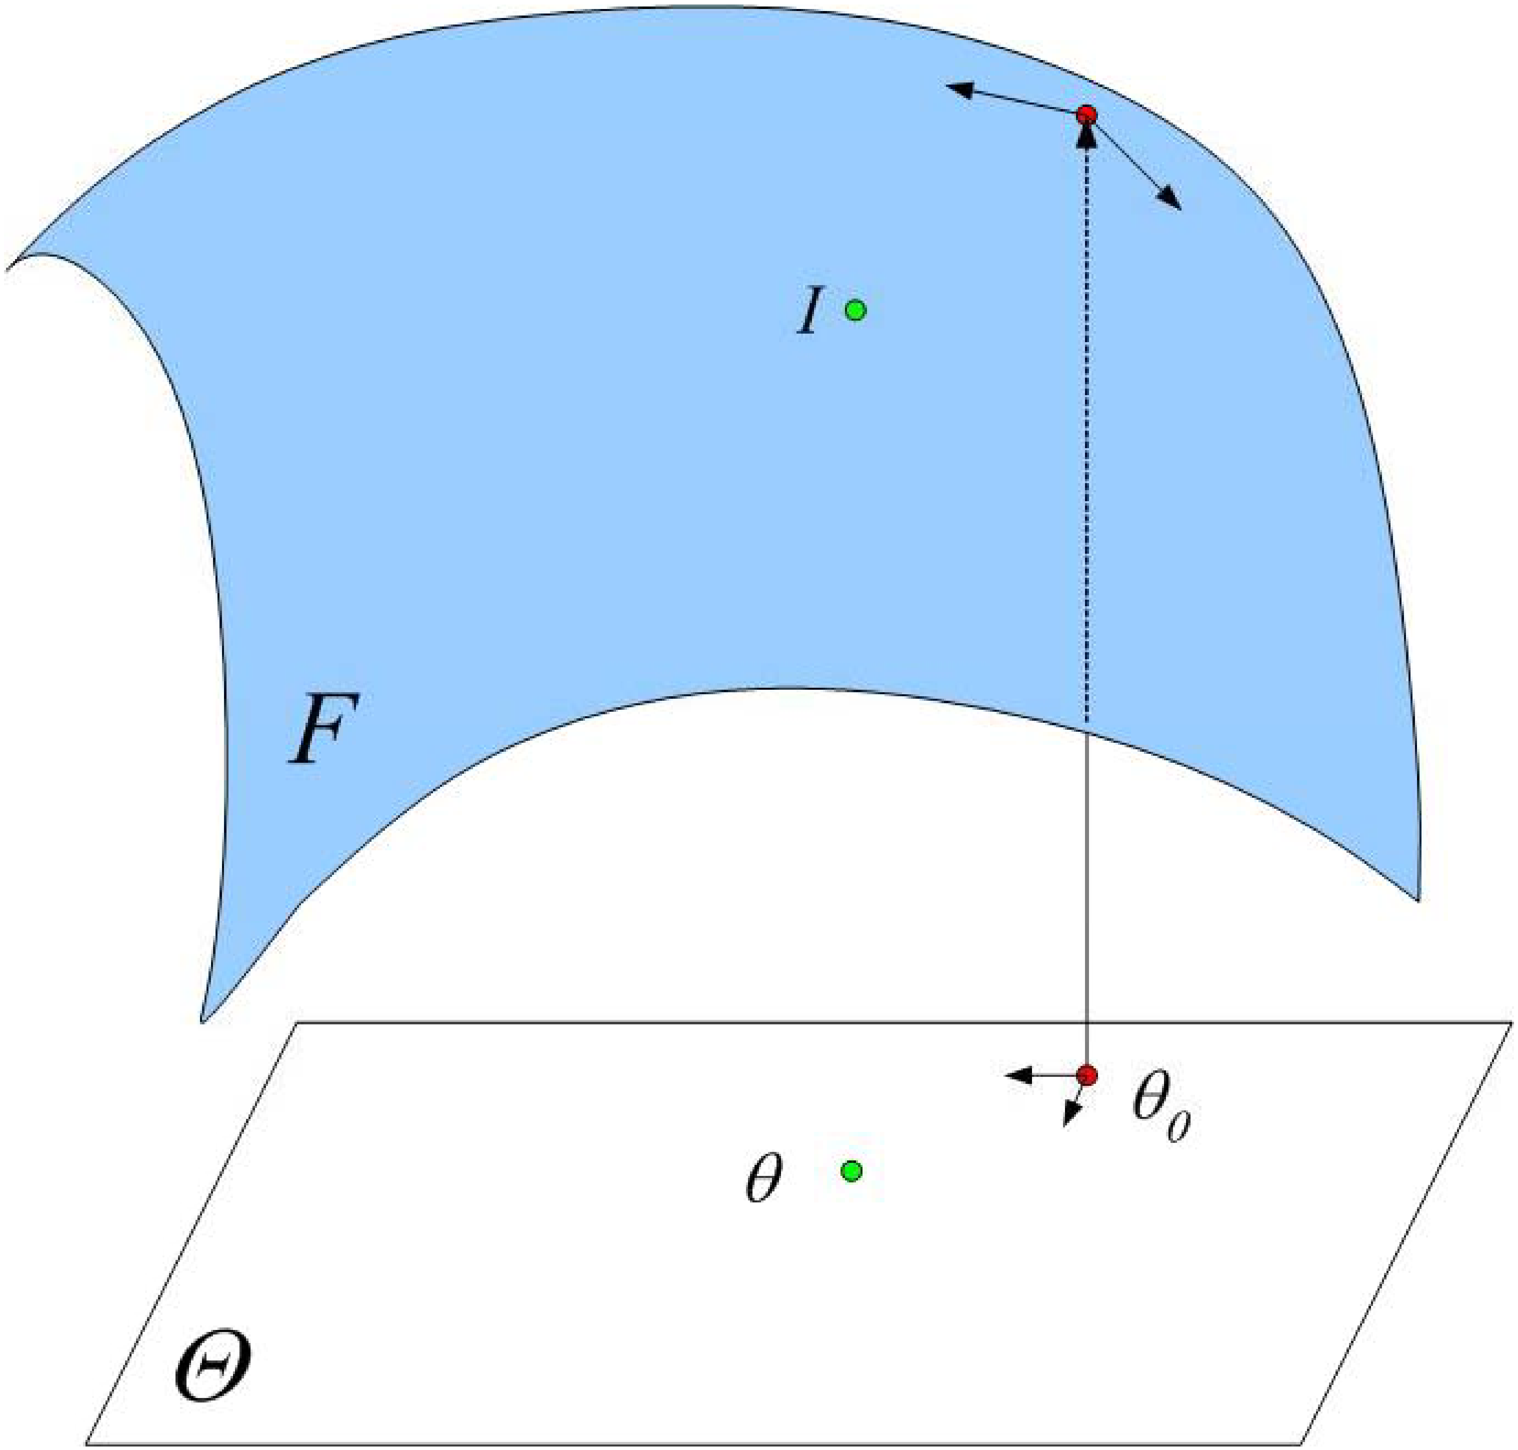
I
F
θ
θ0
Θ

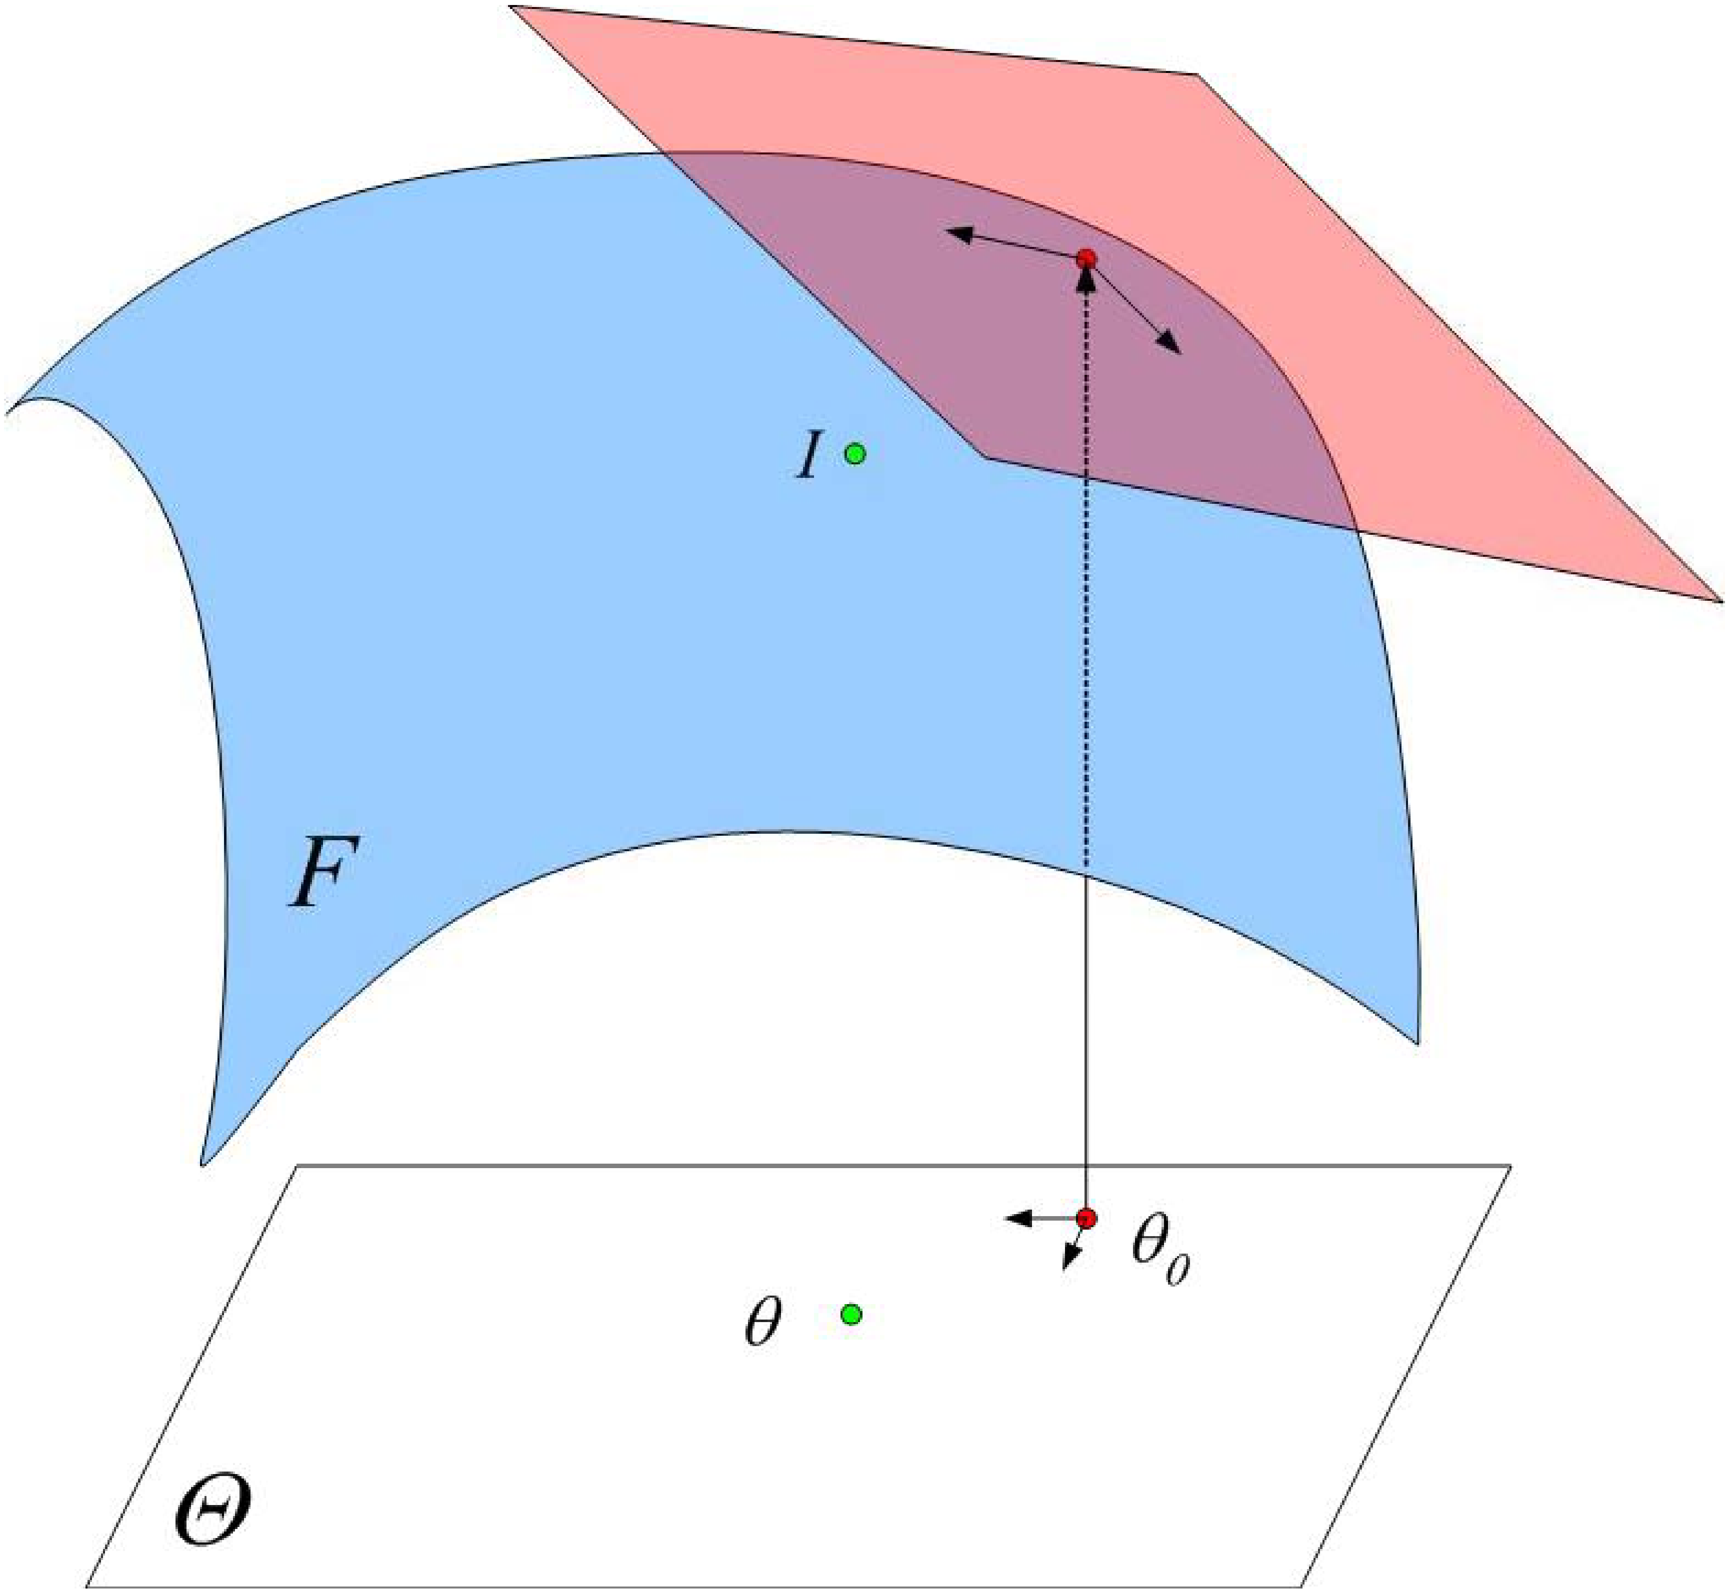
I
F
θ₀
θ
Θ

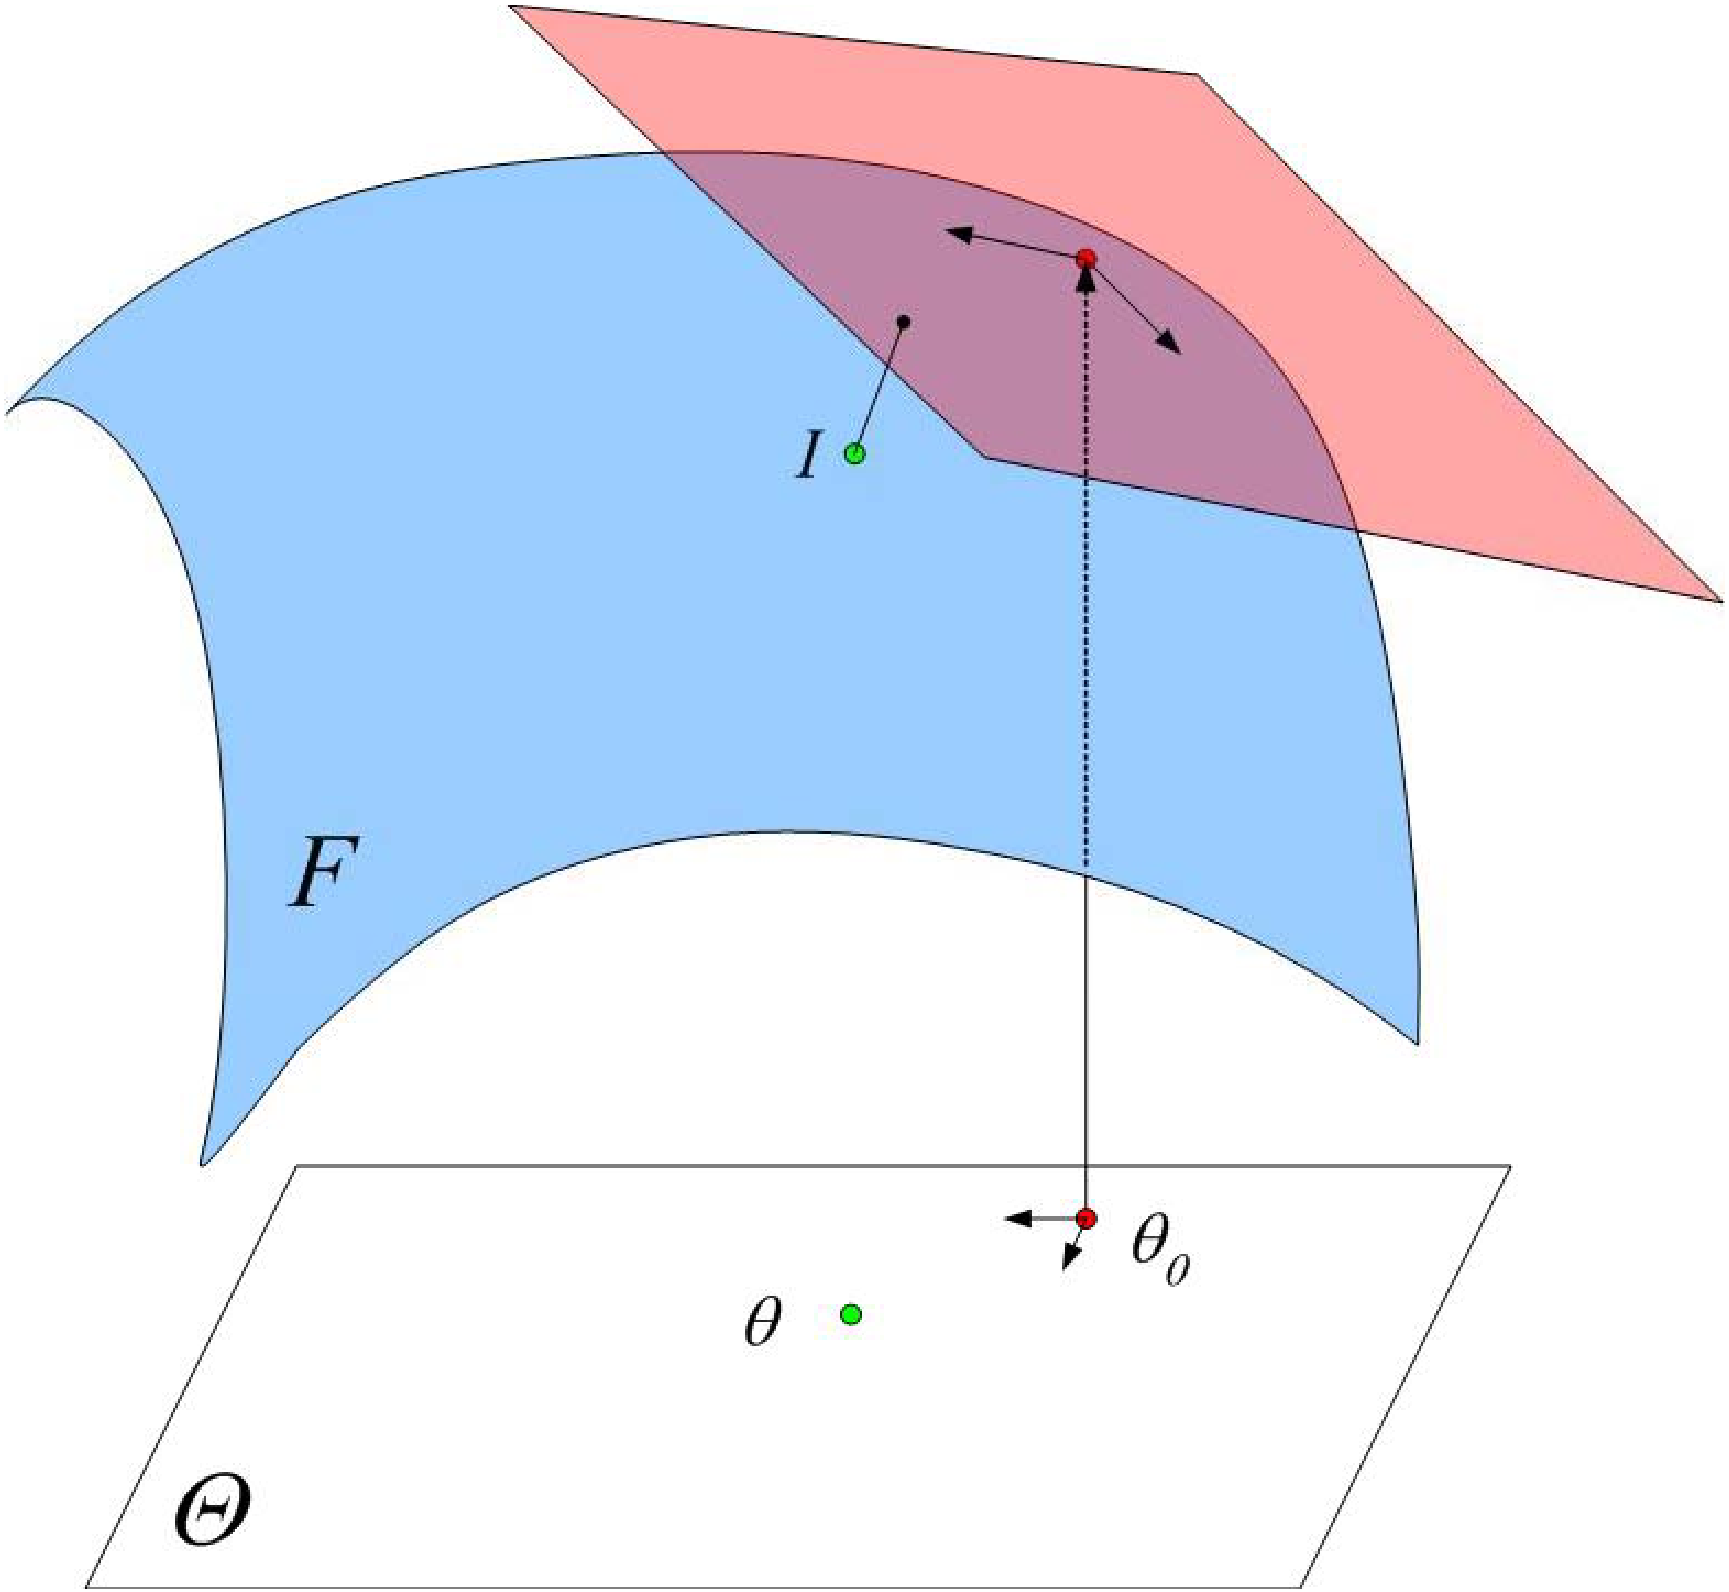
I
F
θ_0
θ
Θ

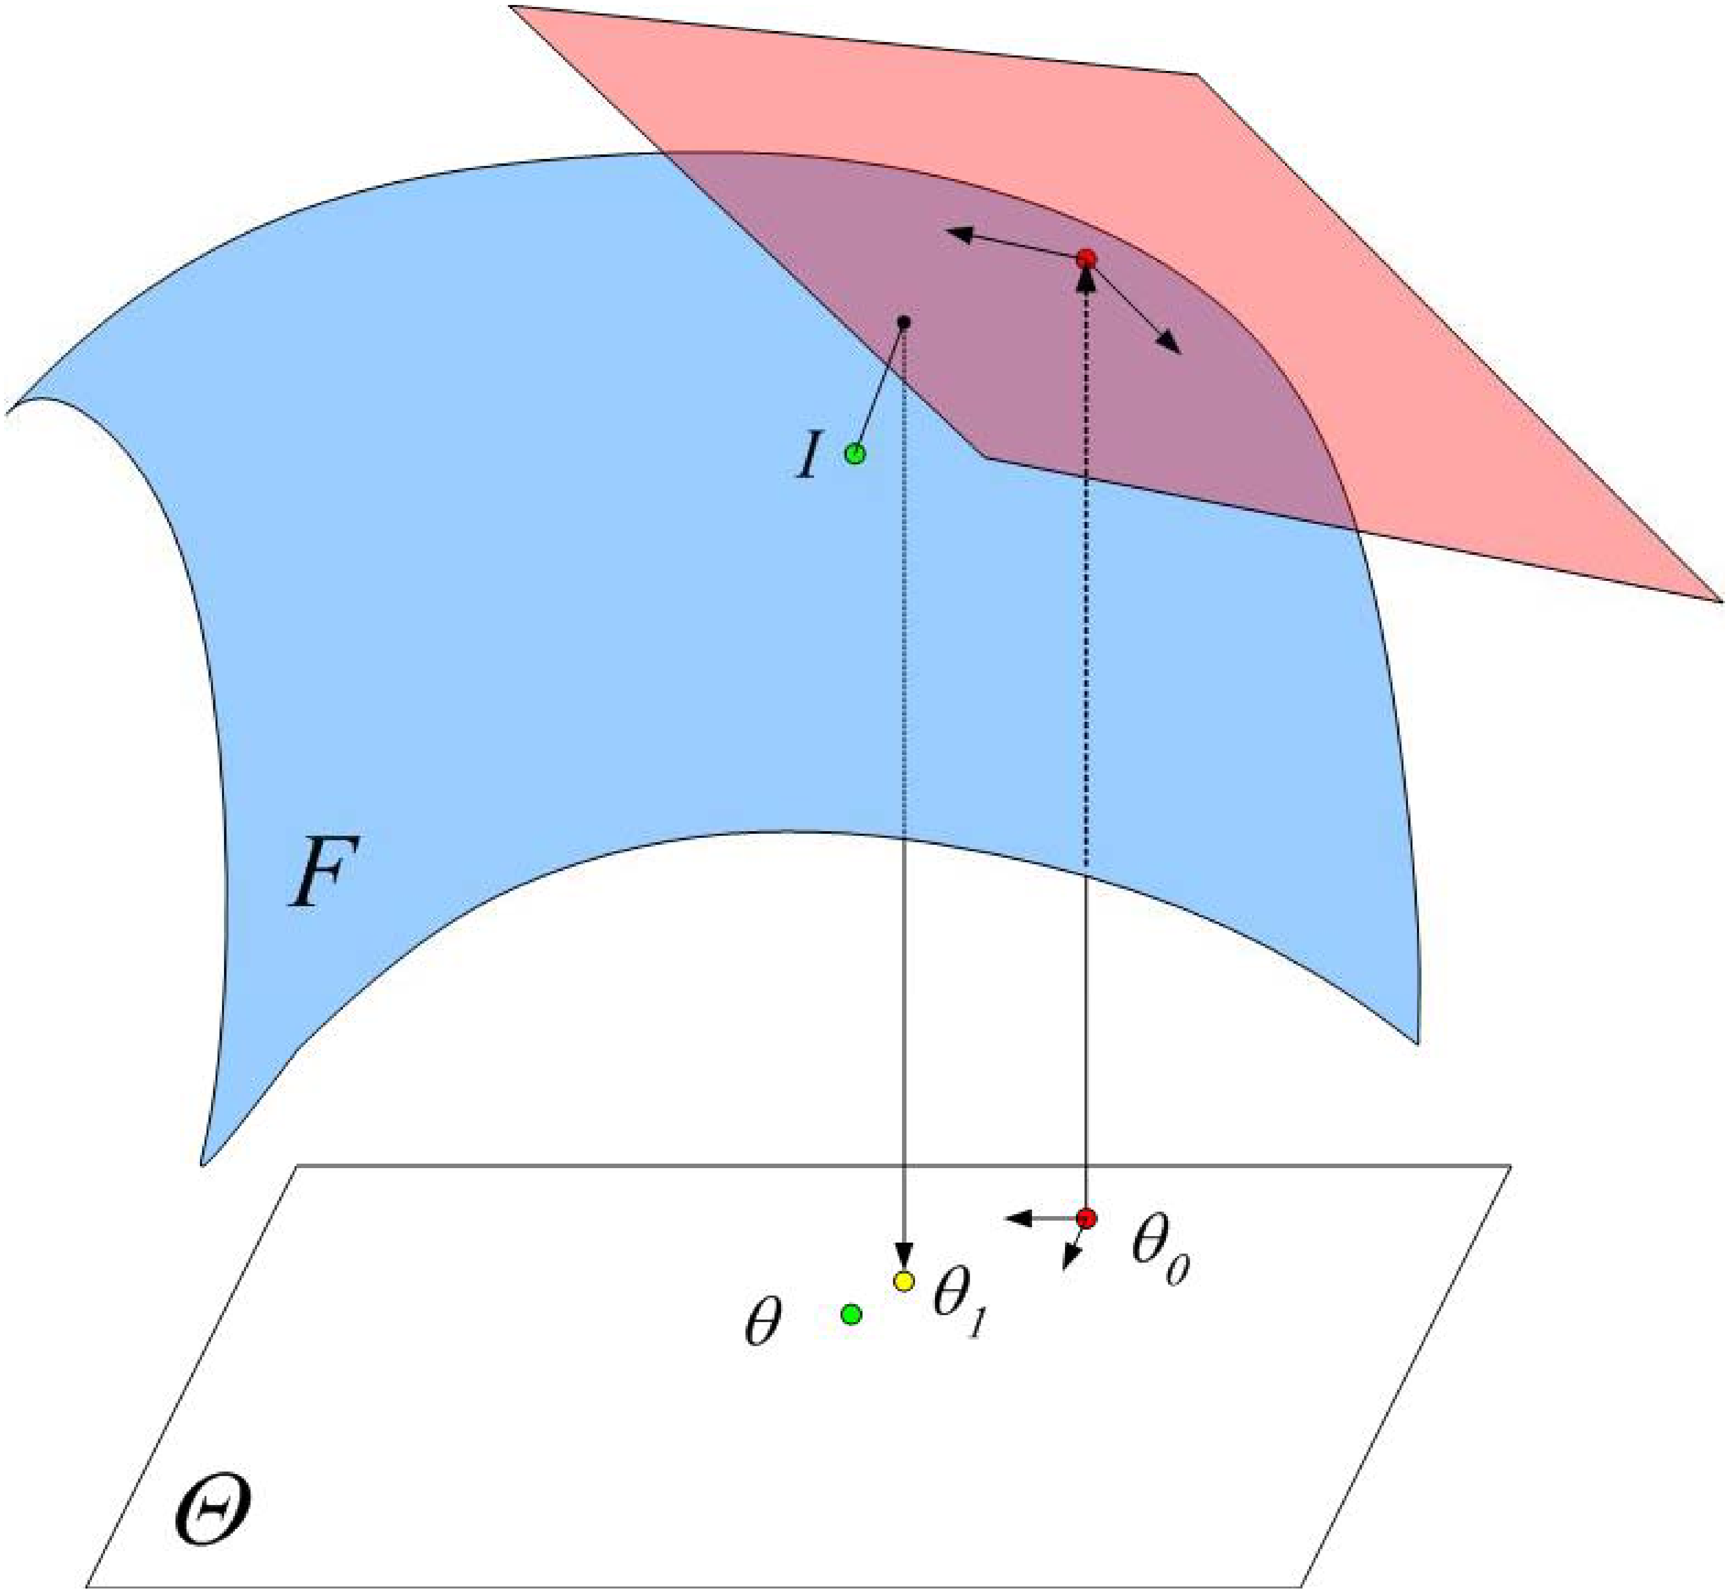

I
F
θ
θ1
θ0
Θ

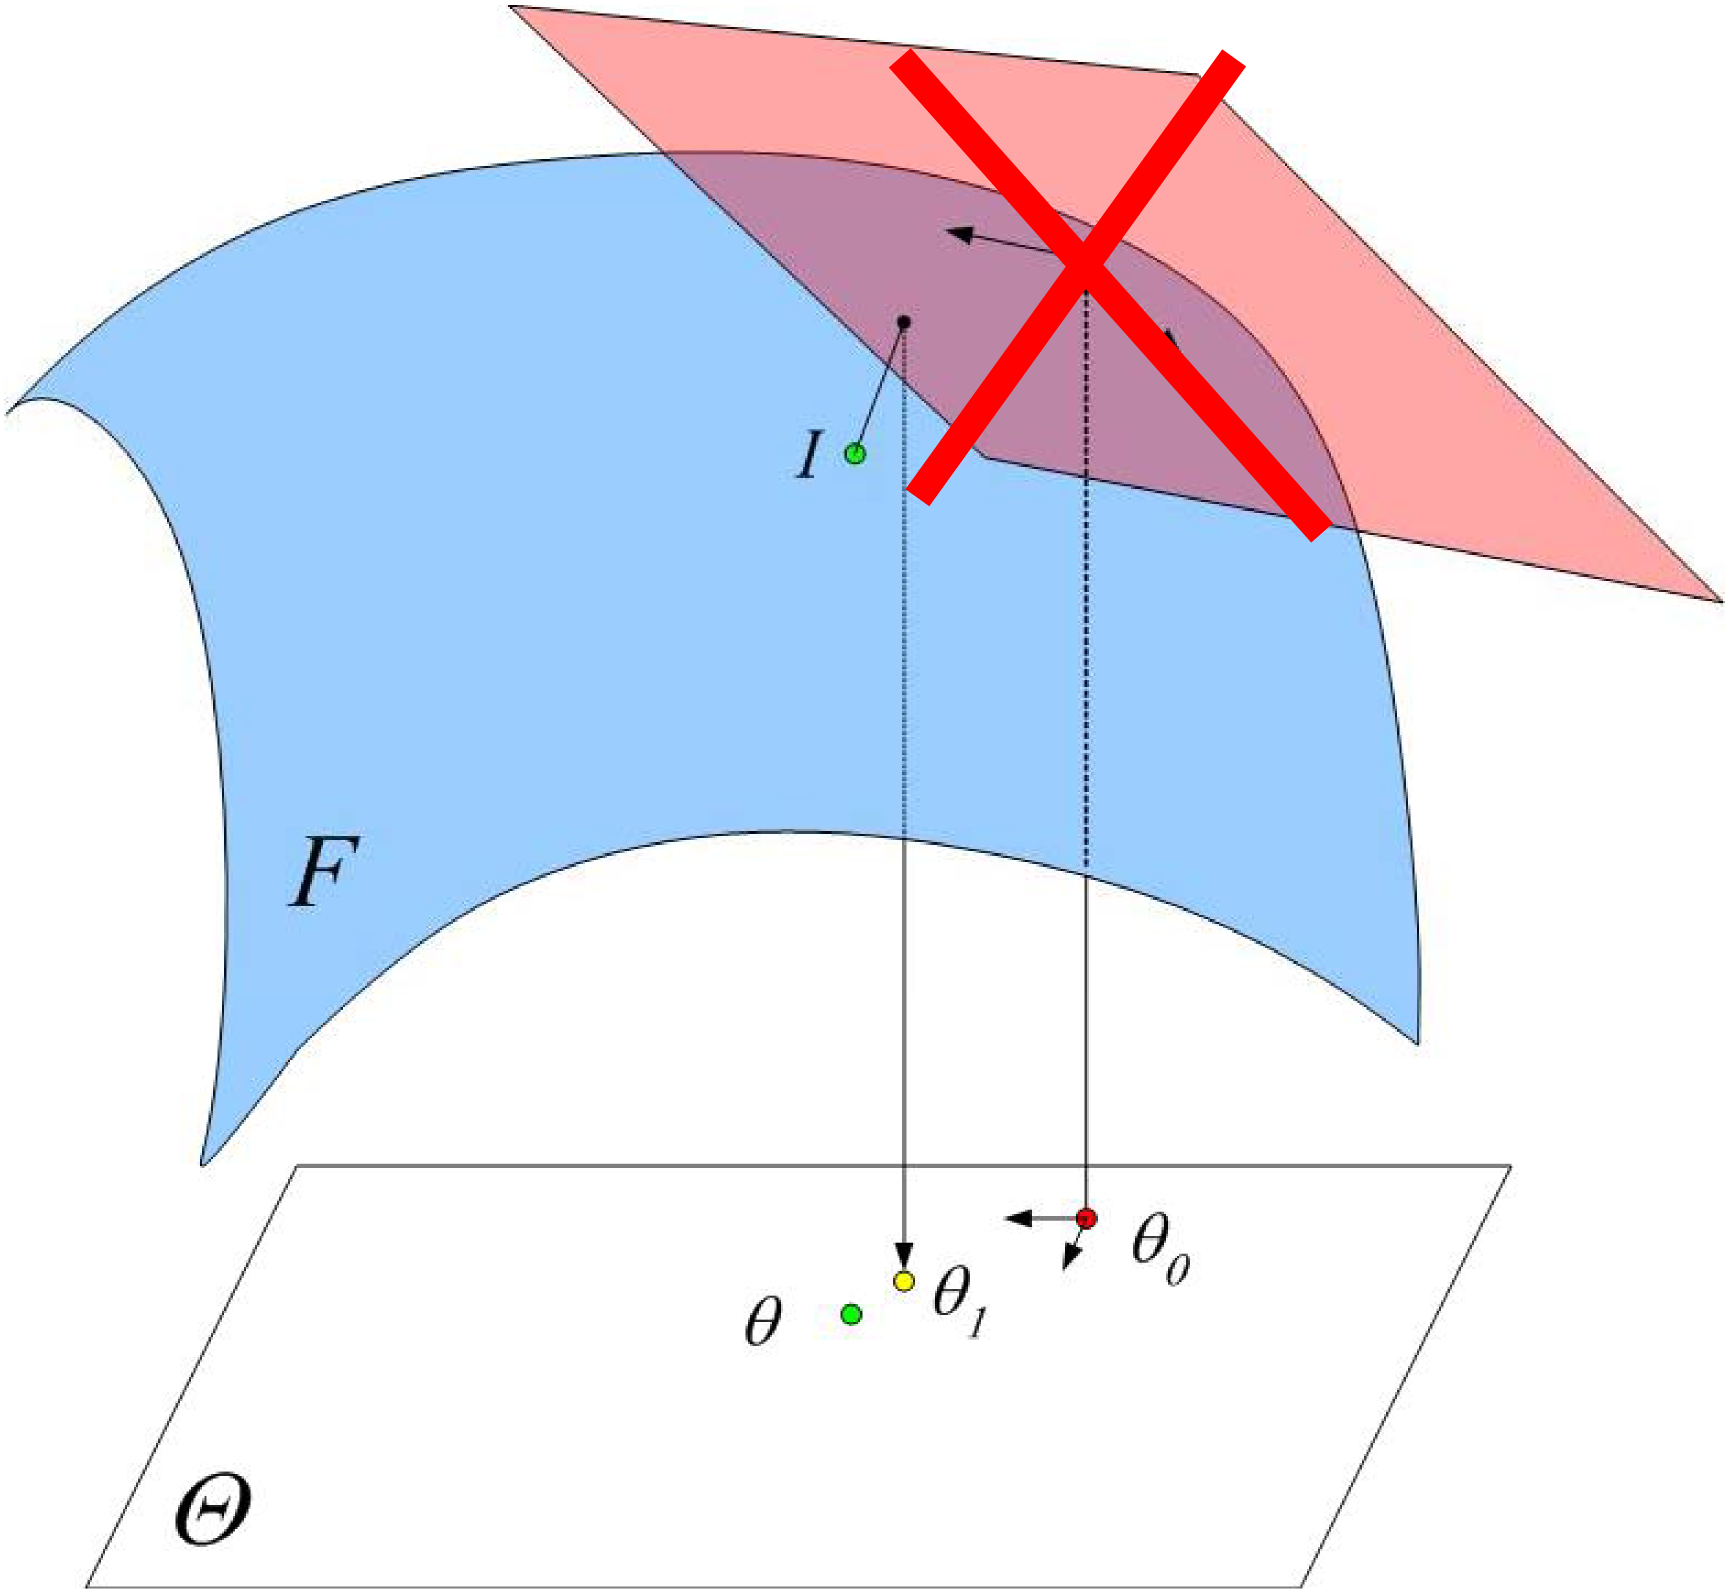
I
F
θ
θ_1
θ_0
Θ

Overview

- Motivating application: parameter estimation
- Non-differentiability from edge migration
- Parameter estimation (revisited)
- Non-differentiability from edge occlusion
- Manifolds in Compressive Sensing

Non-differentiability from Edge Migration

- Problem: movement of *sharp edges*
  - example: shifted disk [Donoho,Grimes]

$$\|f_{\theta+h} - f_\theta\|_2 \sim \|h\|_2^{1/2}, \qquad h \to 0$$

$$\left\| \frac{f_{\theta+h} - f_\theta}{h} \right\|_2 \sim \frac{1}{\|h\|_2^{1/2}} \to \infty$$

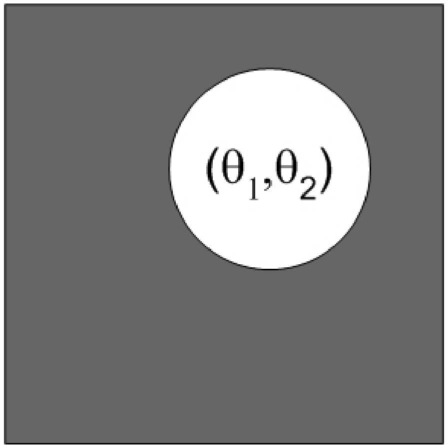

- Tangents *do not exist*

- Visualization: Local PCA experiment

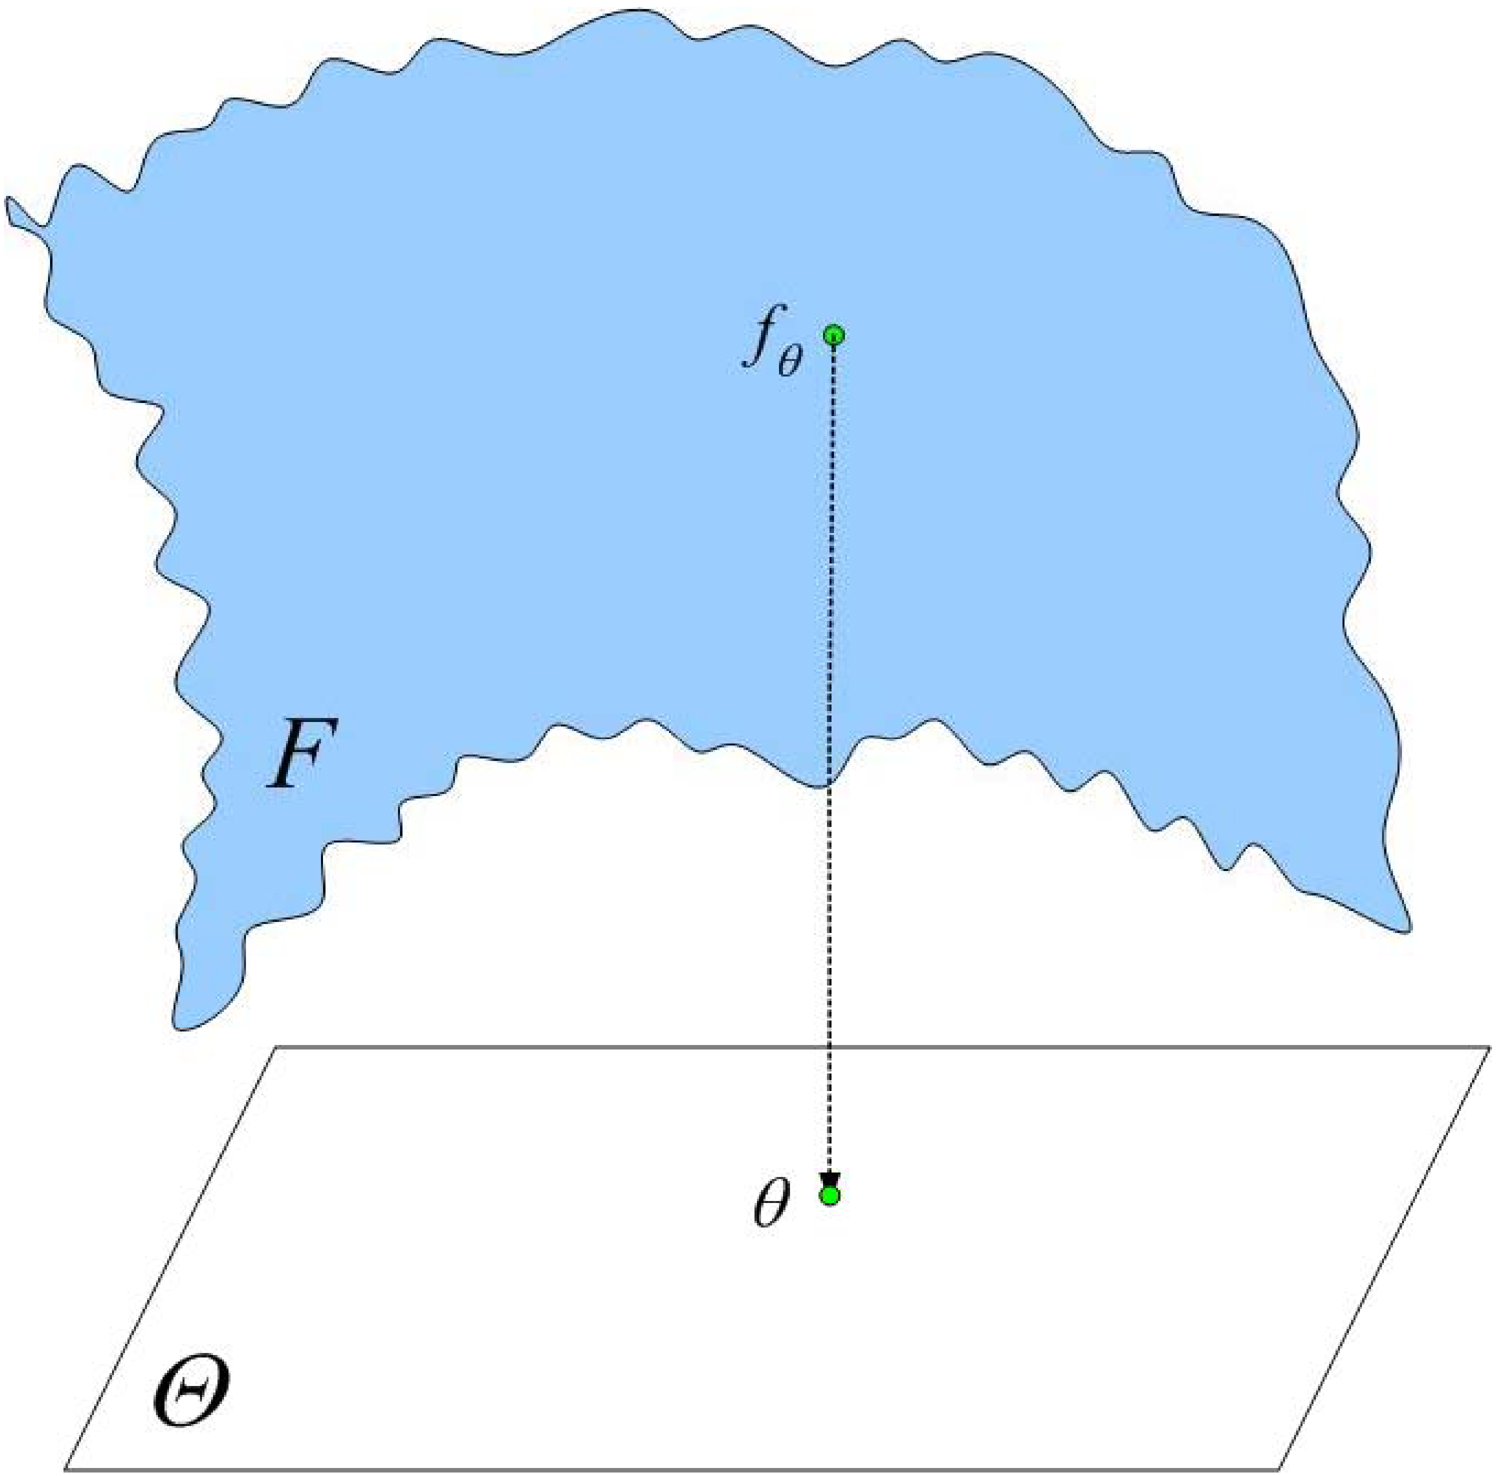
f_θ
F
θ
Θ

“Local PCA” to approximate tangent space

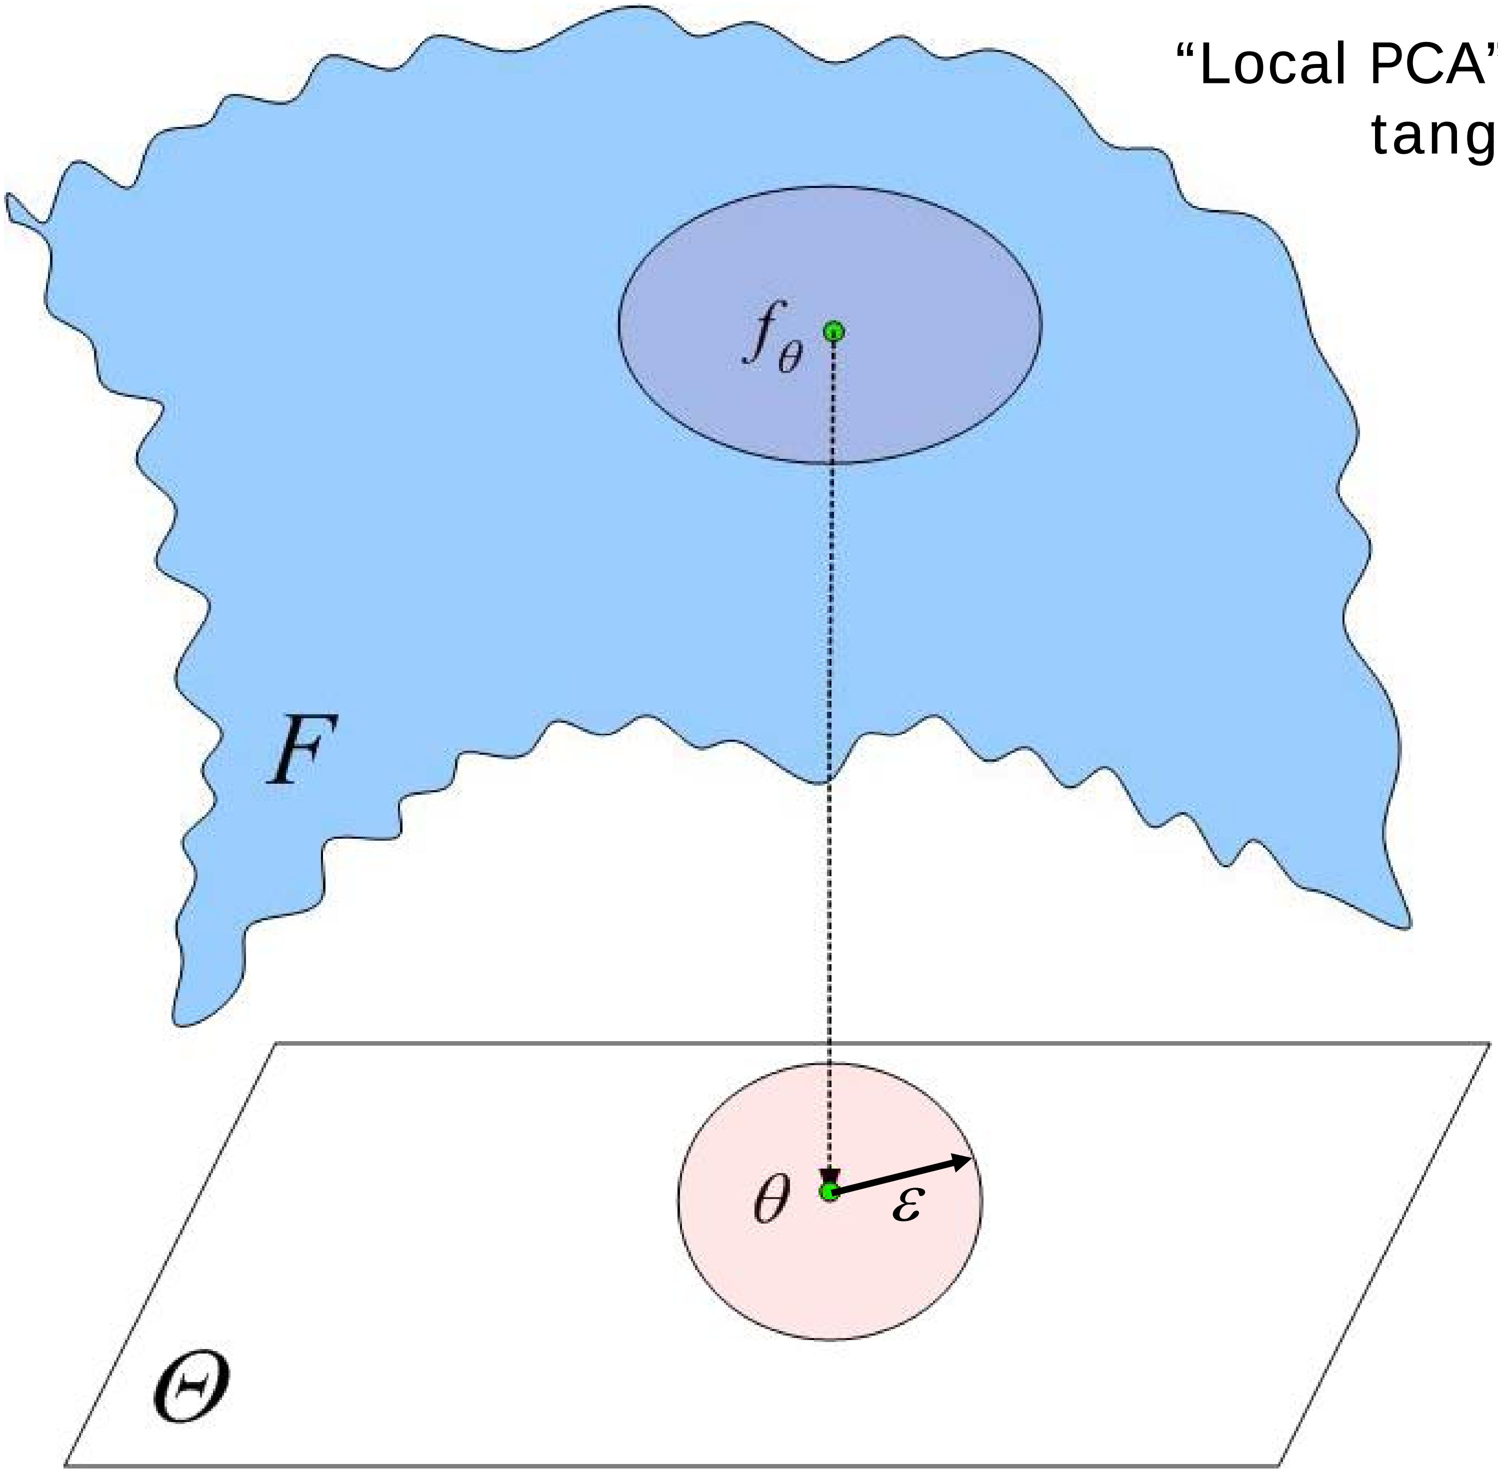

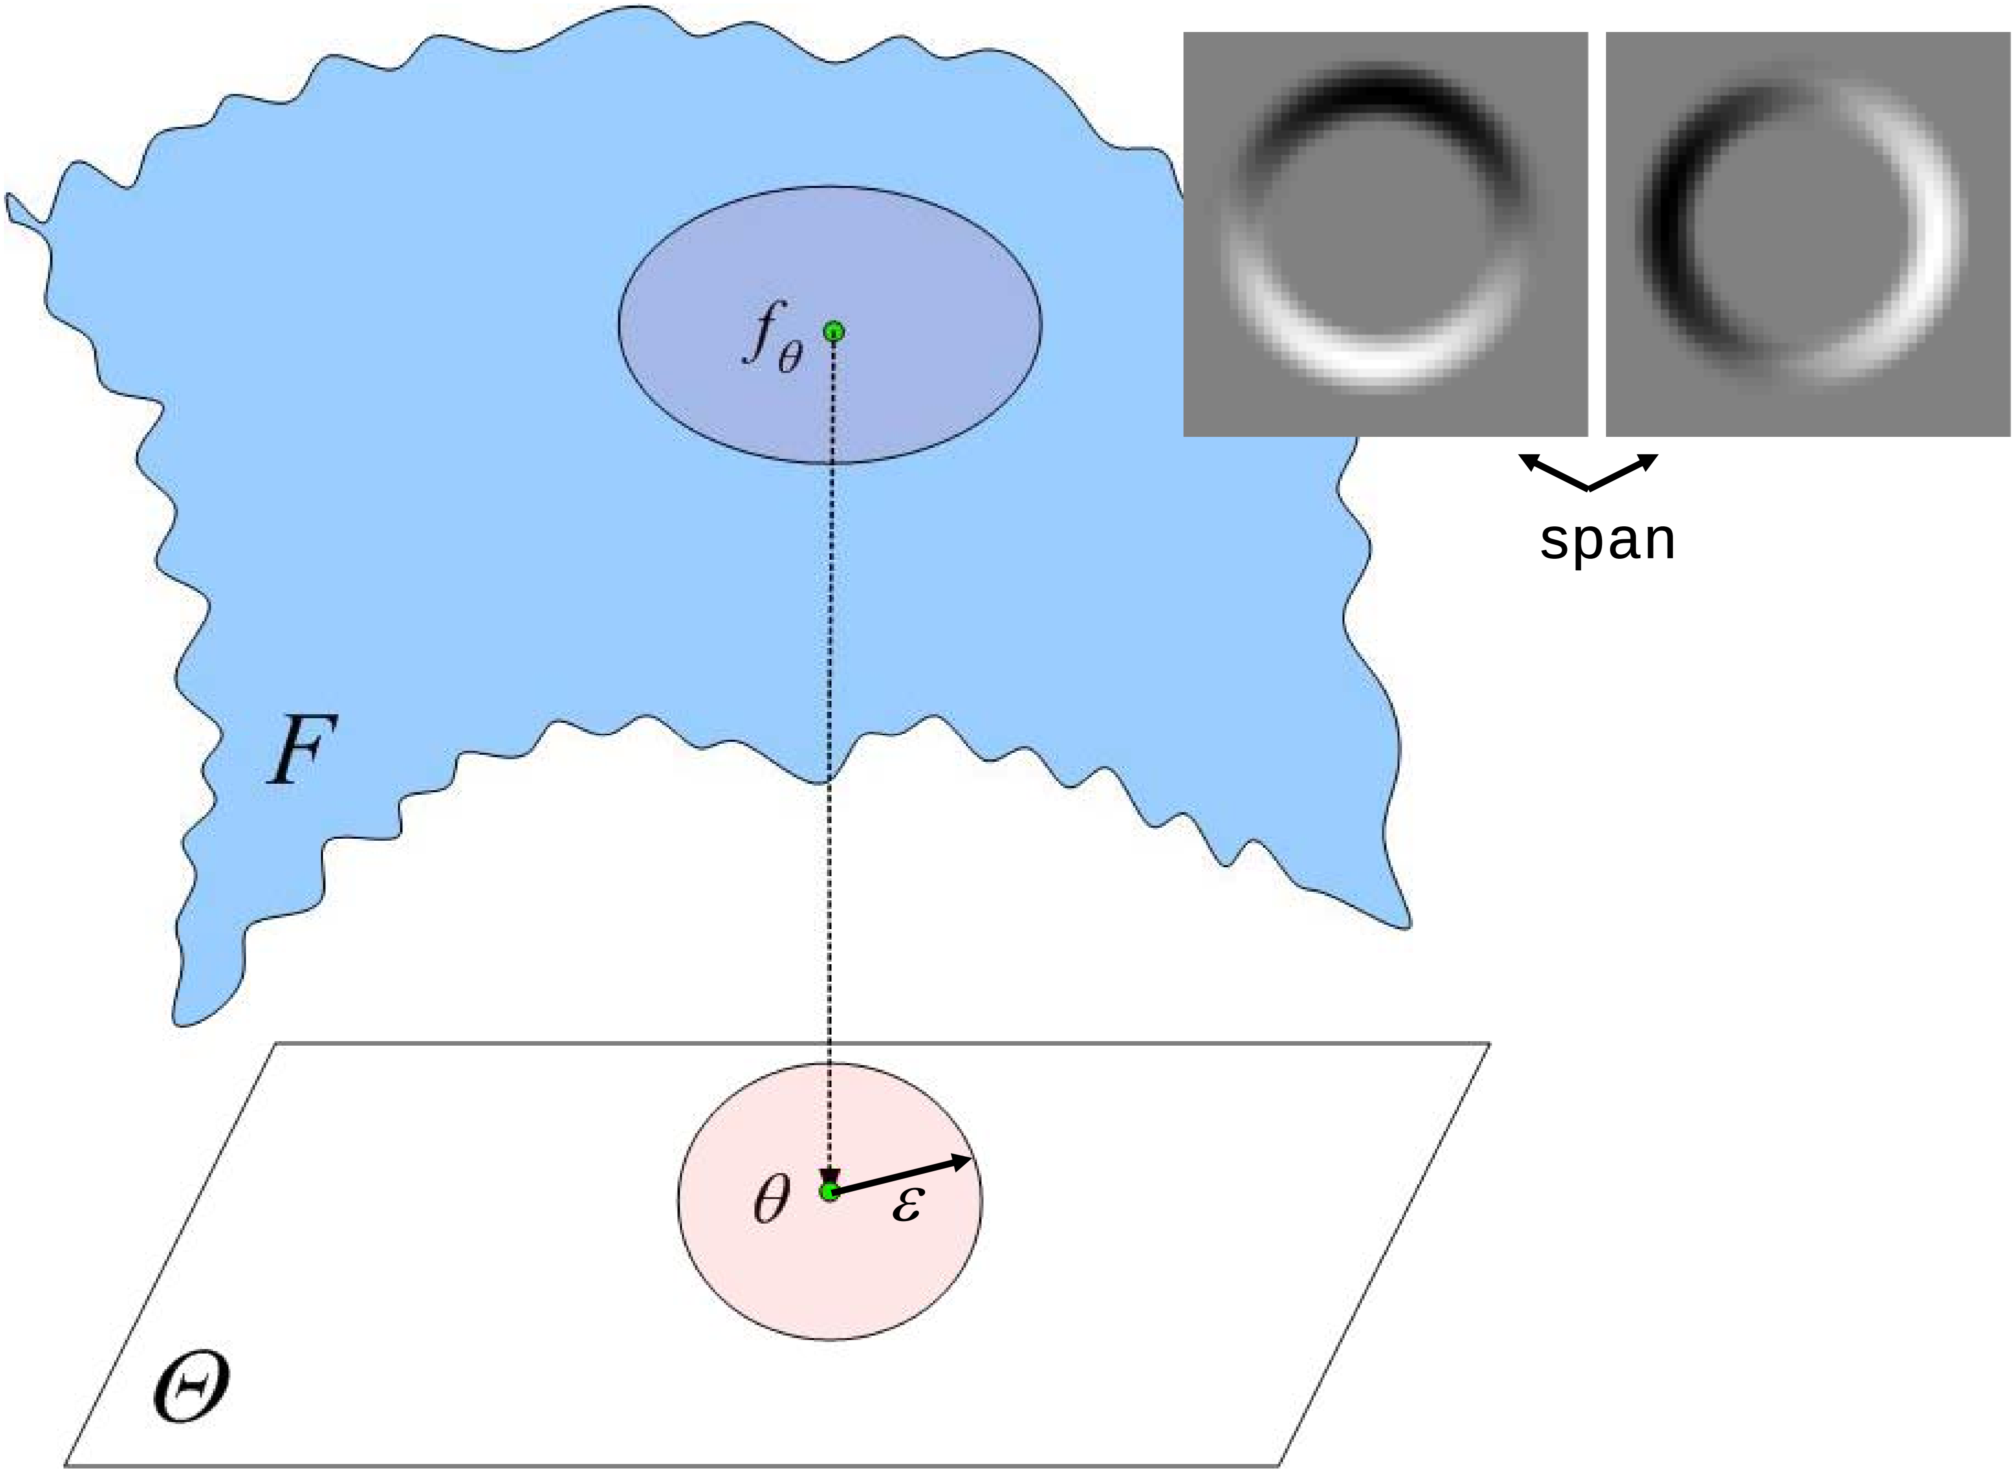

f_θ
span
F
θ
ε
Θ

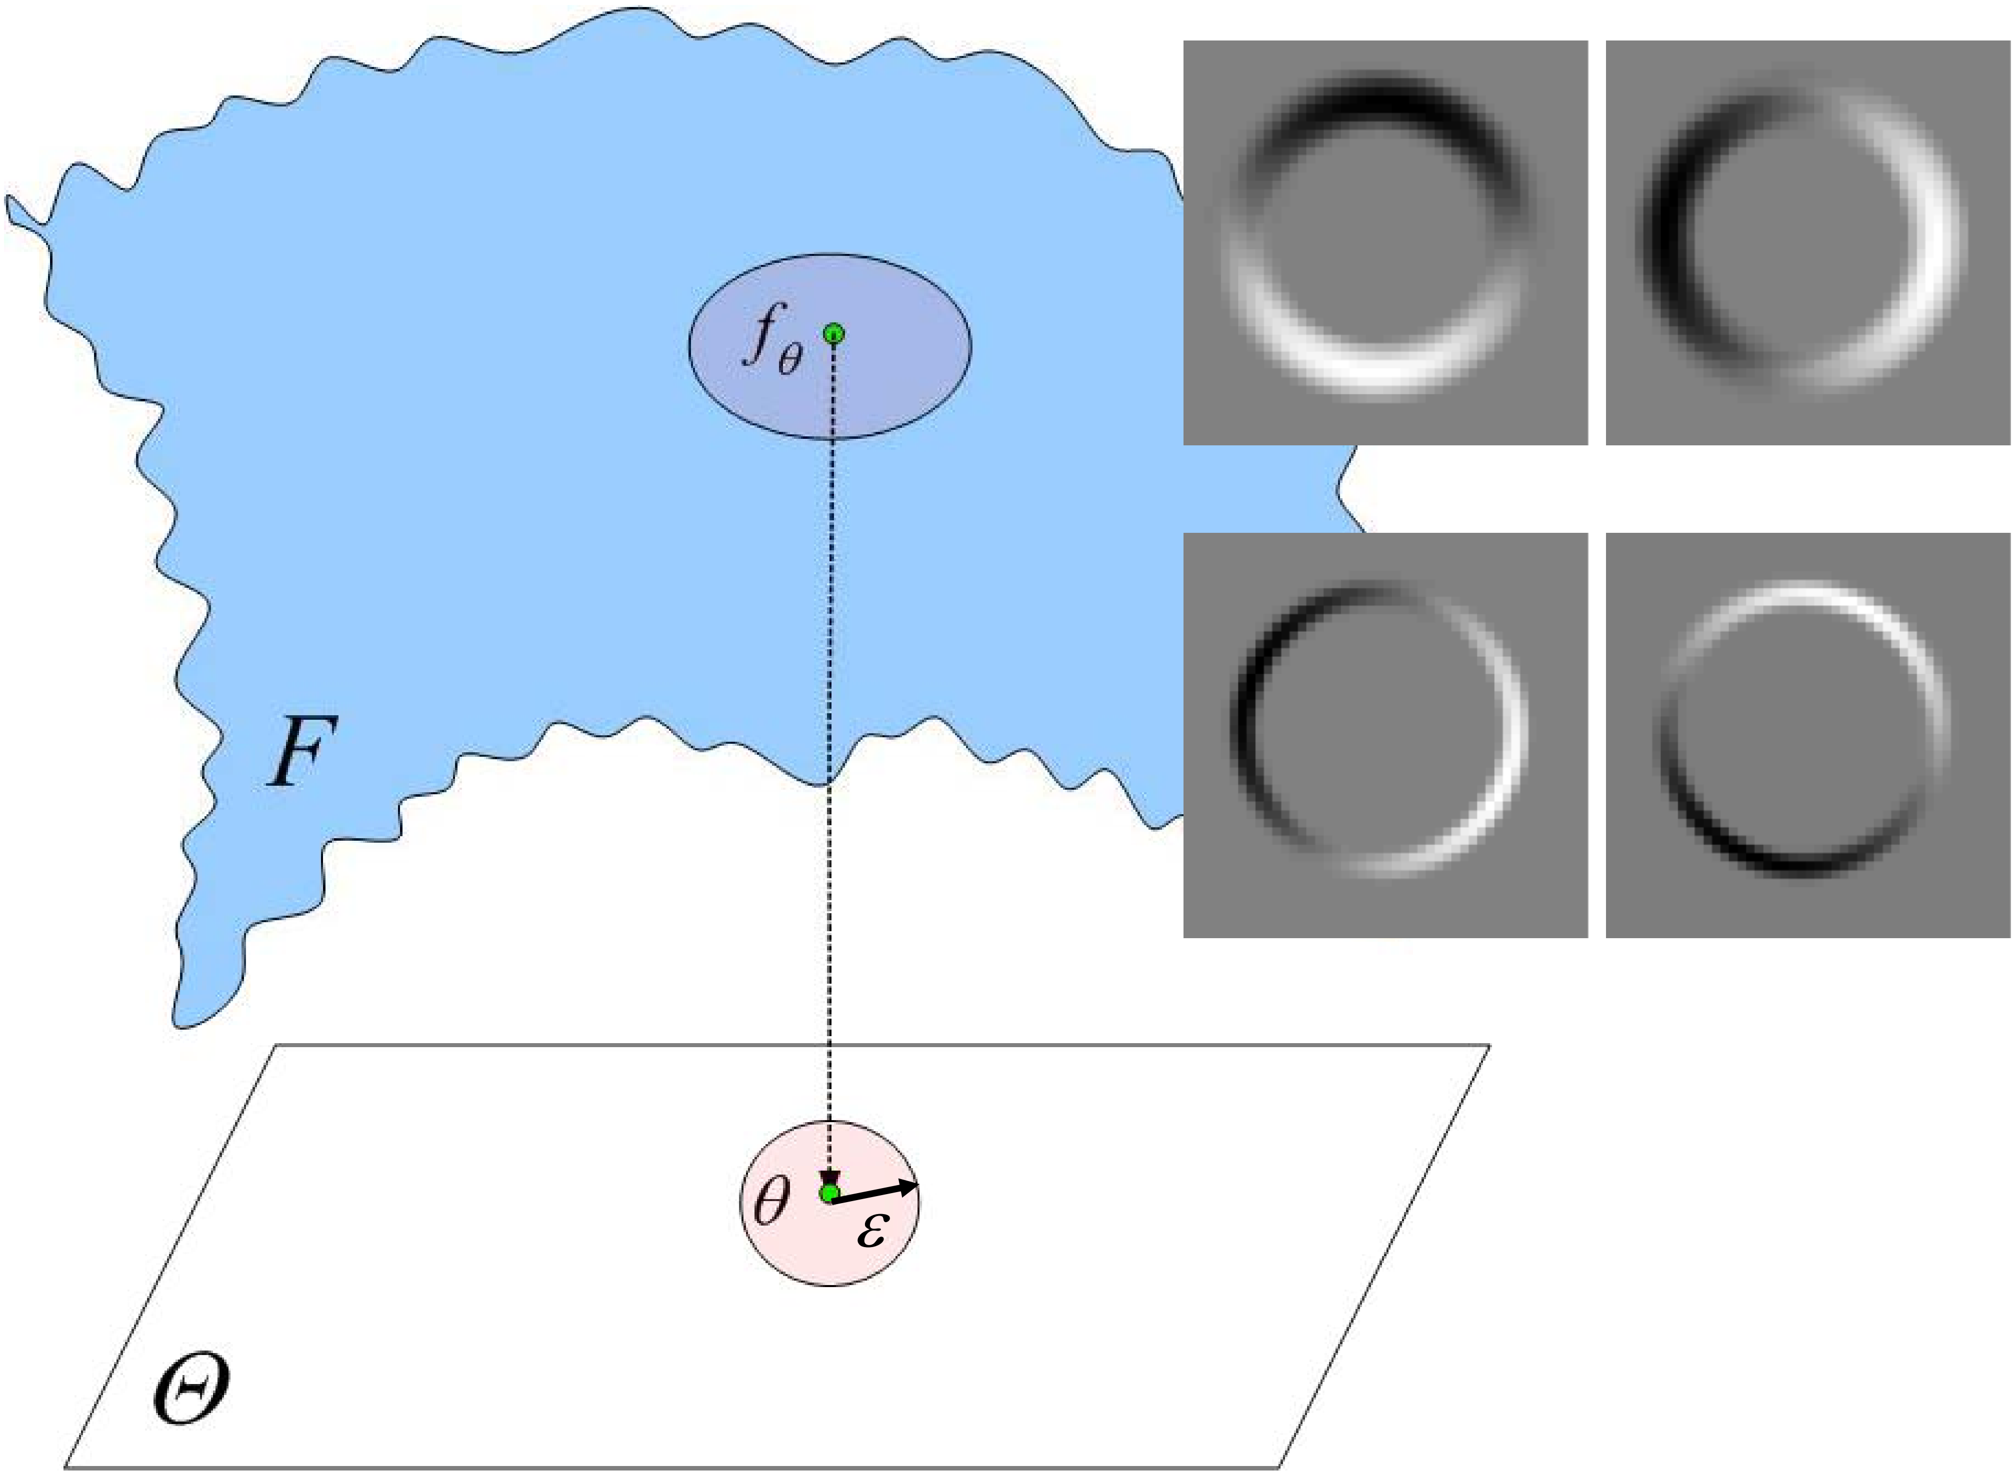

f_θ
F
θ
ε
Θ

Multiscale Tangent Structure

- Family of *approximate tangent planes*
  - $T(\varepsilon, \theta)$ scale, location on manifold
- *If manifold F were differentiable:*

$$\lim_{\epsilon \to 0} T(\epsilon, \theta) = T_\theta(F)$$

- Does not happen when edges exist:

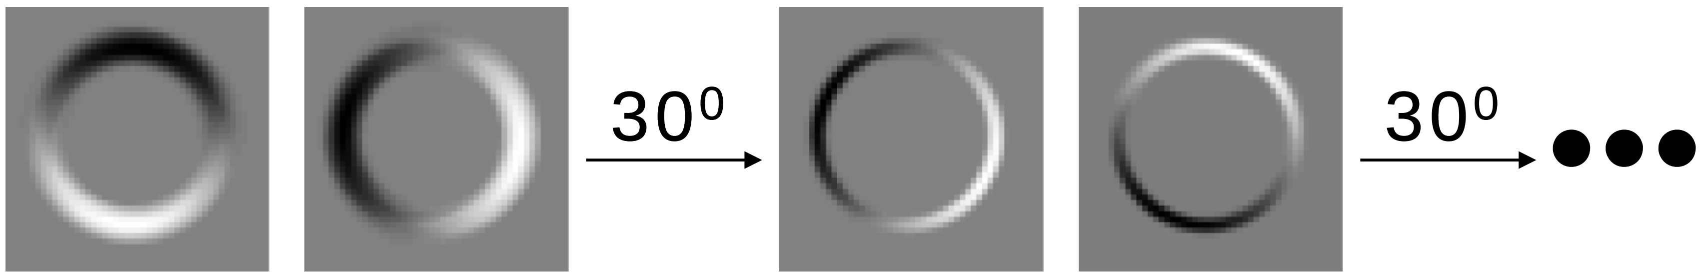

- Tangent spaces do not converge
  - twisting into new dimensions
- But we can study and *exploit* this multiscale structure
  ~ wavelets for non-differentiable functions

Shortcut to Multiscale Structure via Regularization

- Smoothing the *images* smoothes the *manifold*
  - more smoothing gives smoother manifold
- Example: convolution with Gaussian, width s

$$f_{\theta,s} = \phi_s * f_\theta \qquad F_s = \{f_{\theta,s} : \theta \in \Theta\}$$

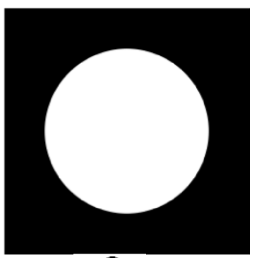

f_θ

$f_{\theta,s}$

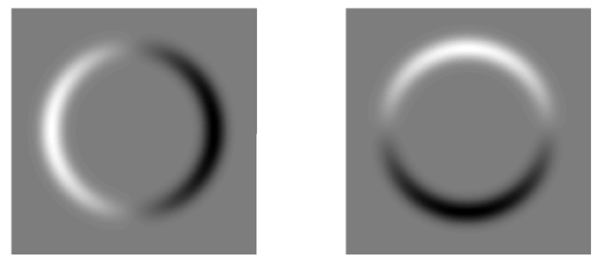

$$T(s,\theta) := T_\theta(F_s)$$

- Alternate family of multiscale tangent planes
  - tangent planes well defined, analogous to PCA

Wavelet-like Characterization

- Family $T(s, \theta)$ like continuous wavelet transform
  - discretization: $T(s_i, \theta_j) \quad (i, j)$

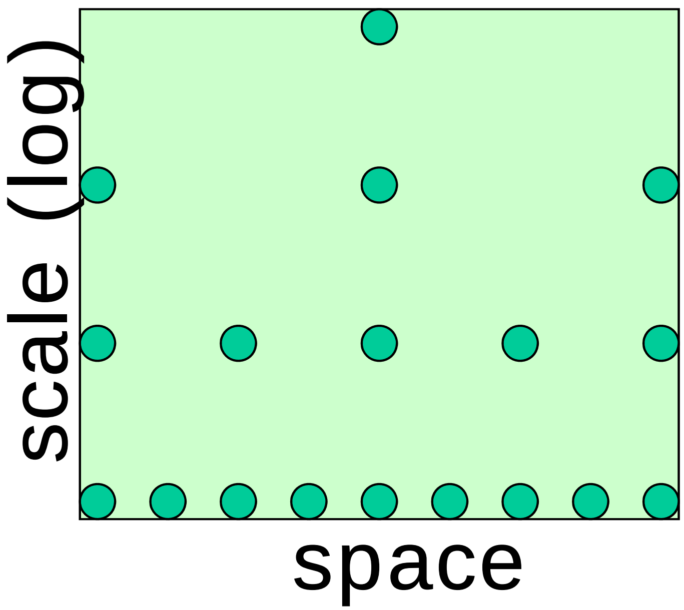

Fixed angle of twist between samples

- Sampling is manifold-dependent

Overview

- Motivating application: parameter estimation
- Non-differentiability from edge migration
- Parameter estimation (revisited)
- Non-differentiability from edge occlusion
- Manifolds in Compressive Sensing

Recall Newton's Method...

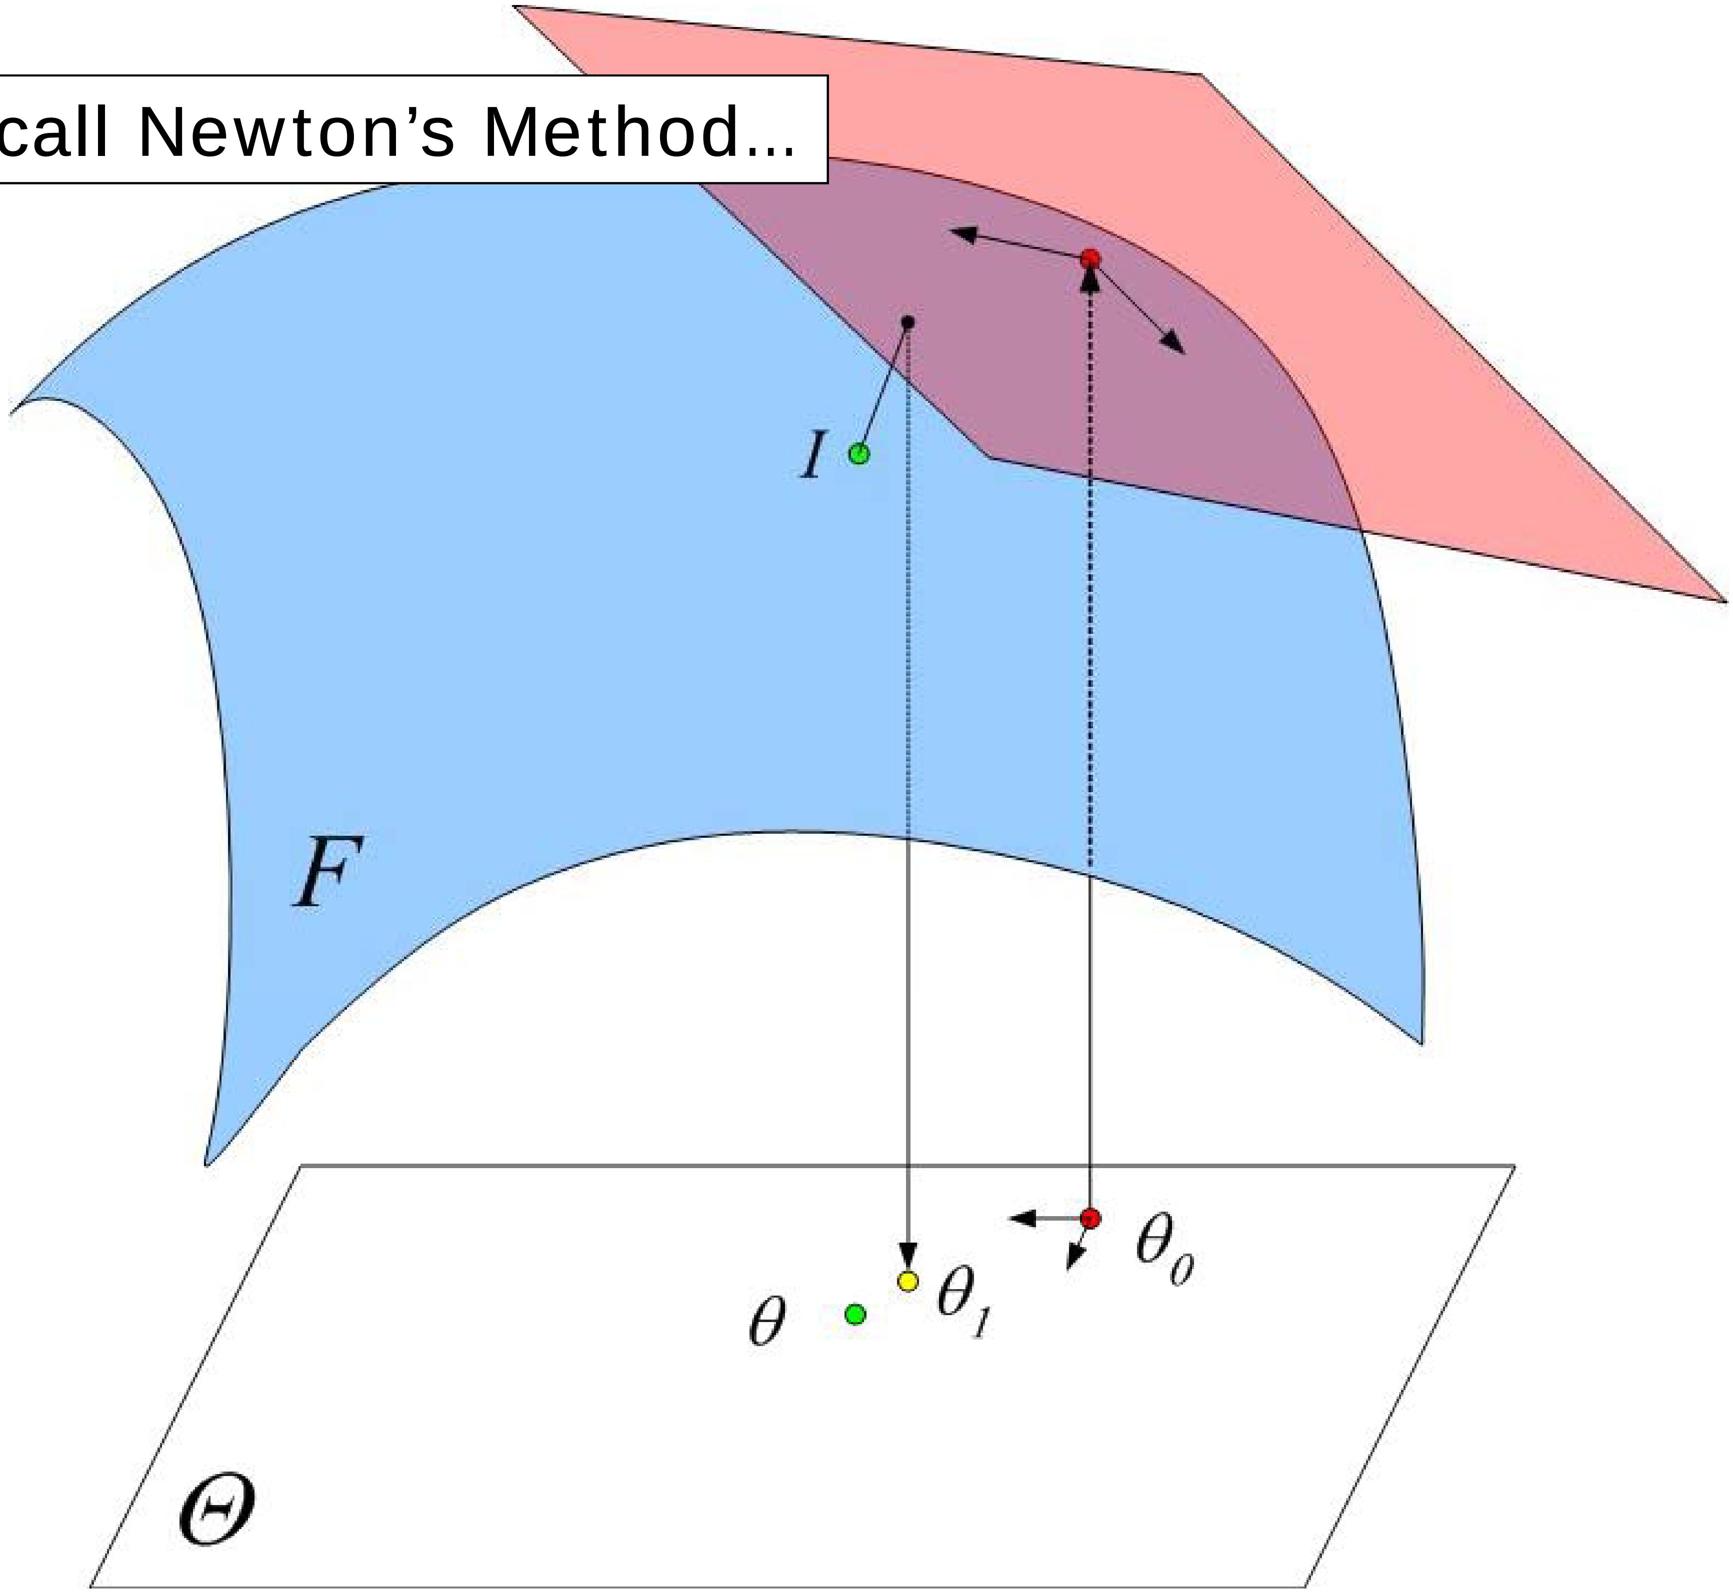

Recall Newton's Method...

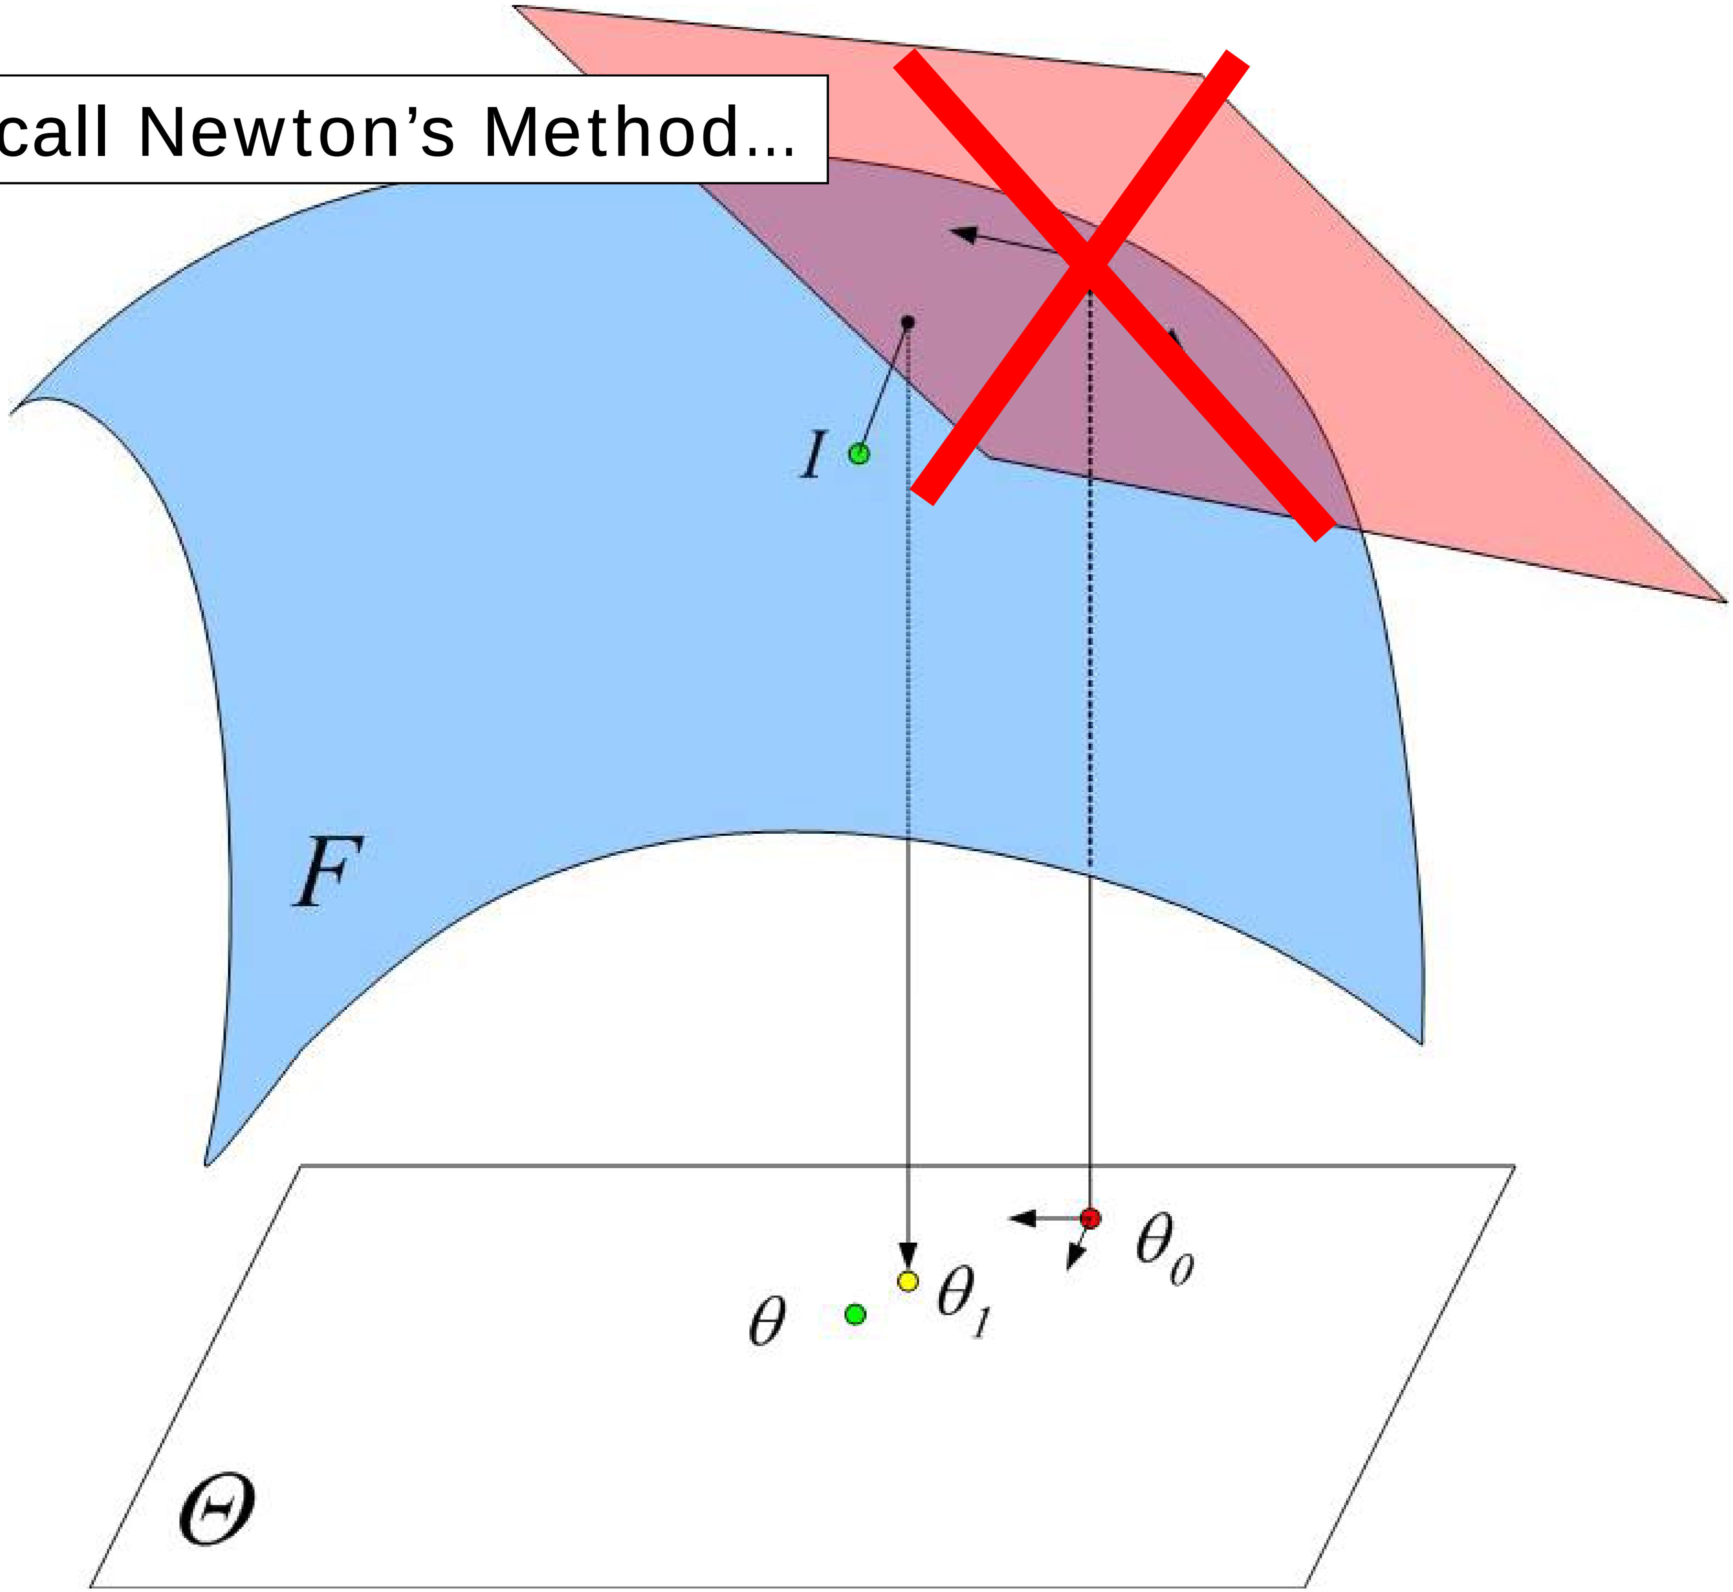

Multiscale Newton Algorithm

- Construct a coarse-to-fine *sequence* $\{F_s\}$ of manifolds that converge to F

$$\phi_s * f_\theta \to f_\theta, \qquad s \to 0$$

$$F_s \to F, \qquad s \to 0$$

- Take one Newton step at each scale

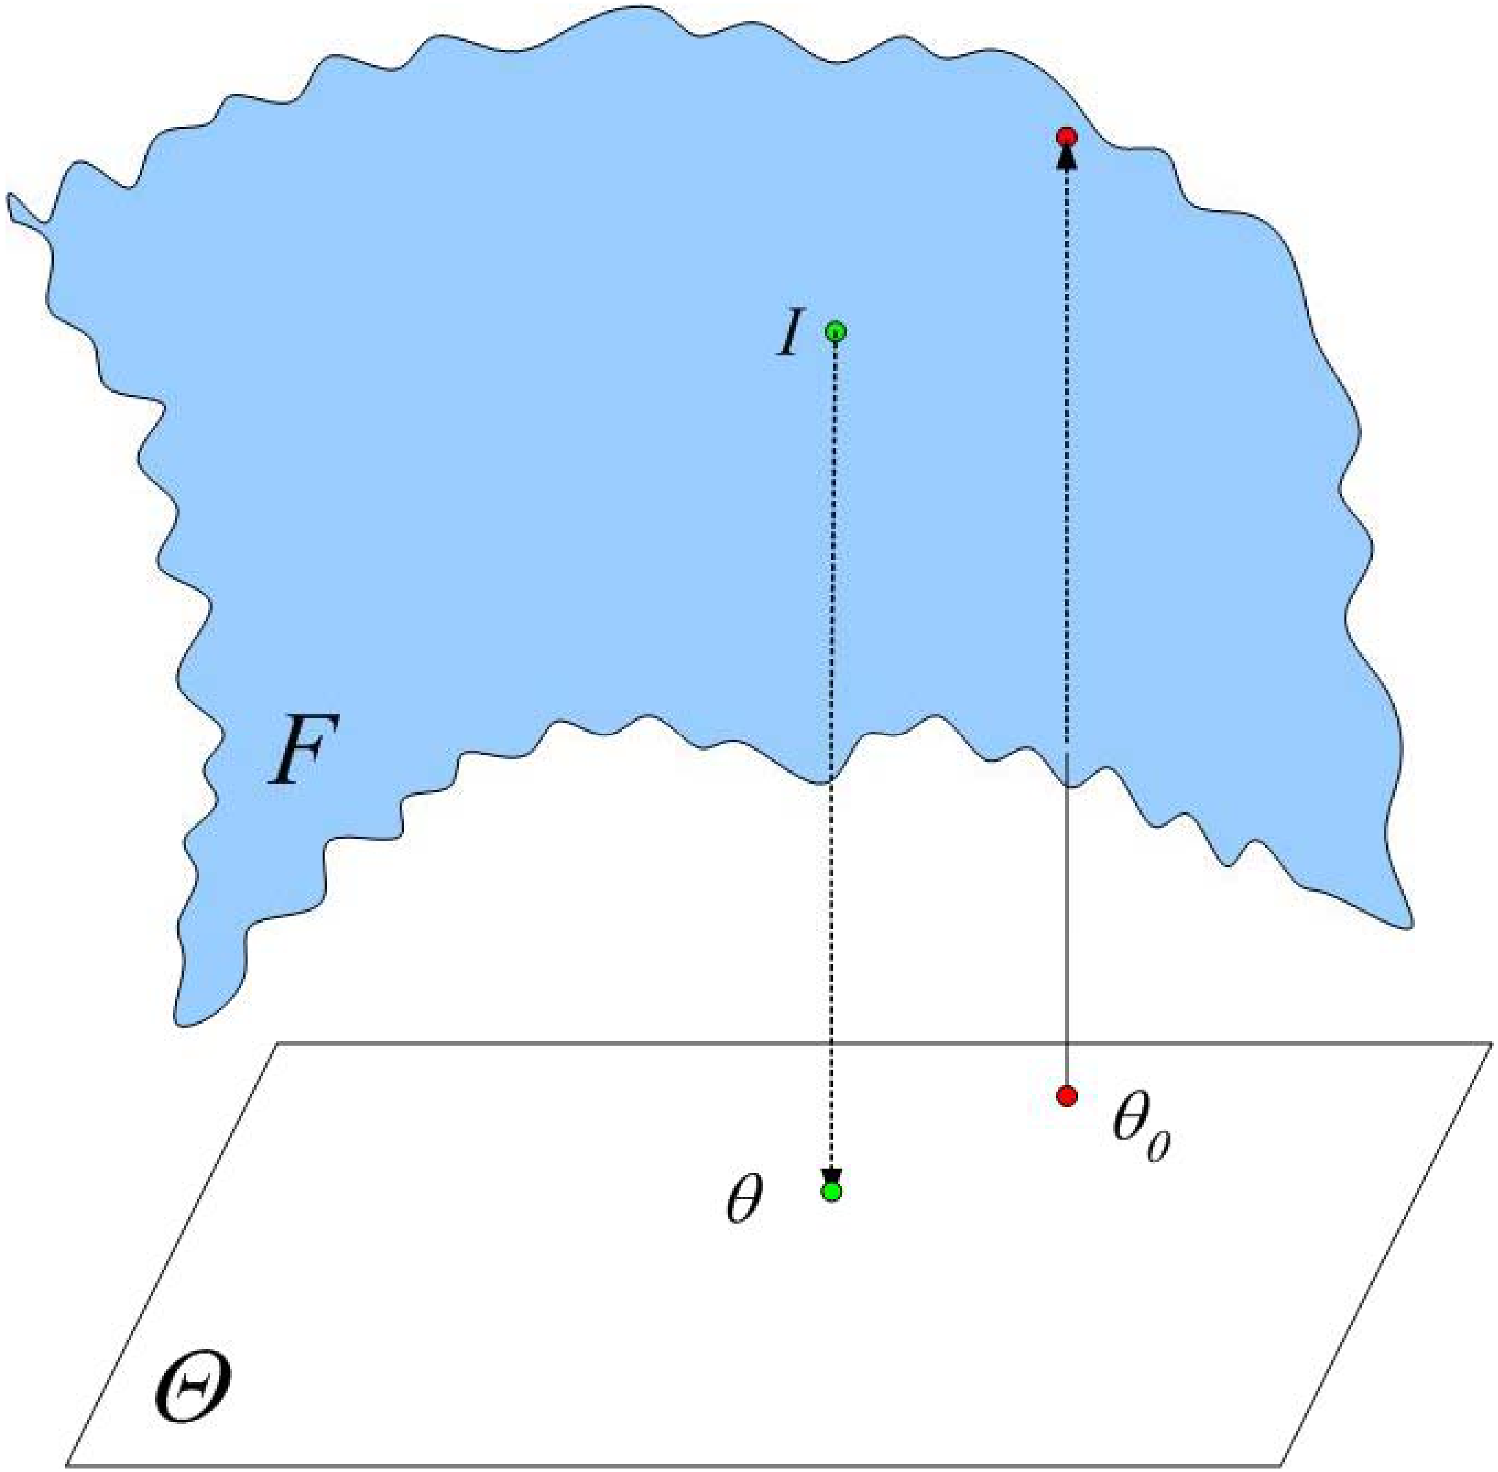
I
F
θ
θ_0
Θ

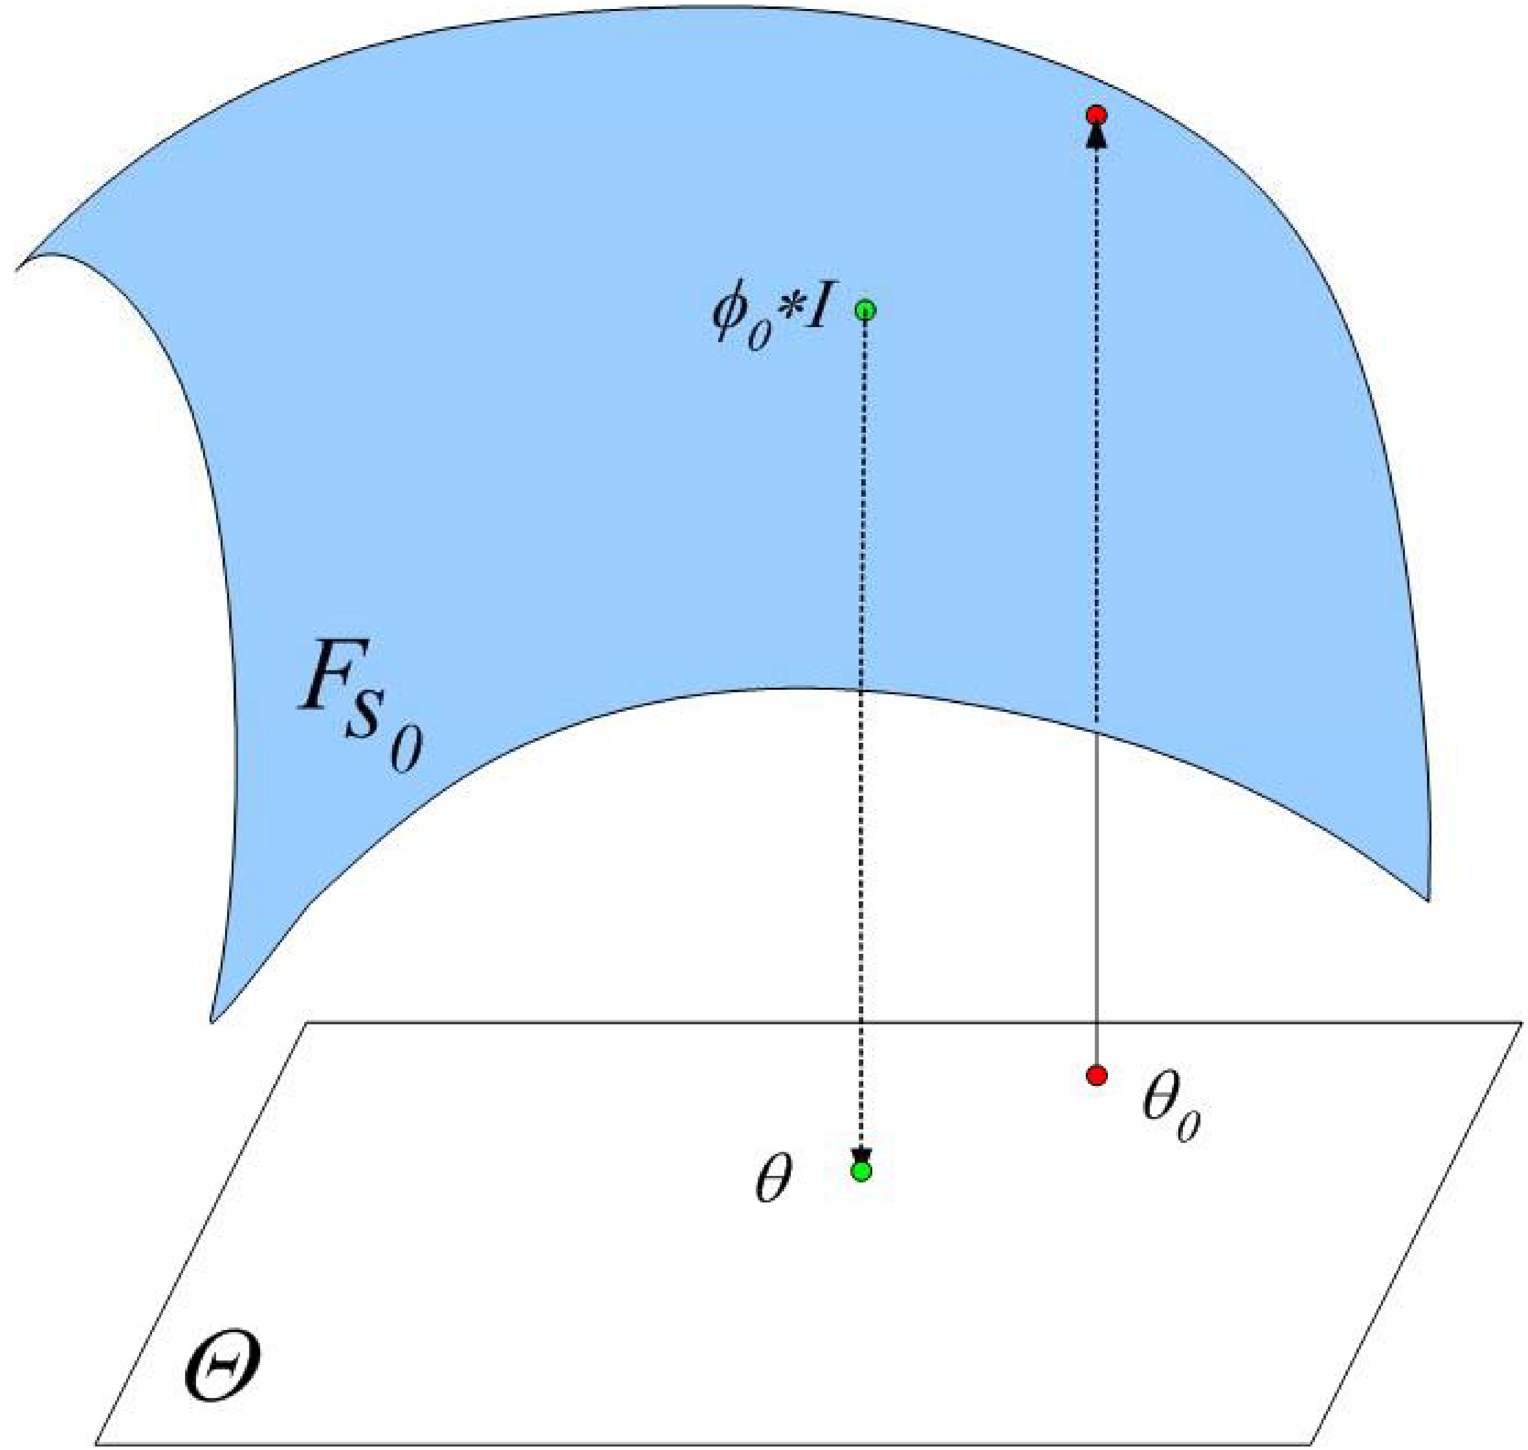

$\phi_0 * I$
F_{S_0}
θ
θ_0
Θ

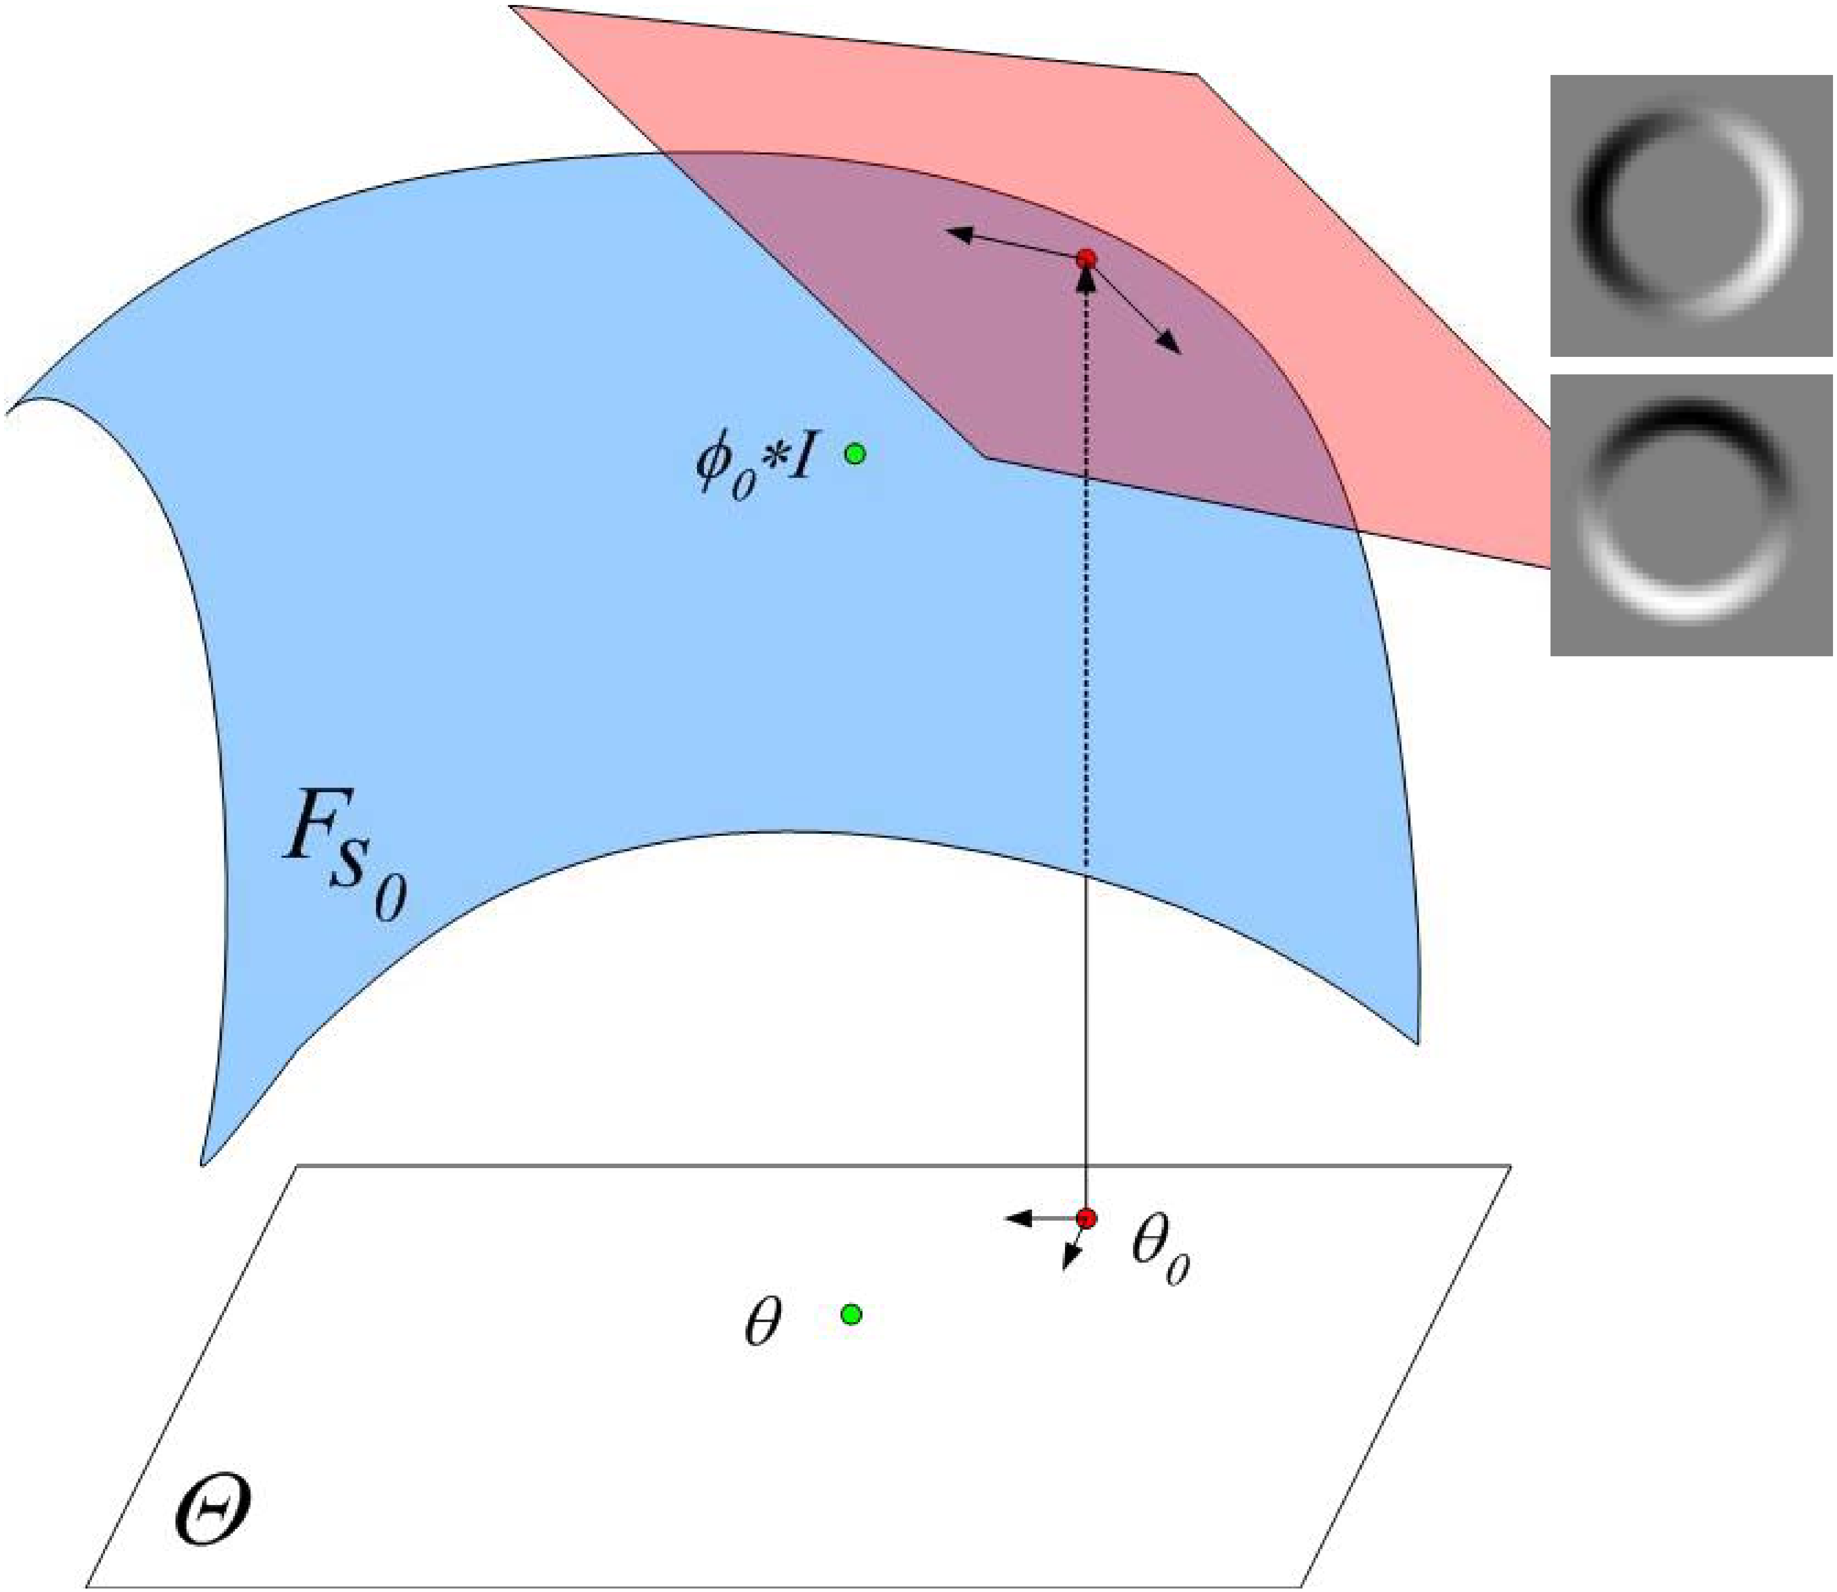

$\phi_0 * I$
F_{s_0}
θ_0
θ
Θ

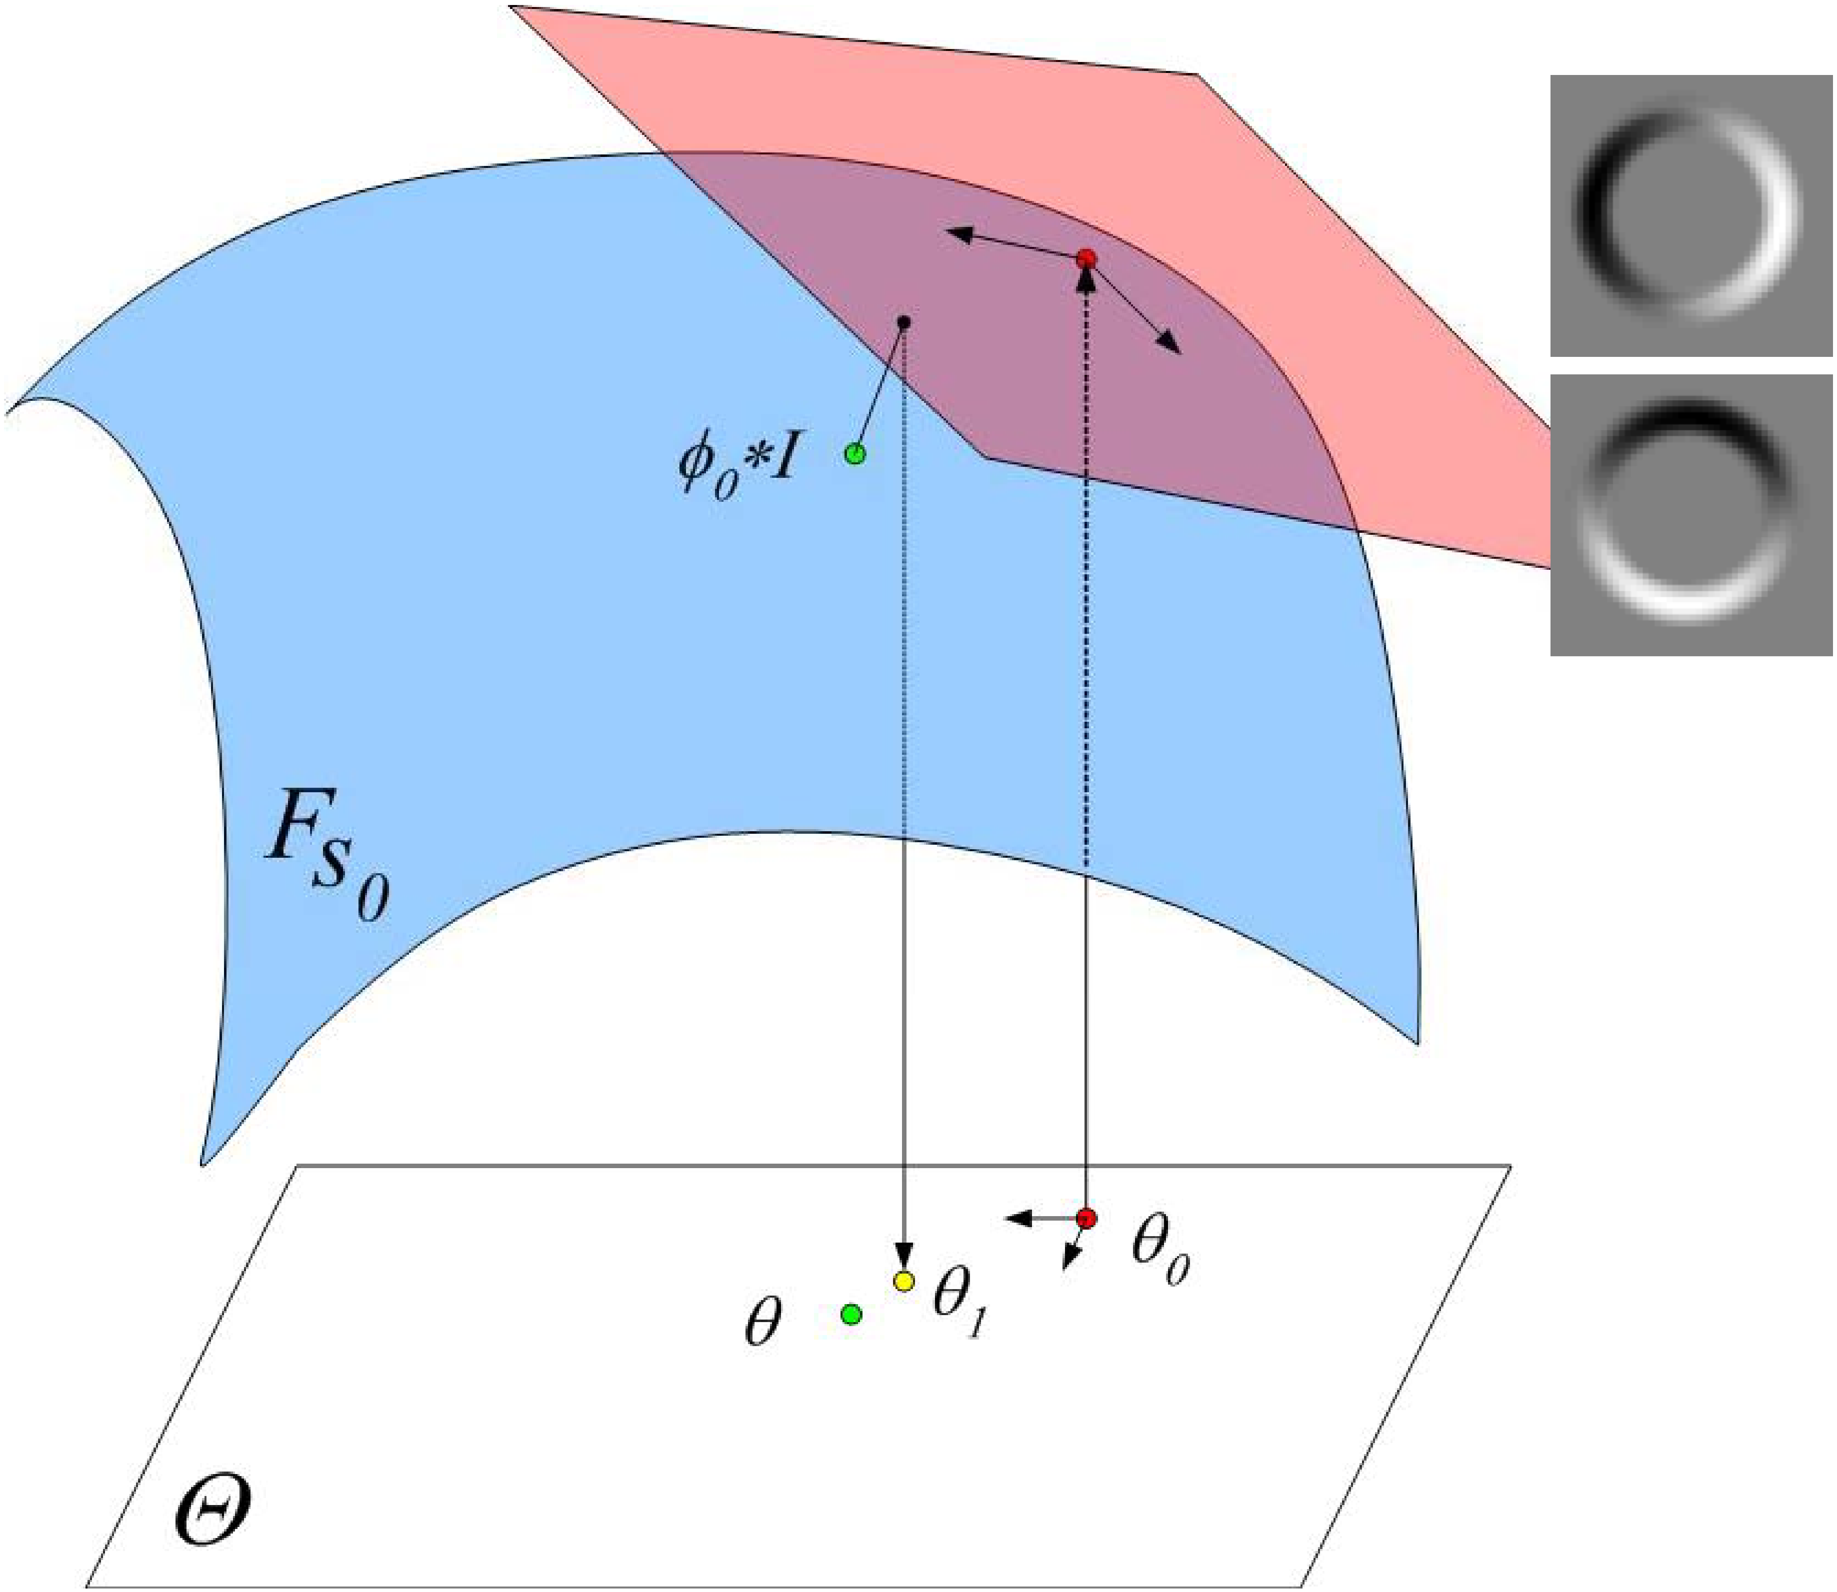

$\phi_0 * I$
F_{s_0}
θ_0
θ_1
θ
Θ

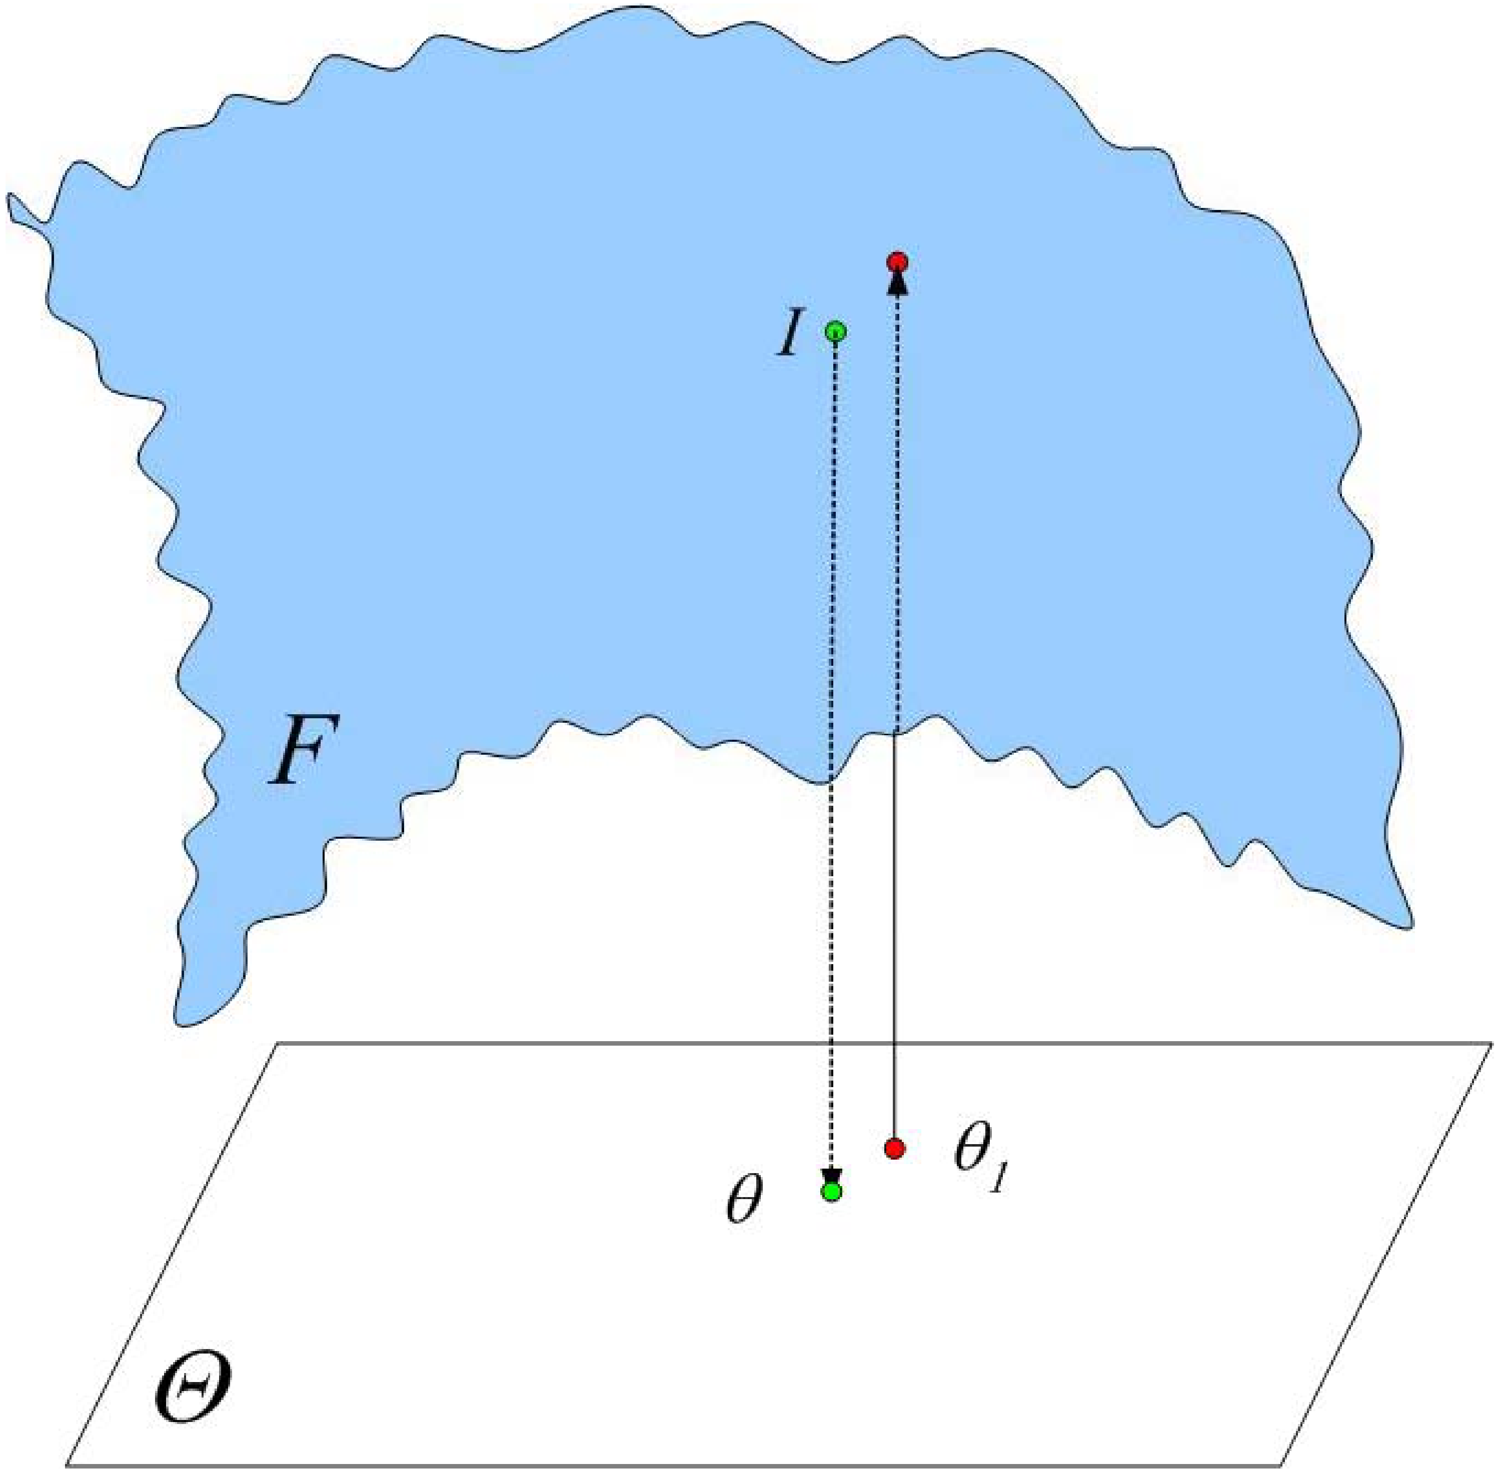
I
F
θ
θ1
Θ

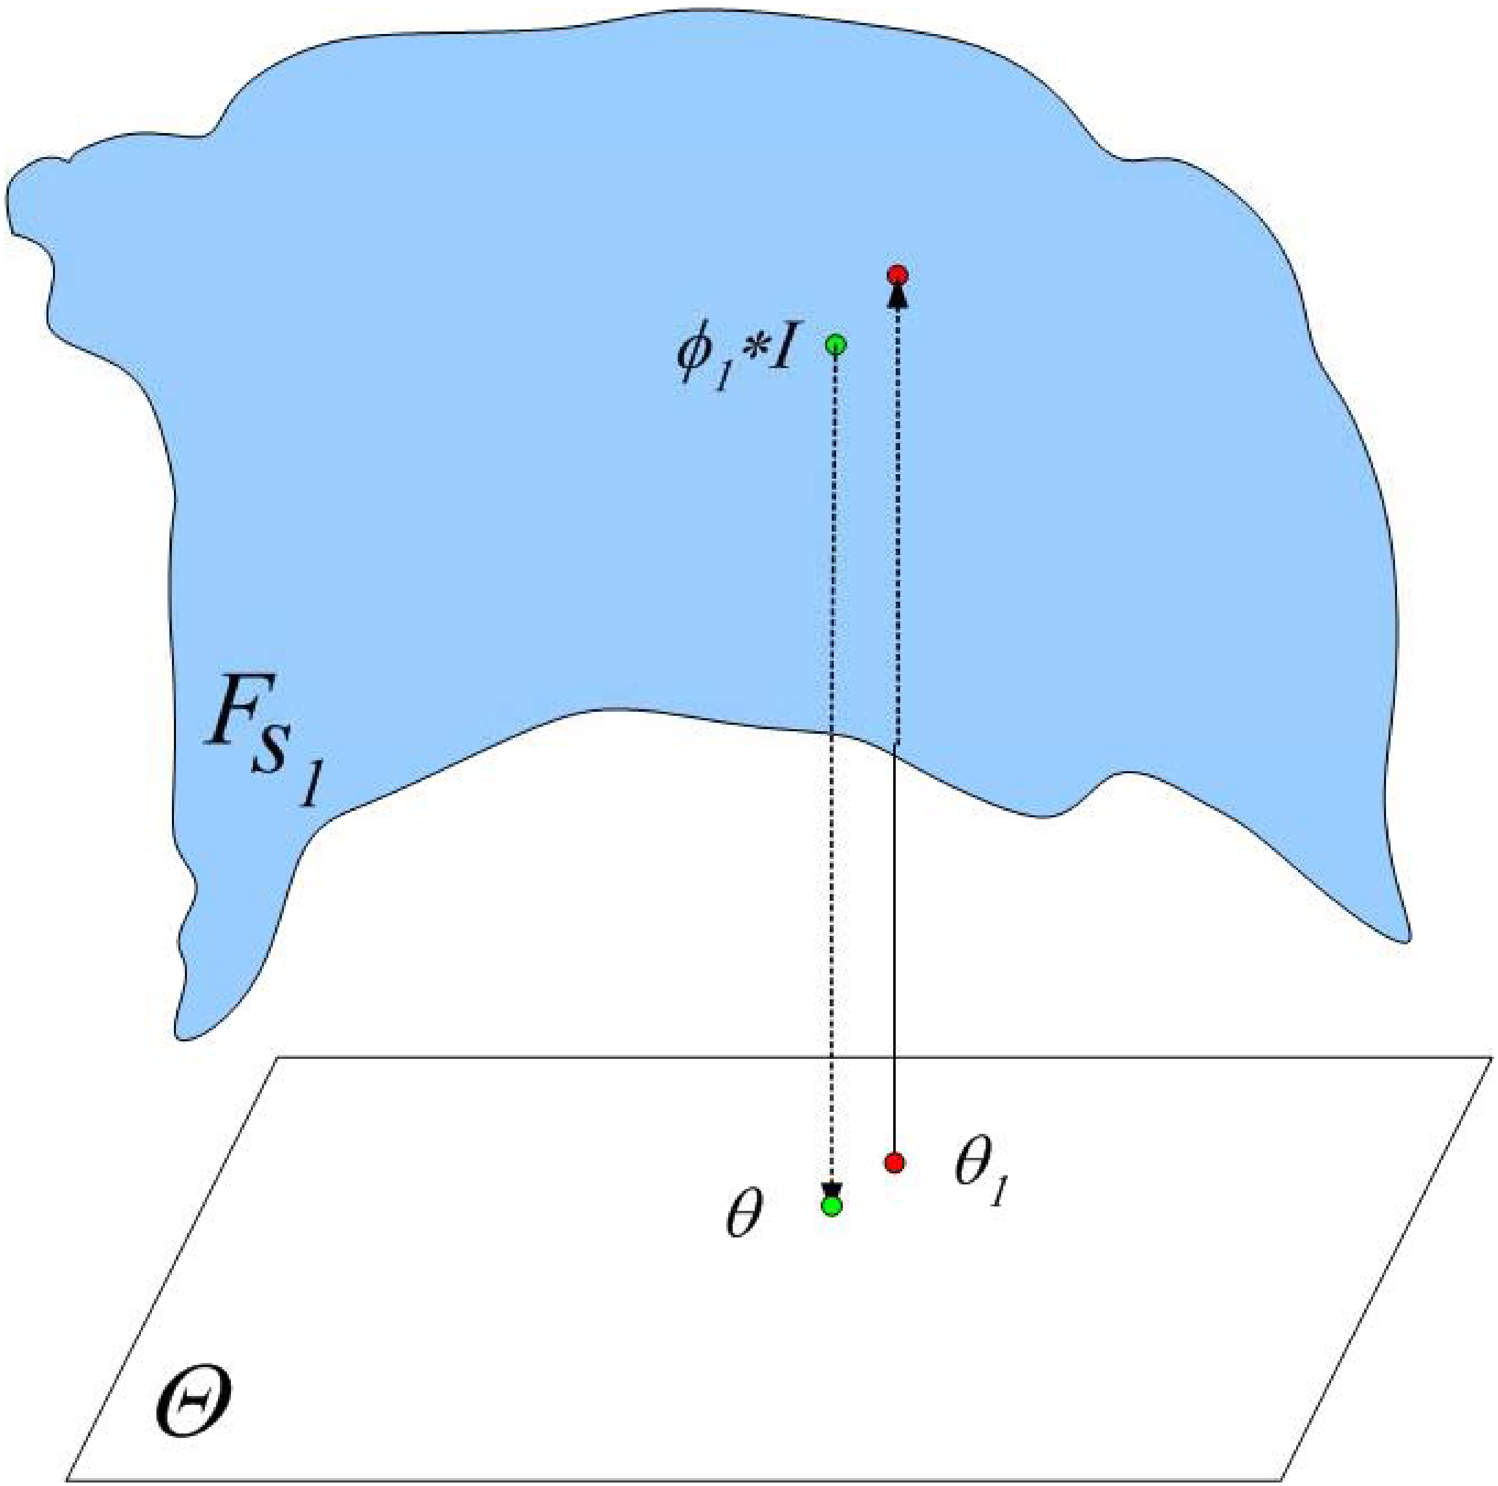

$\phi_1 * I$
F_{s_1}
θ
θ_1
Θ

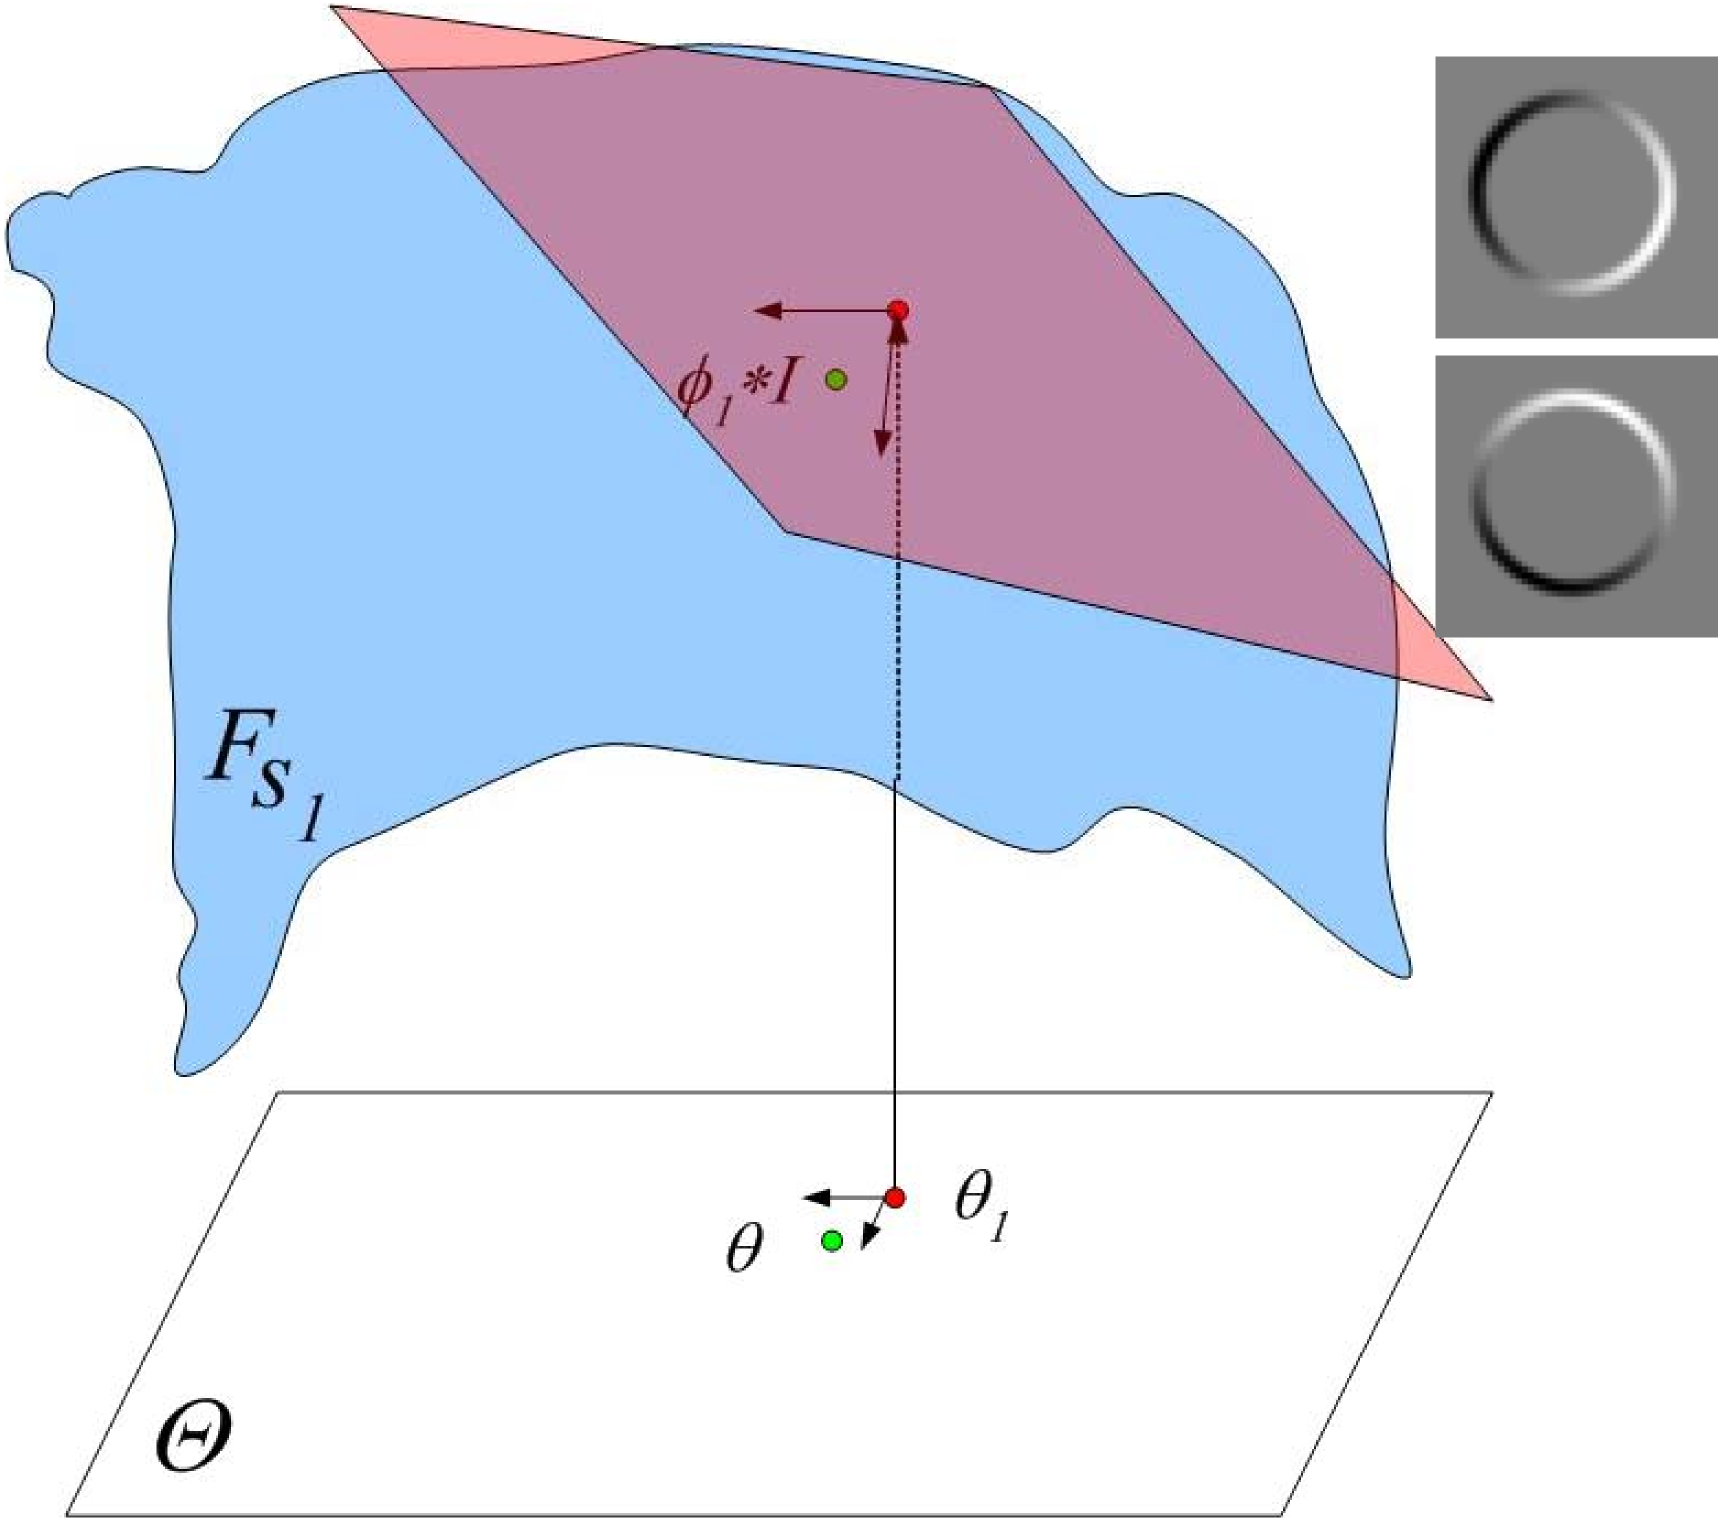

$\phi_1 * I$
F_{s_1}
θ_1
θ
Θ

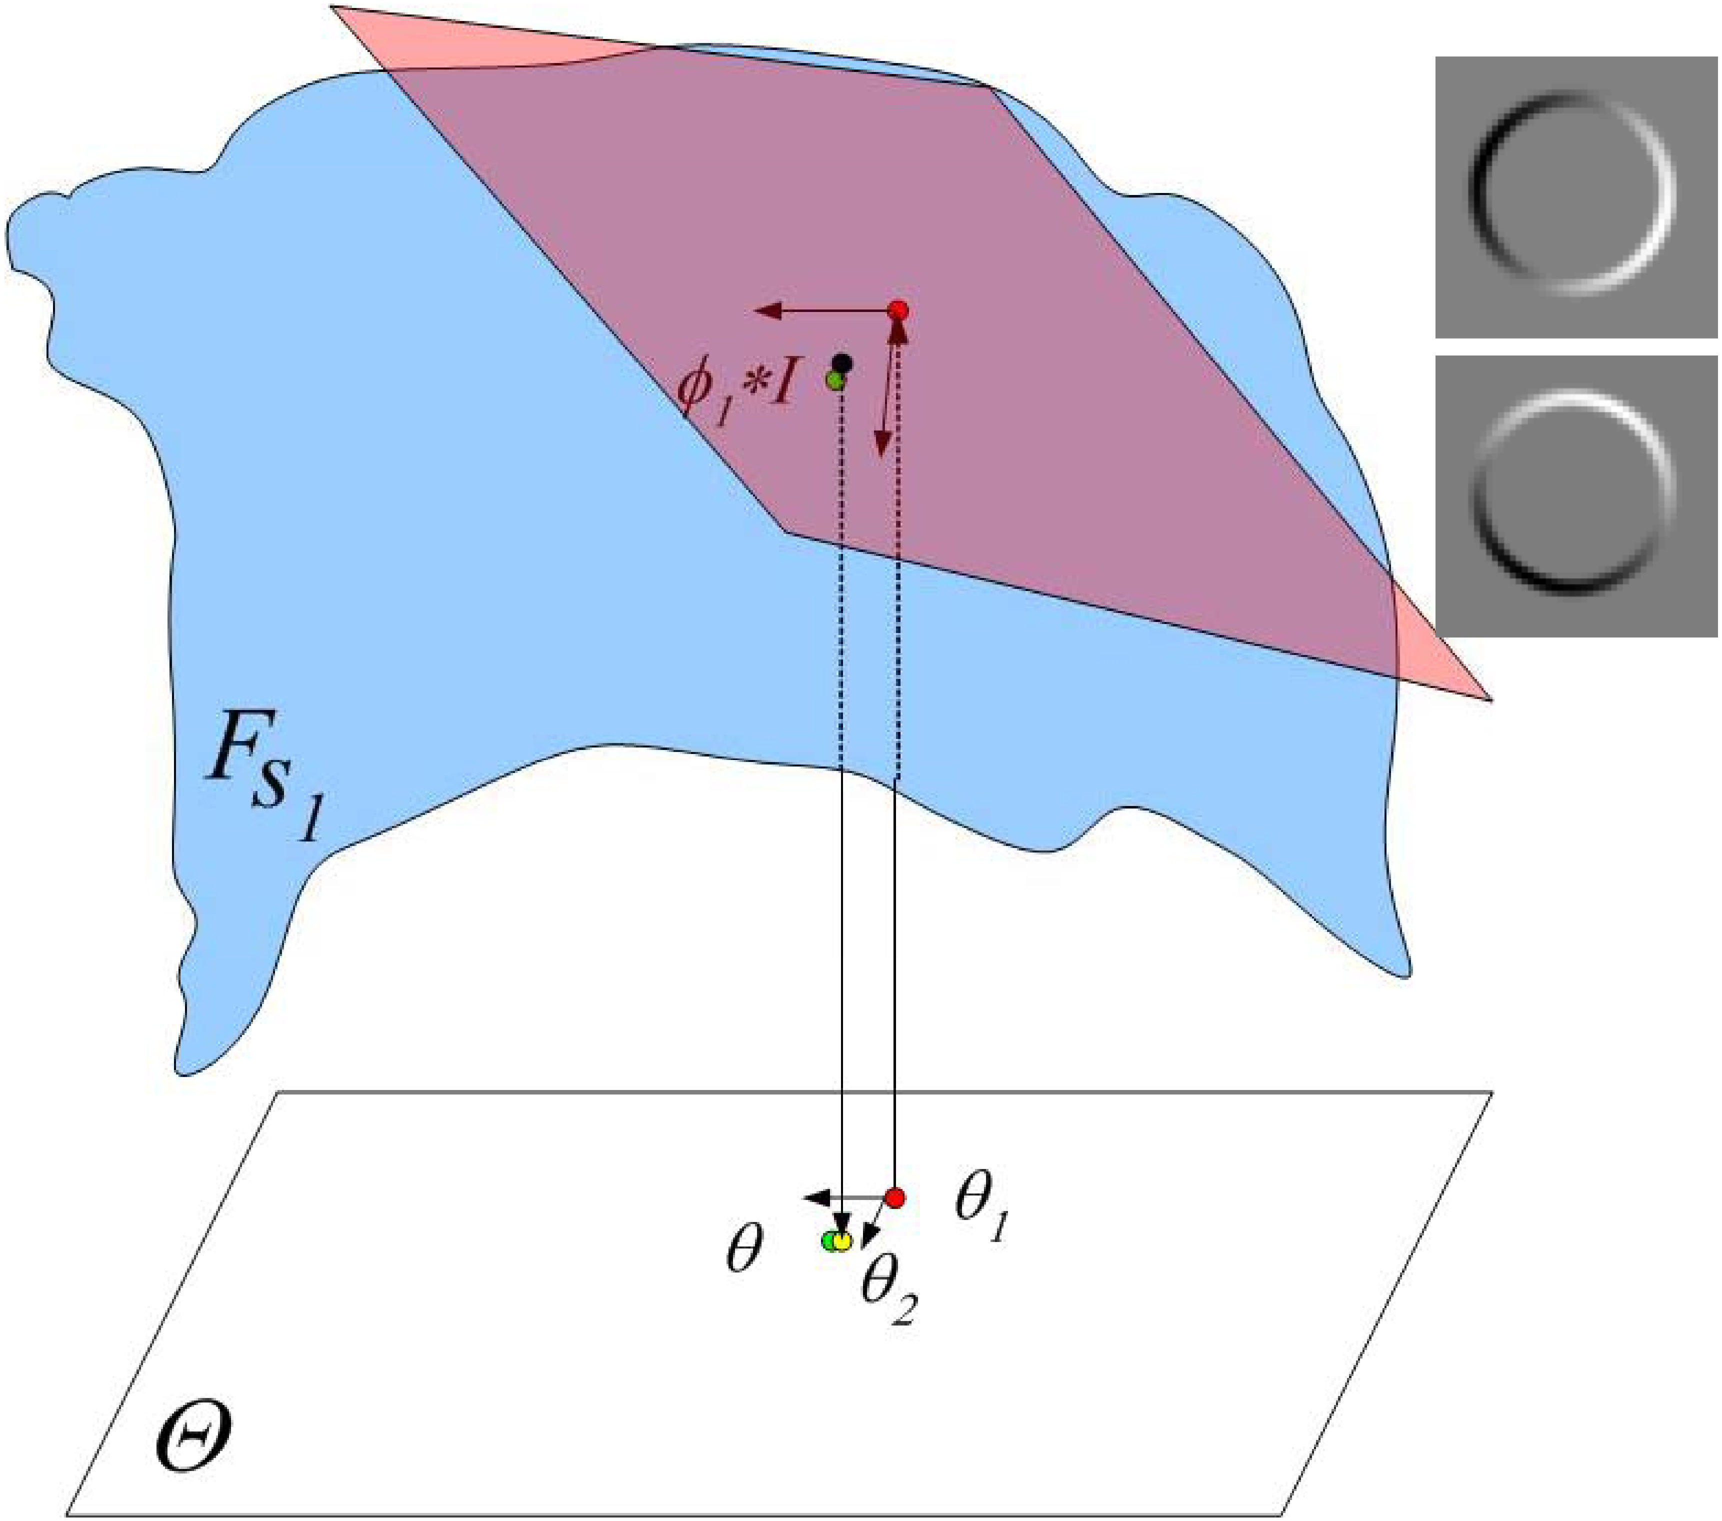

$\phi_1 * I$
F_{s_1}
θ_1
θ
θ_2
Θ

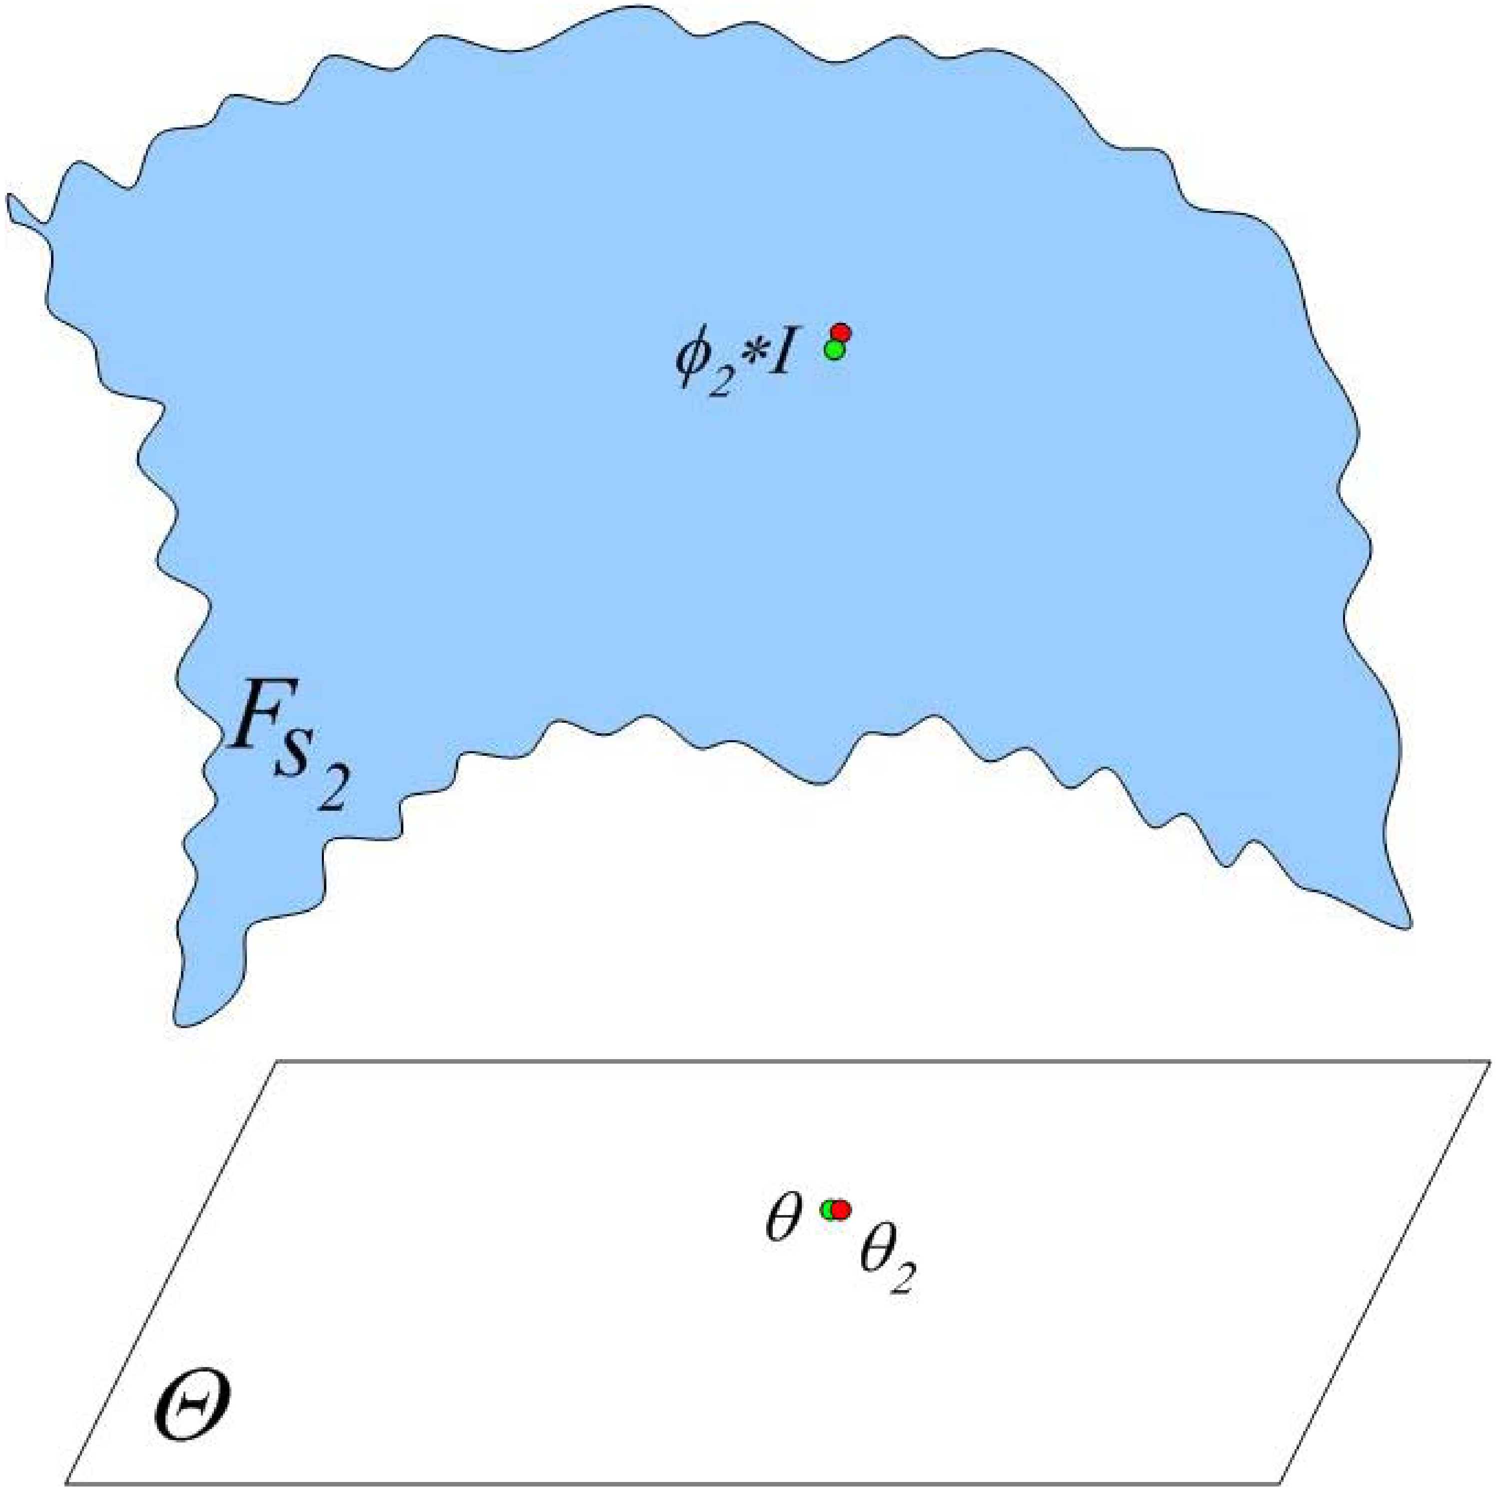
ϕ_2*I
F_{s_2}
θ
θ_2
Θ

New Perspective on Multiscale Techniques

- Image Registration & Coarse-to-Fine Differential Estimation
  - Irani/Peleg,
  - Belhumeur/Hager,
  - Keller/Averbach,
  - Simoncelli
  - & many others...

  ...all suggested by the geometry of the manifold

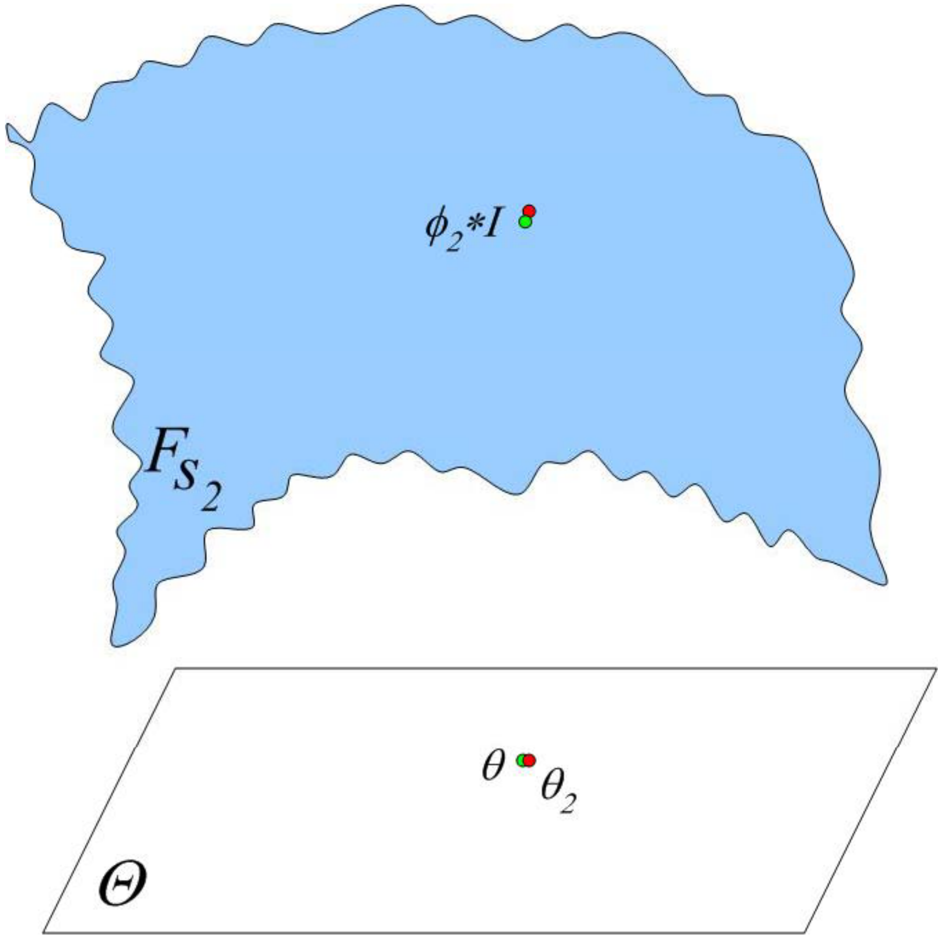

Experiments: Translating Disk

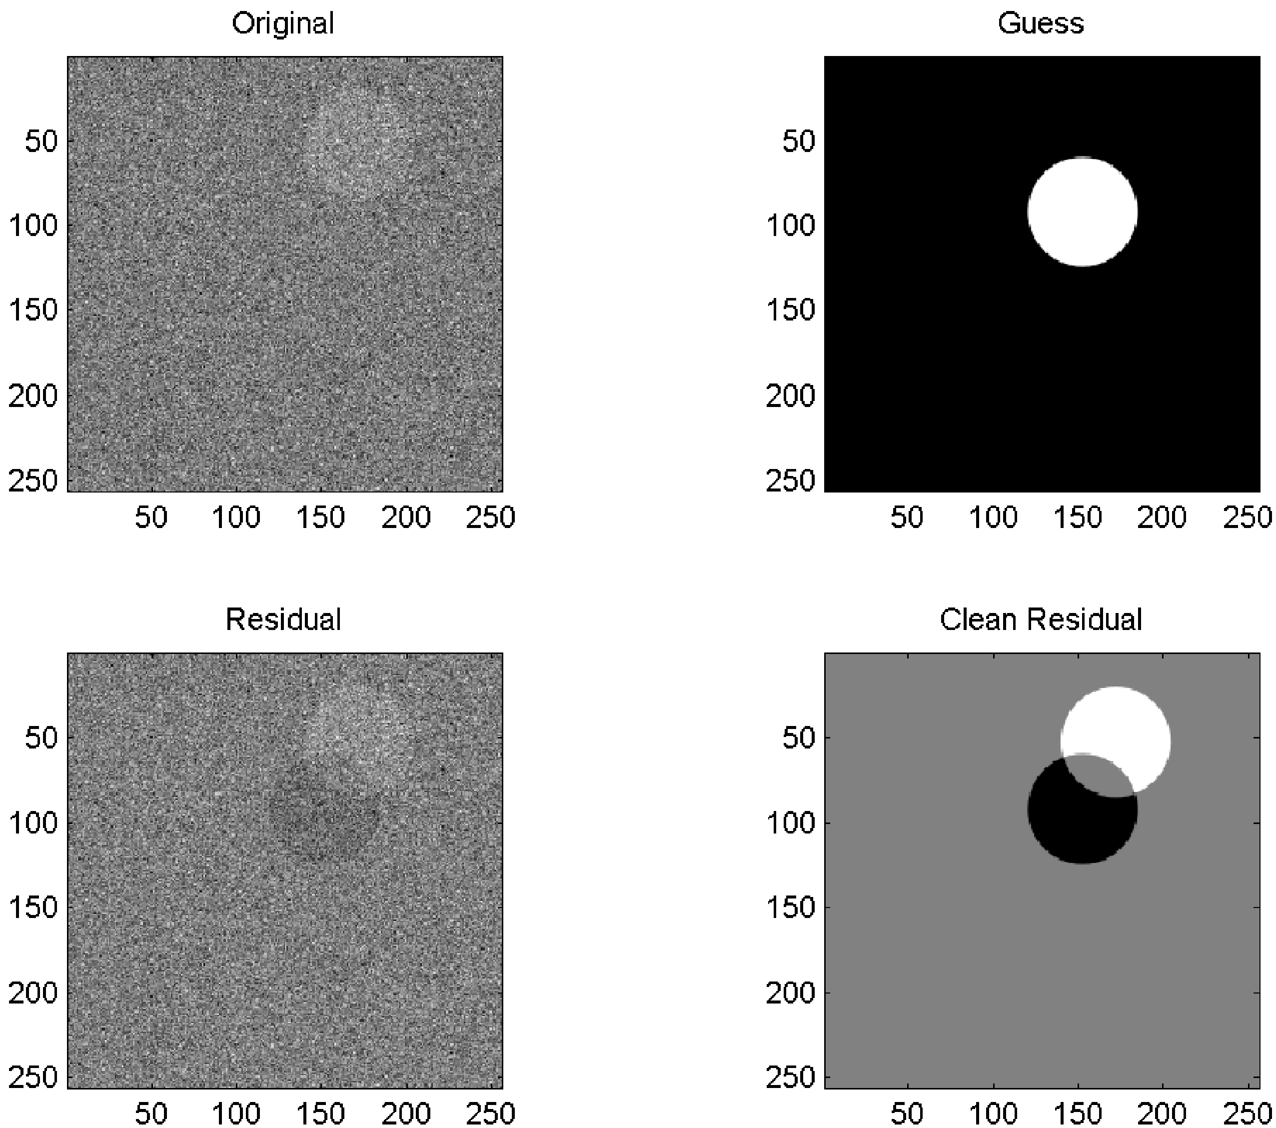

$s = 1/2$

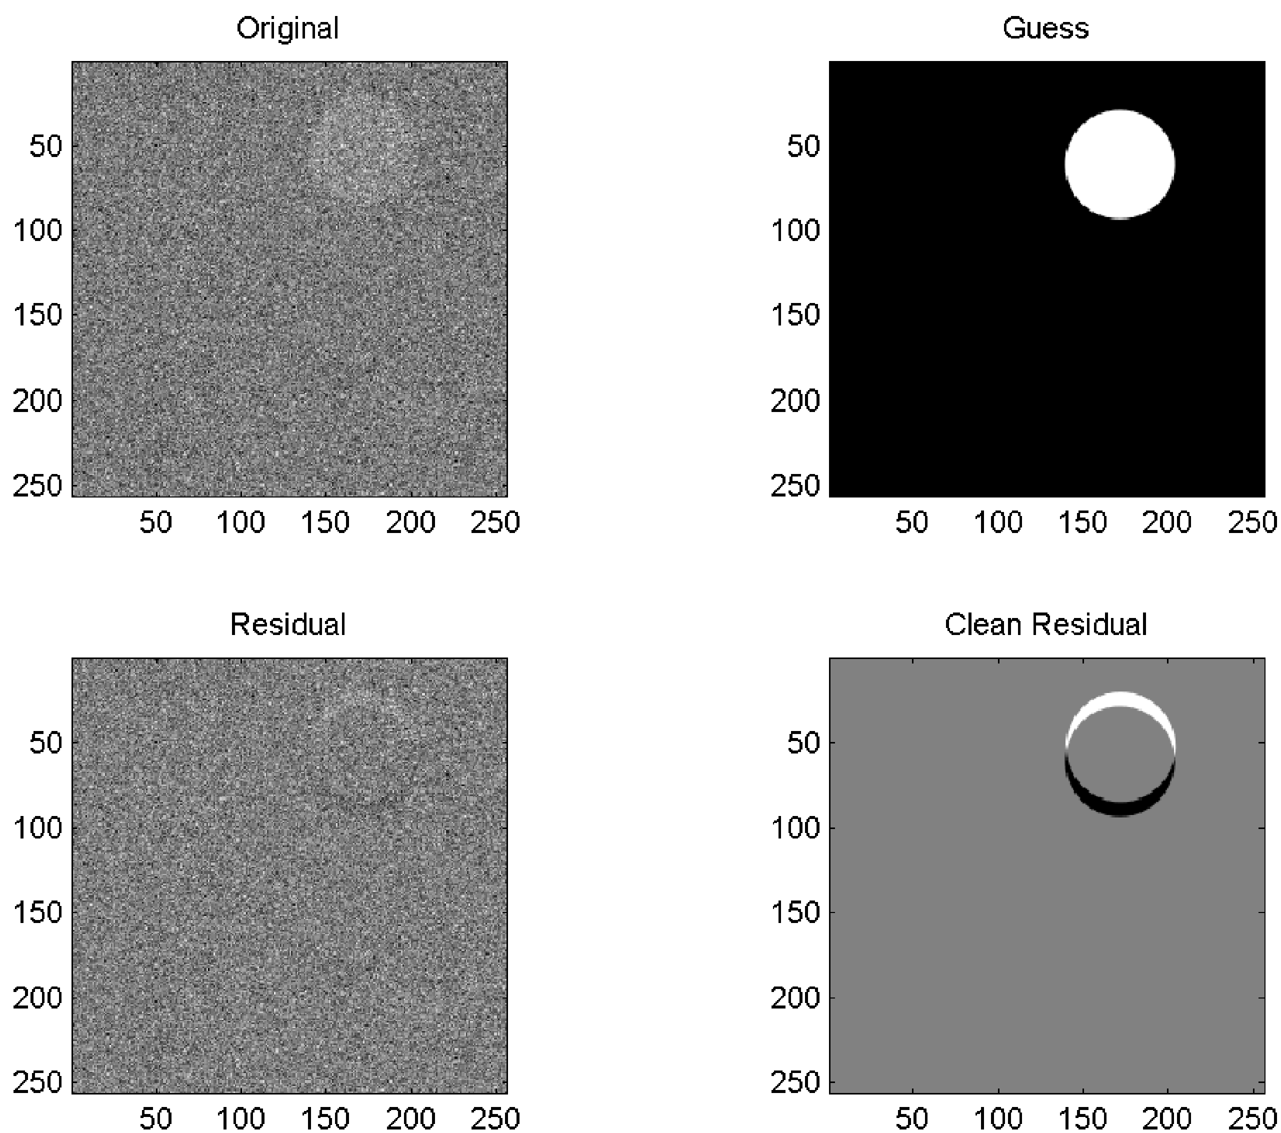

$s = 1/4$

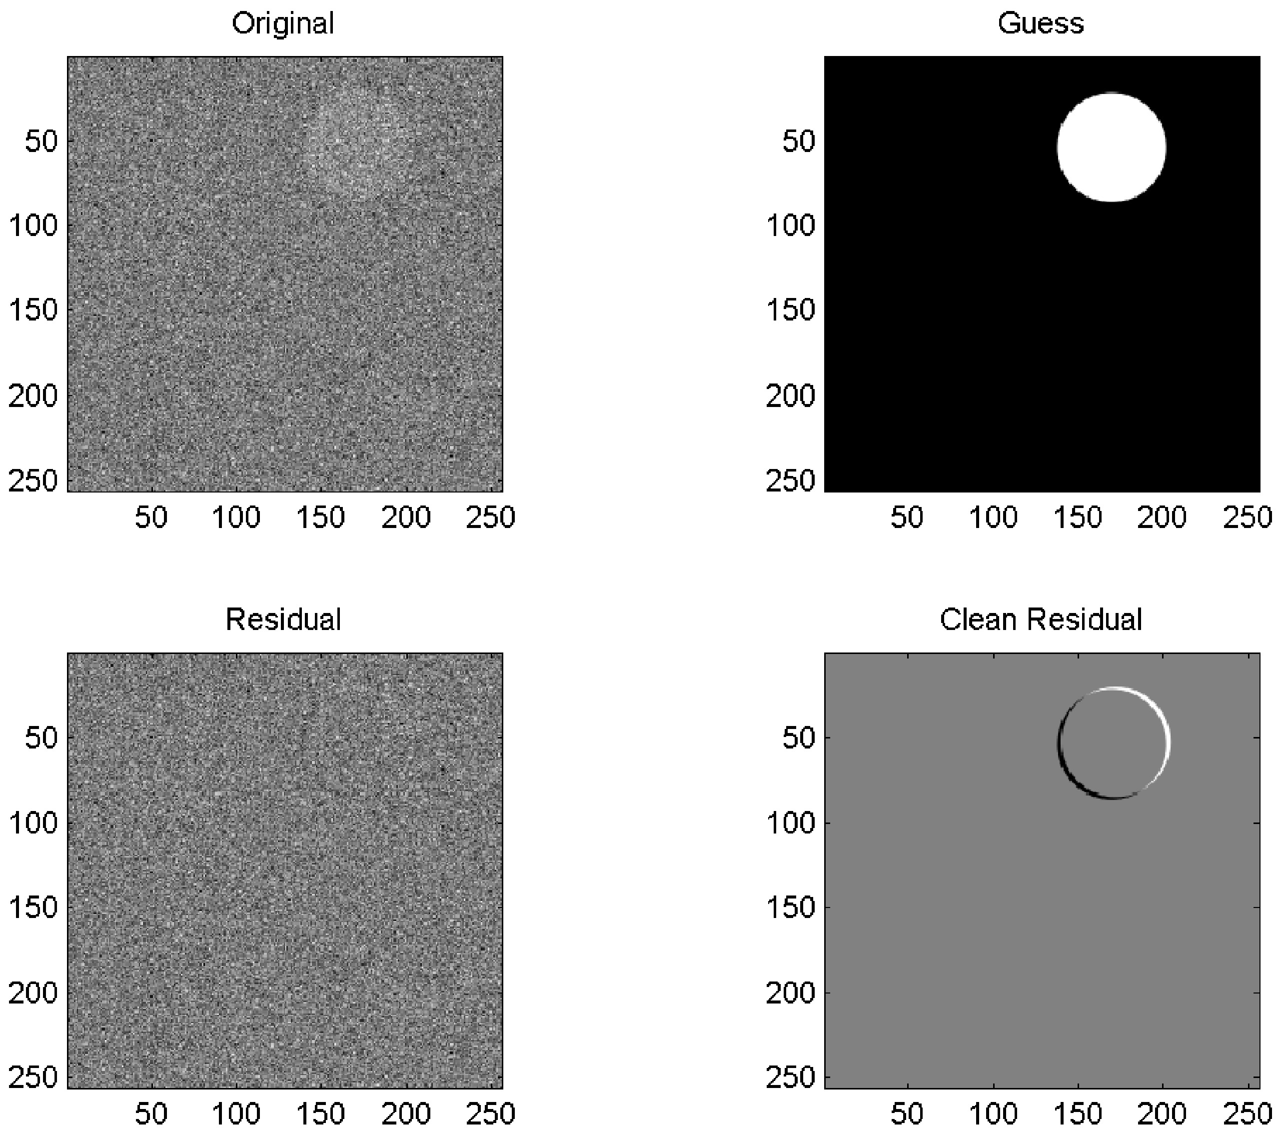
Original
Guess
Residual
Clean Residual
50
100
150
200
250
50
100
150
200
250

$s = 1/16$

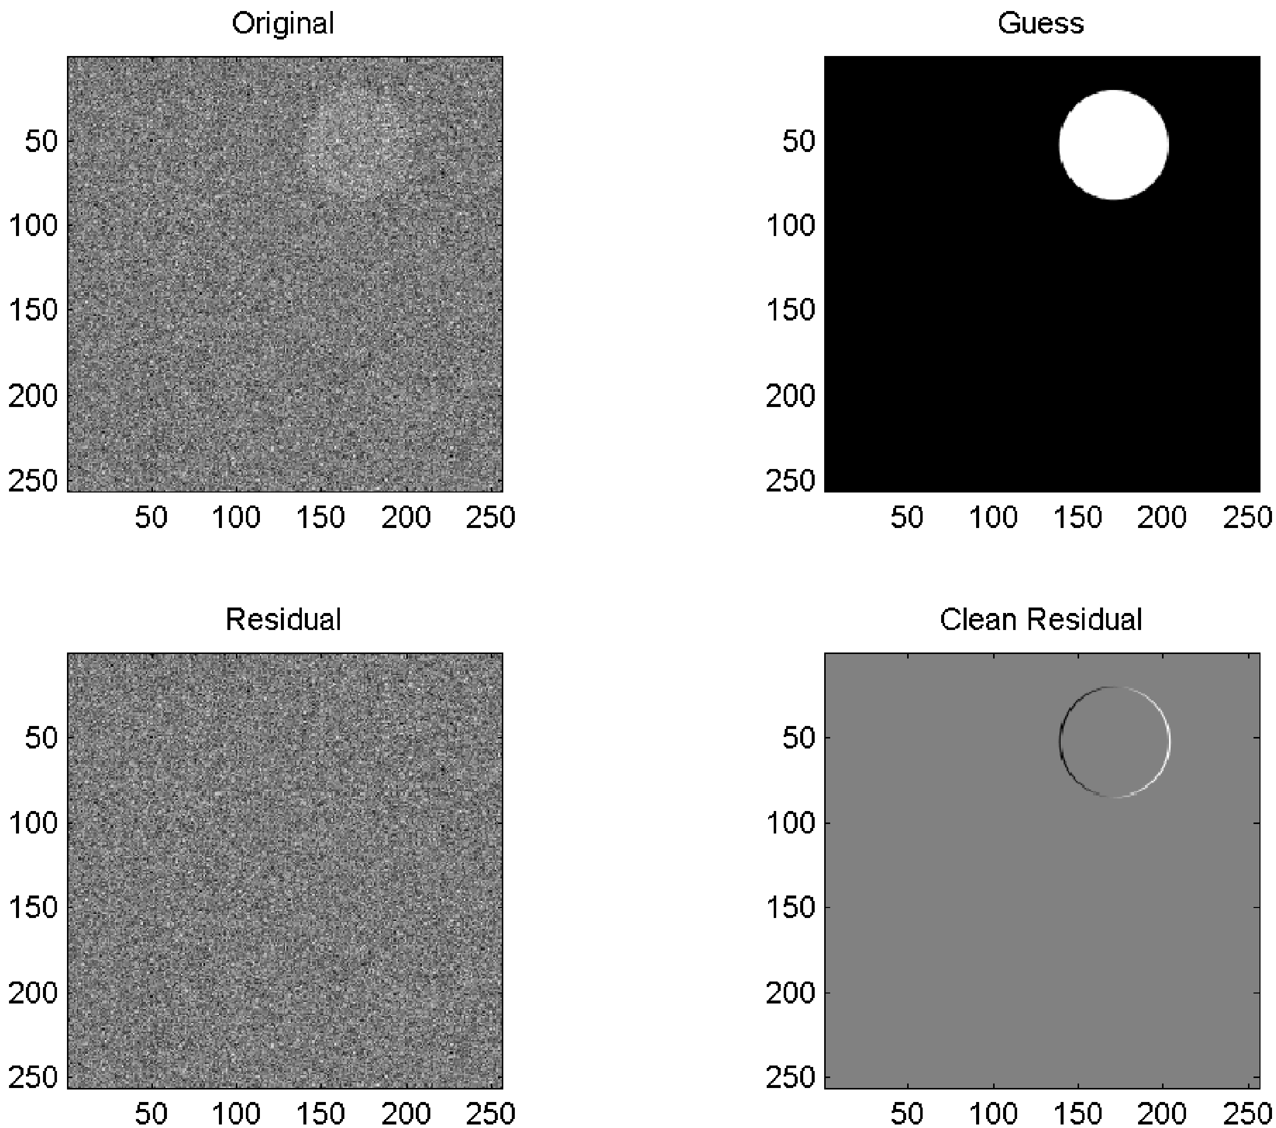

$s = 1/256$

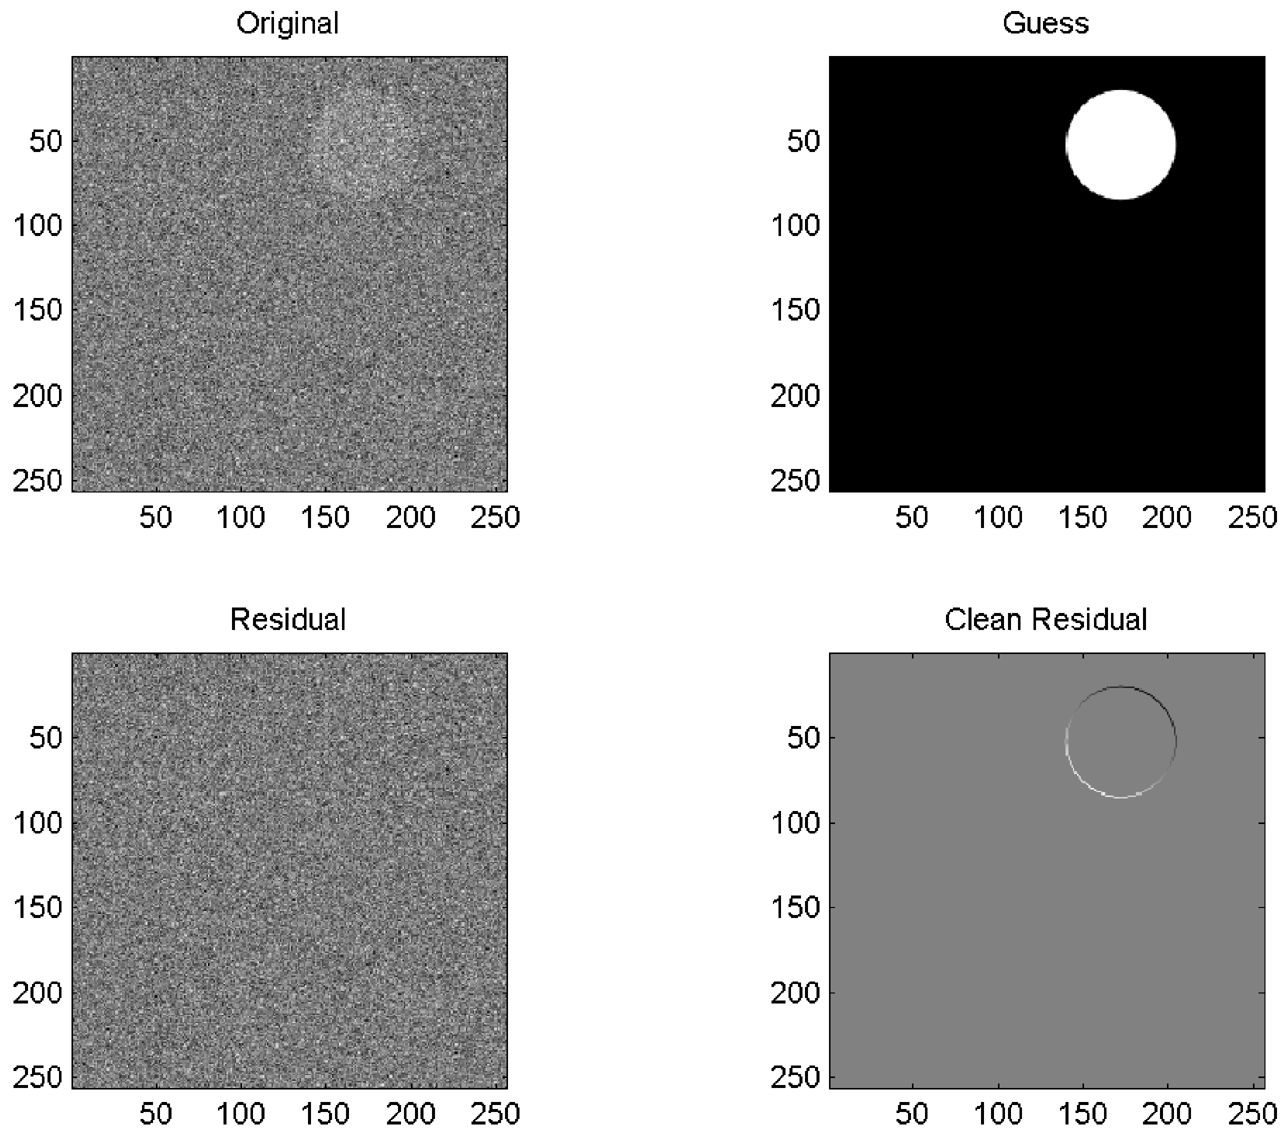

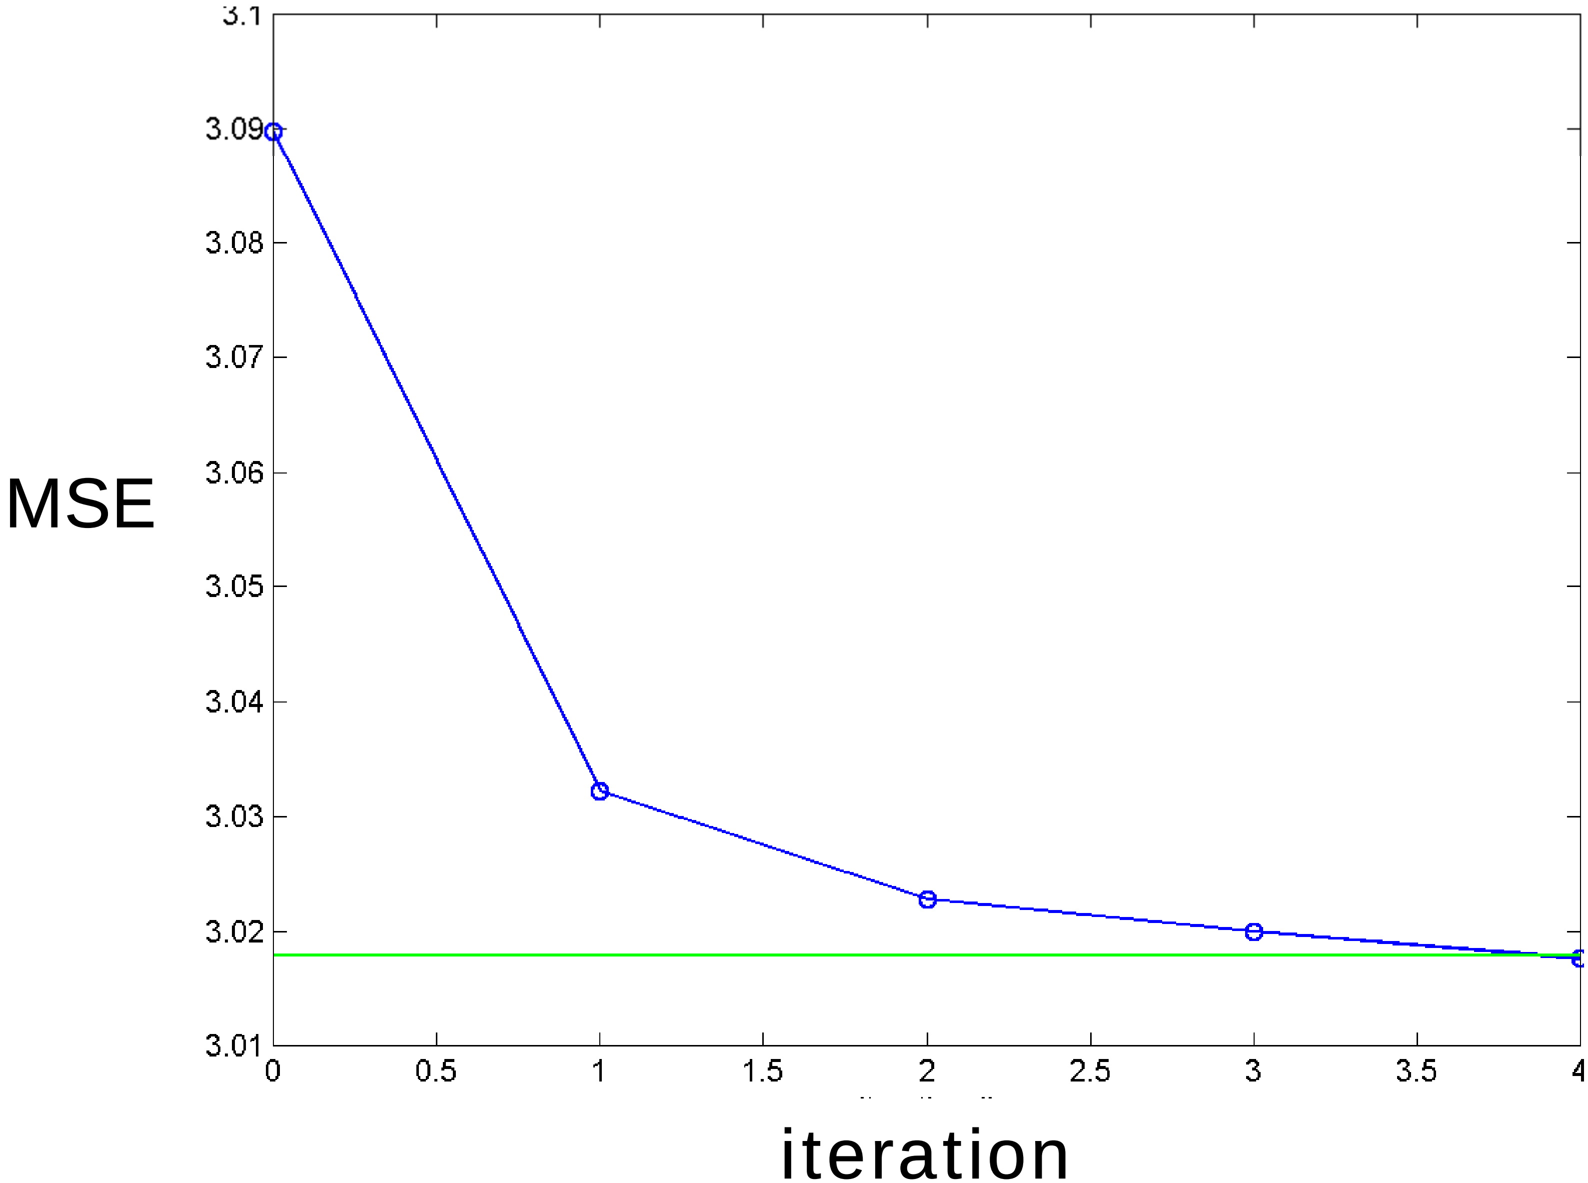

MSE
3.1
3.09
3.08
3.07
3.06
3.05
3.04
3.03
3.02
3.01
0
0.5
1
1.5
2
2.5
3
3.5
4
iteration

Experiments: Rotating 3-D Cube

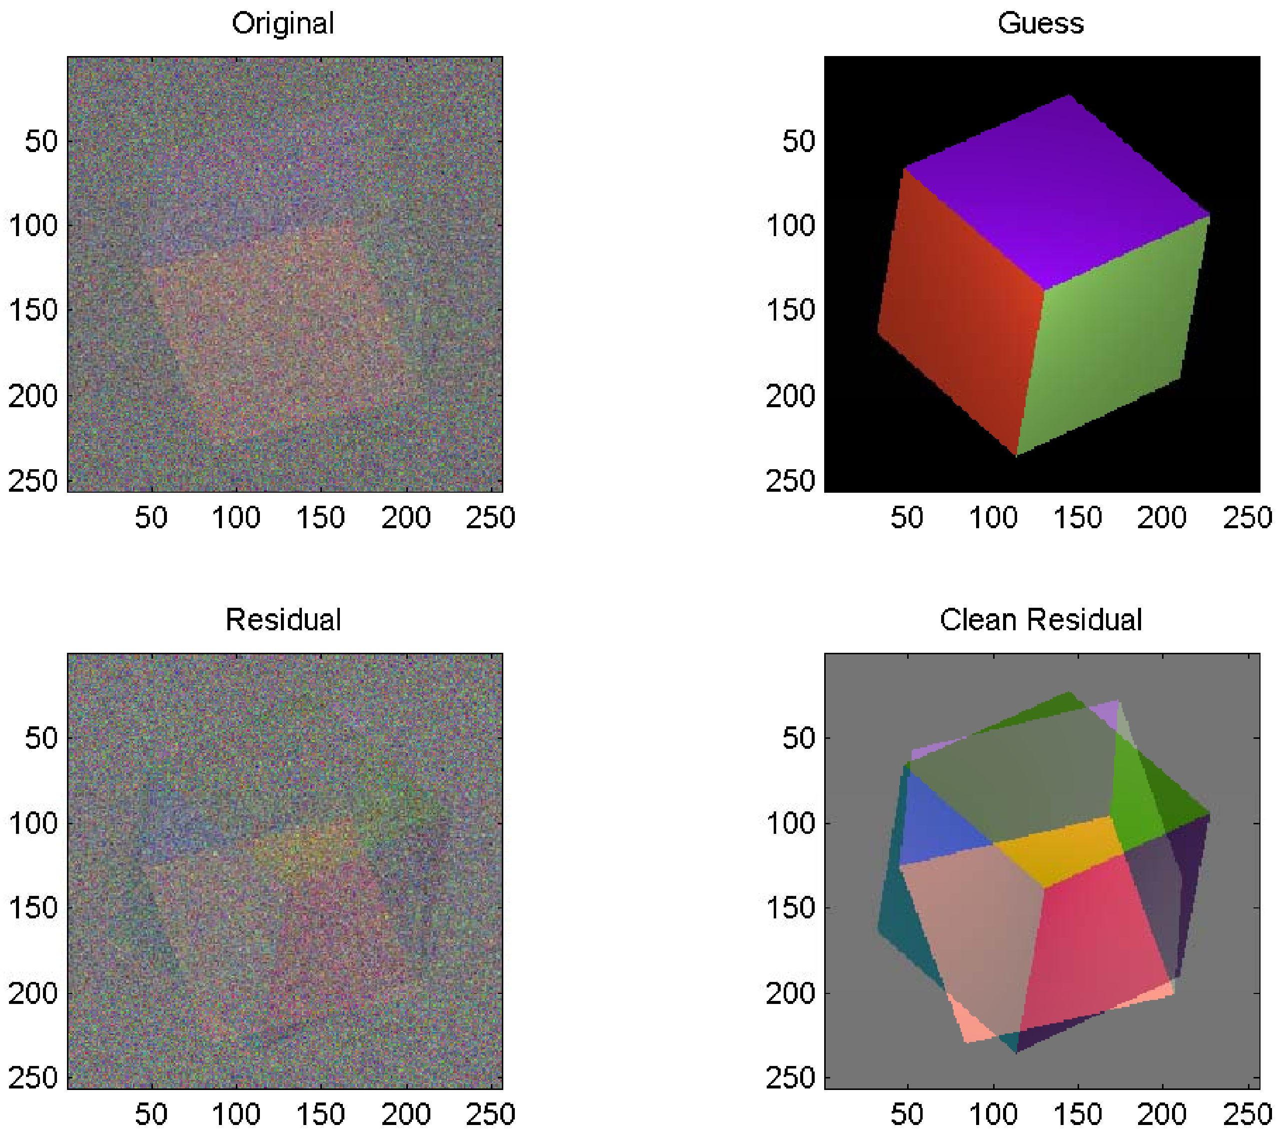

$s = 1/2$

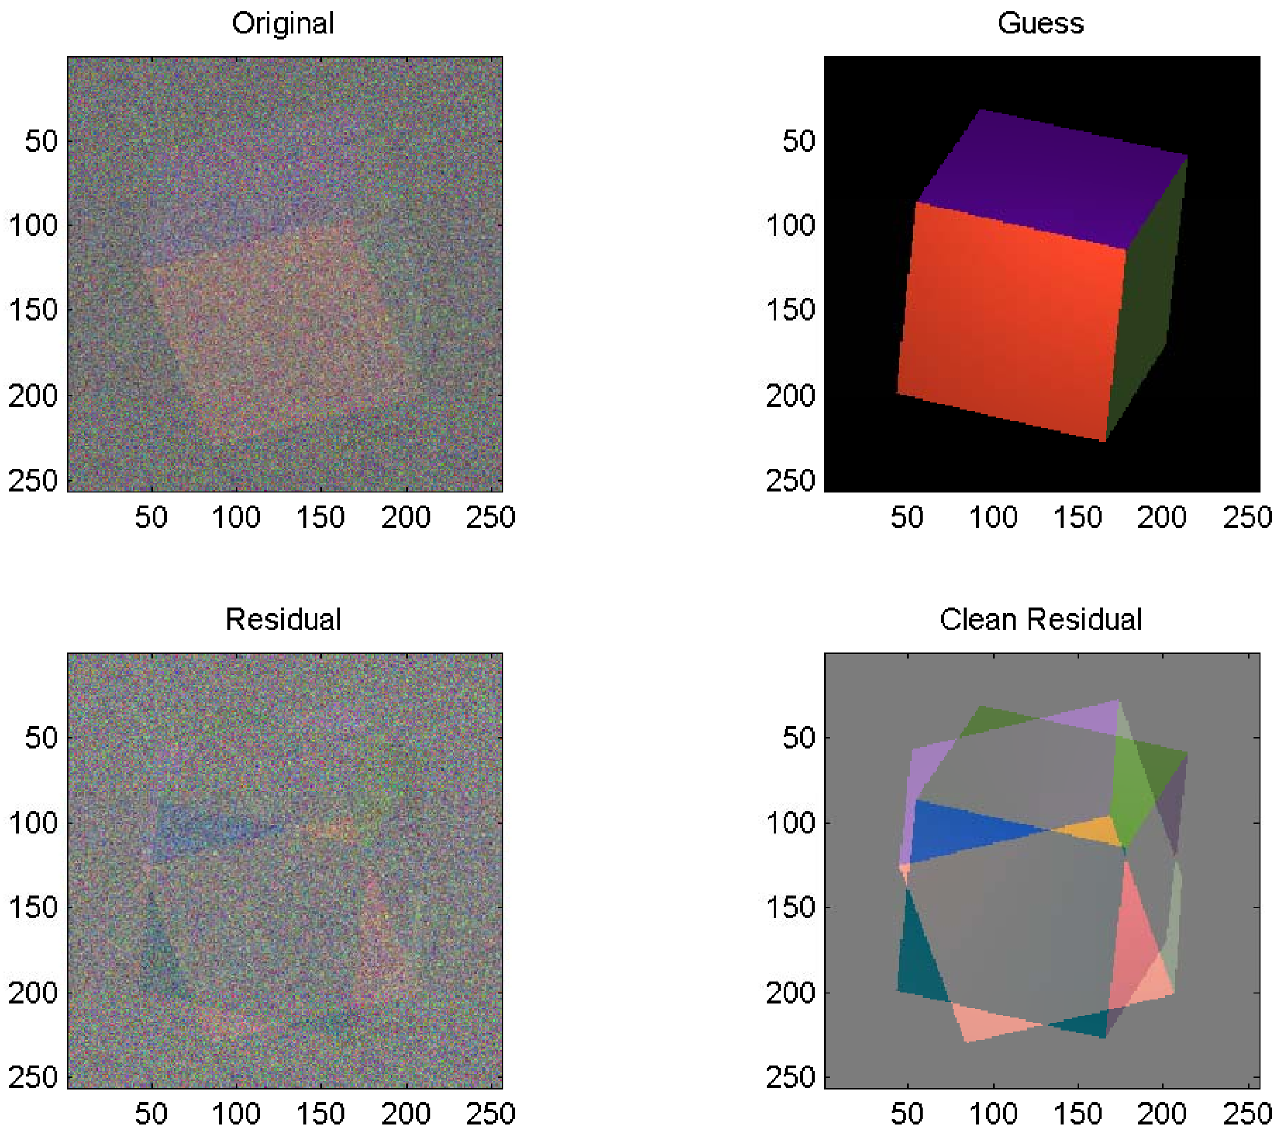

$s = 1/4$

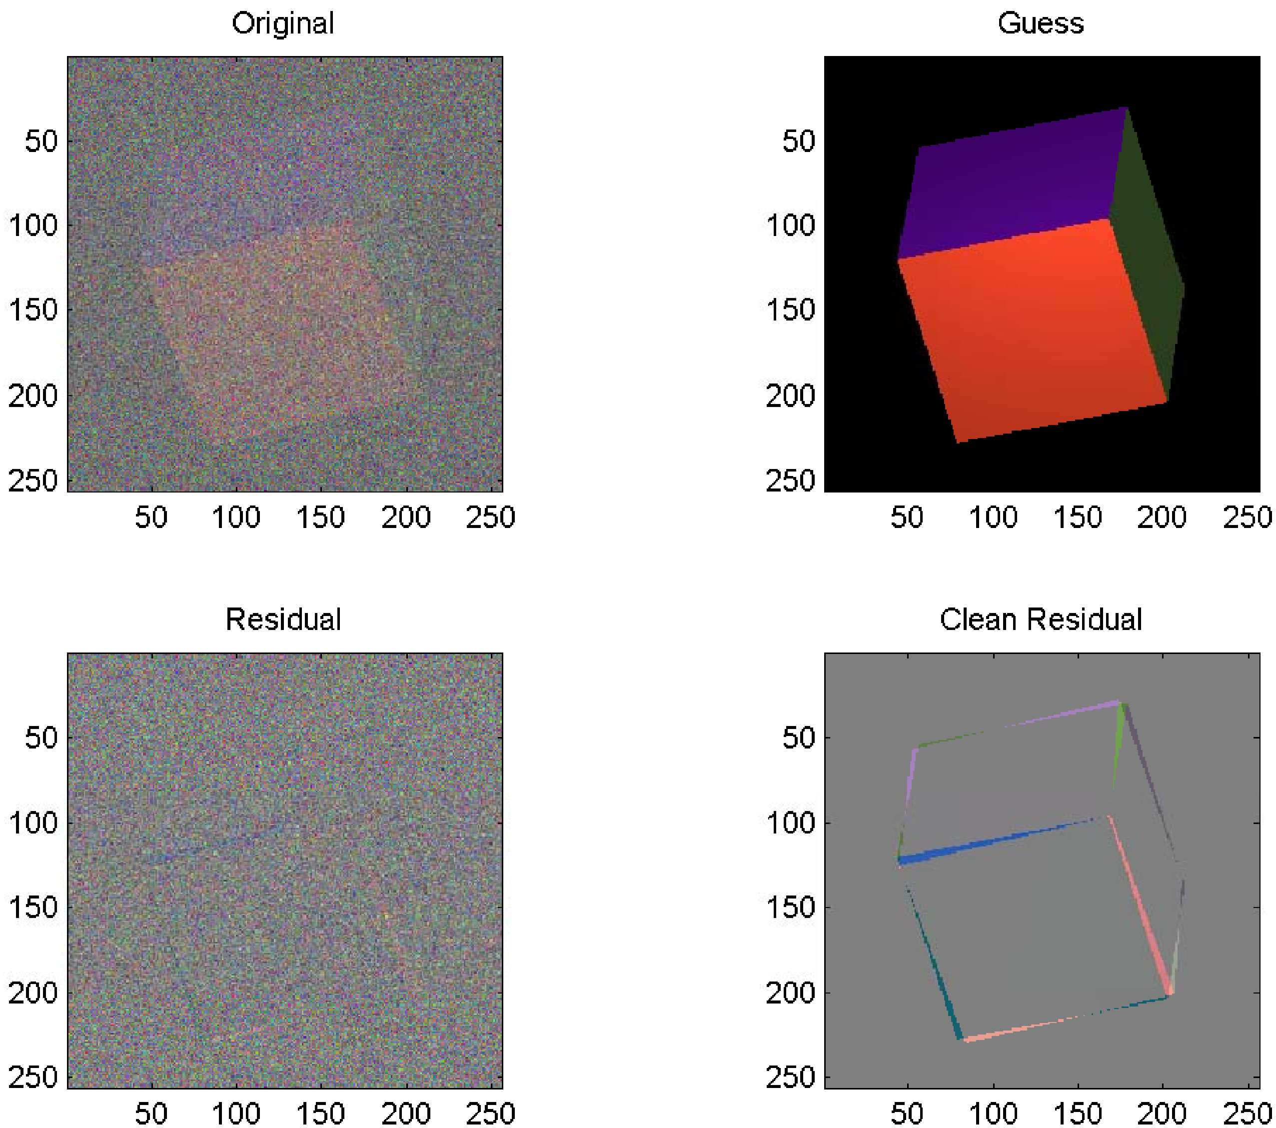

$s = 1/16$

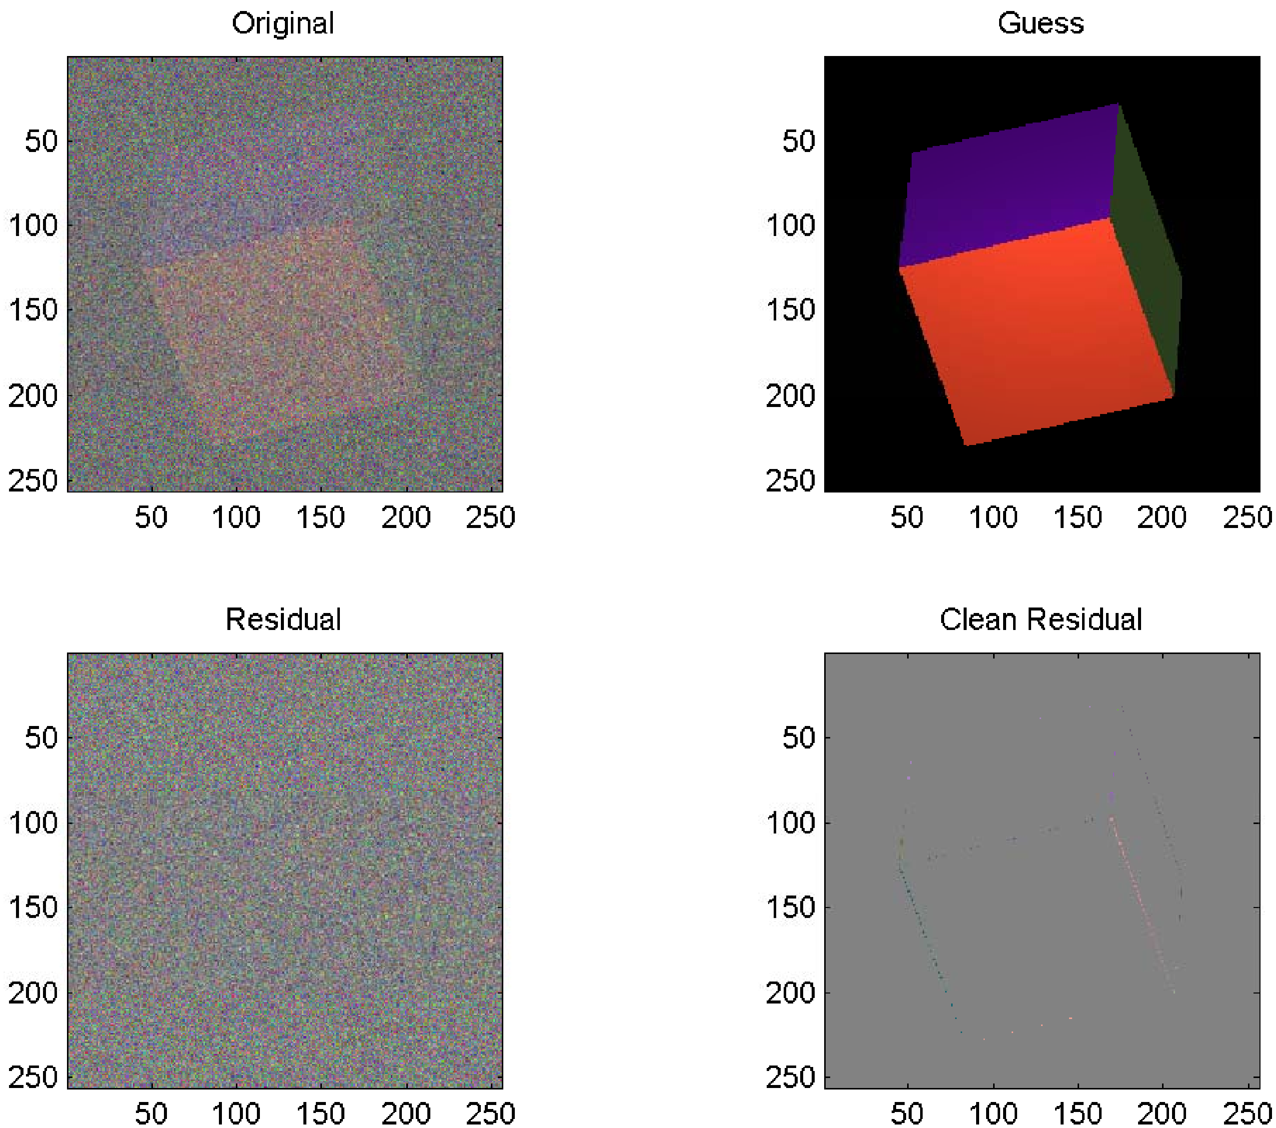

$s = 1/256$

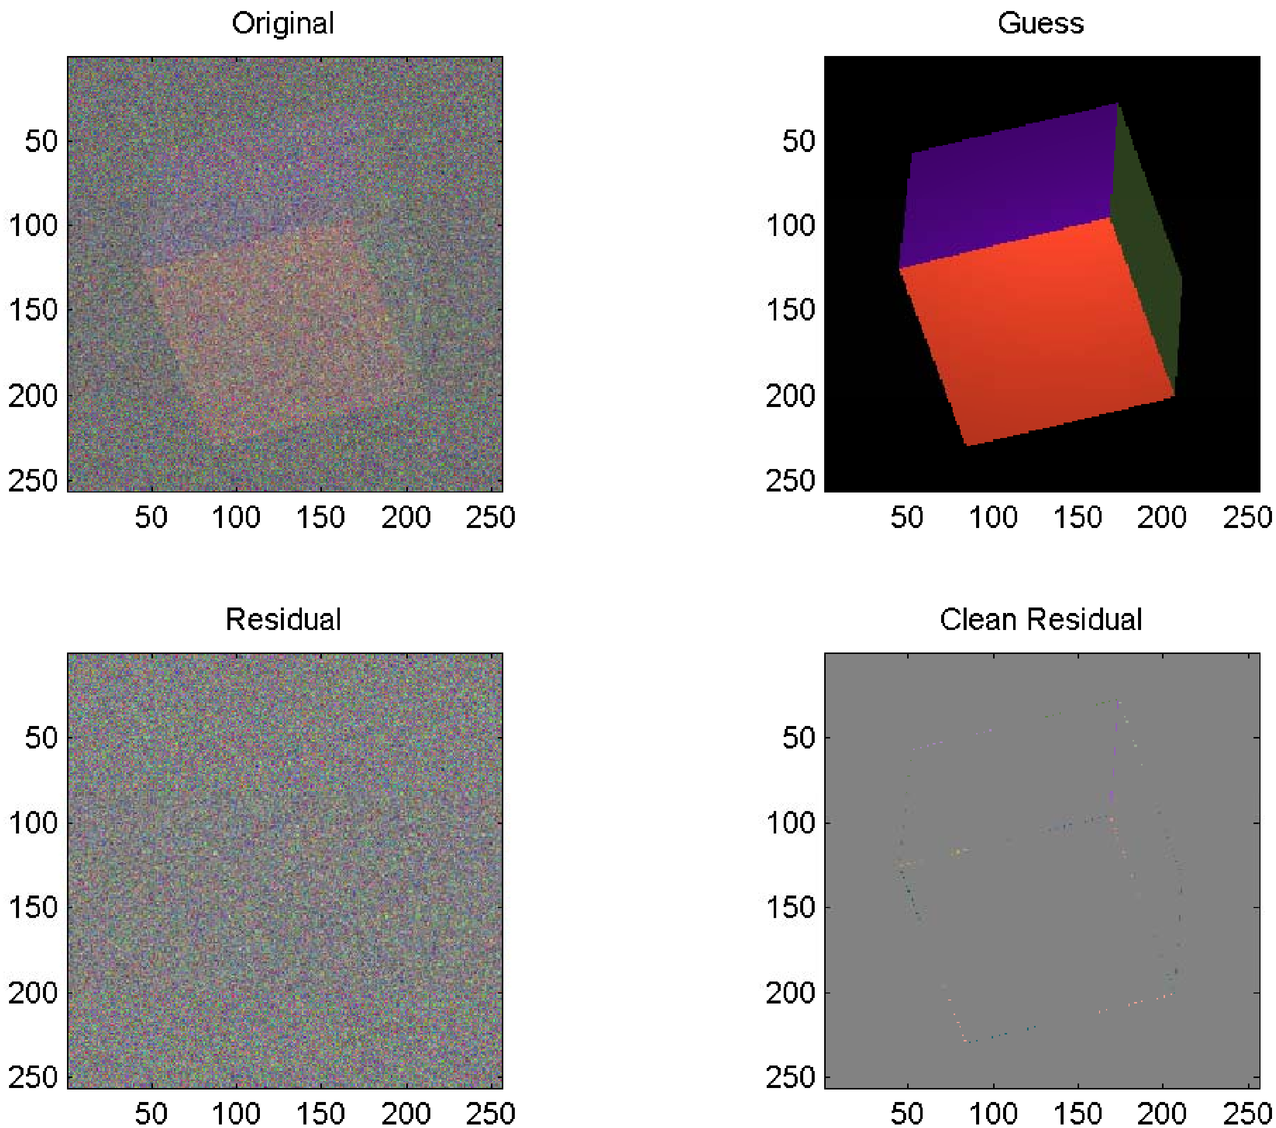

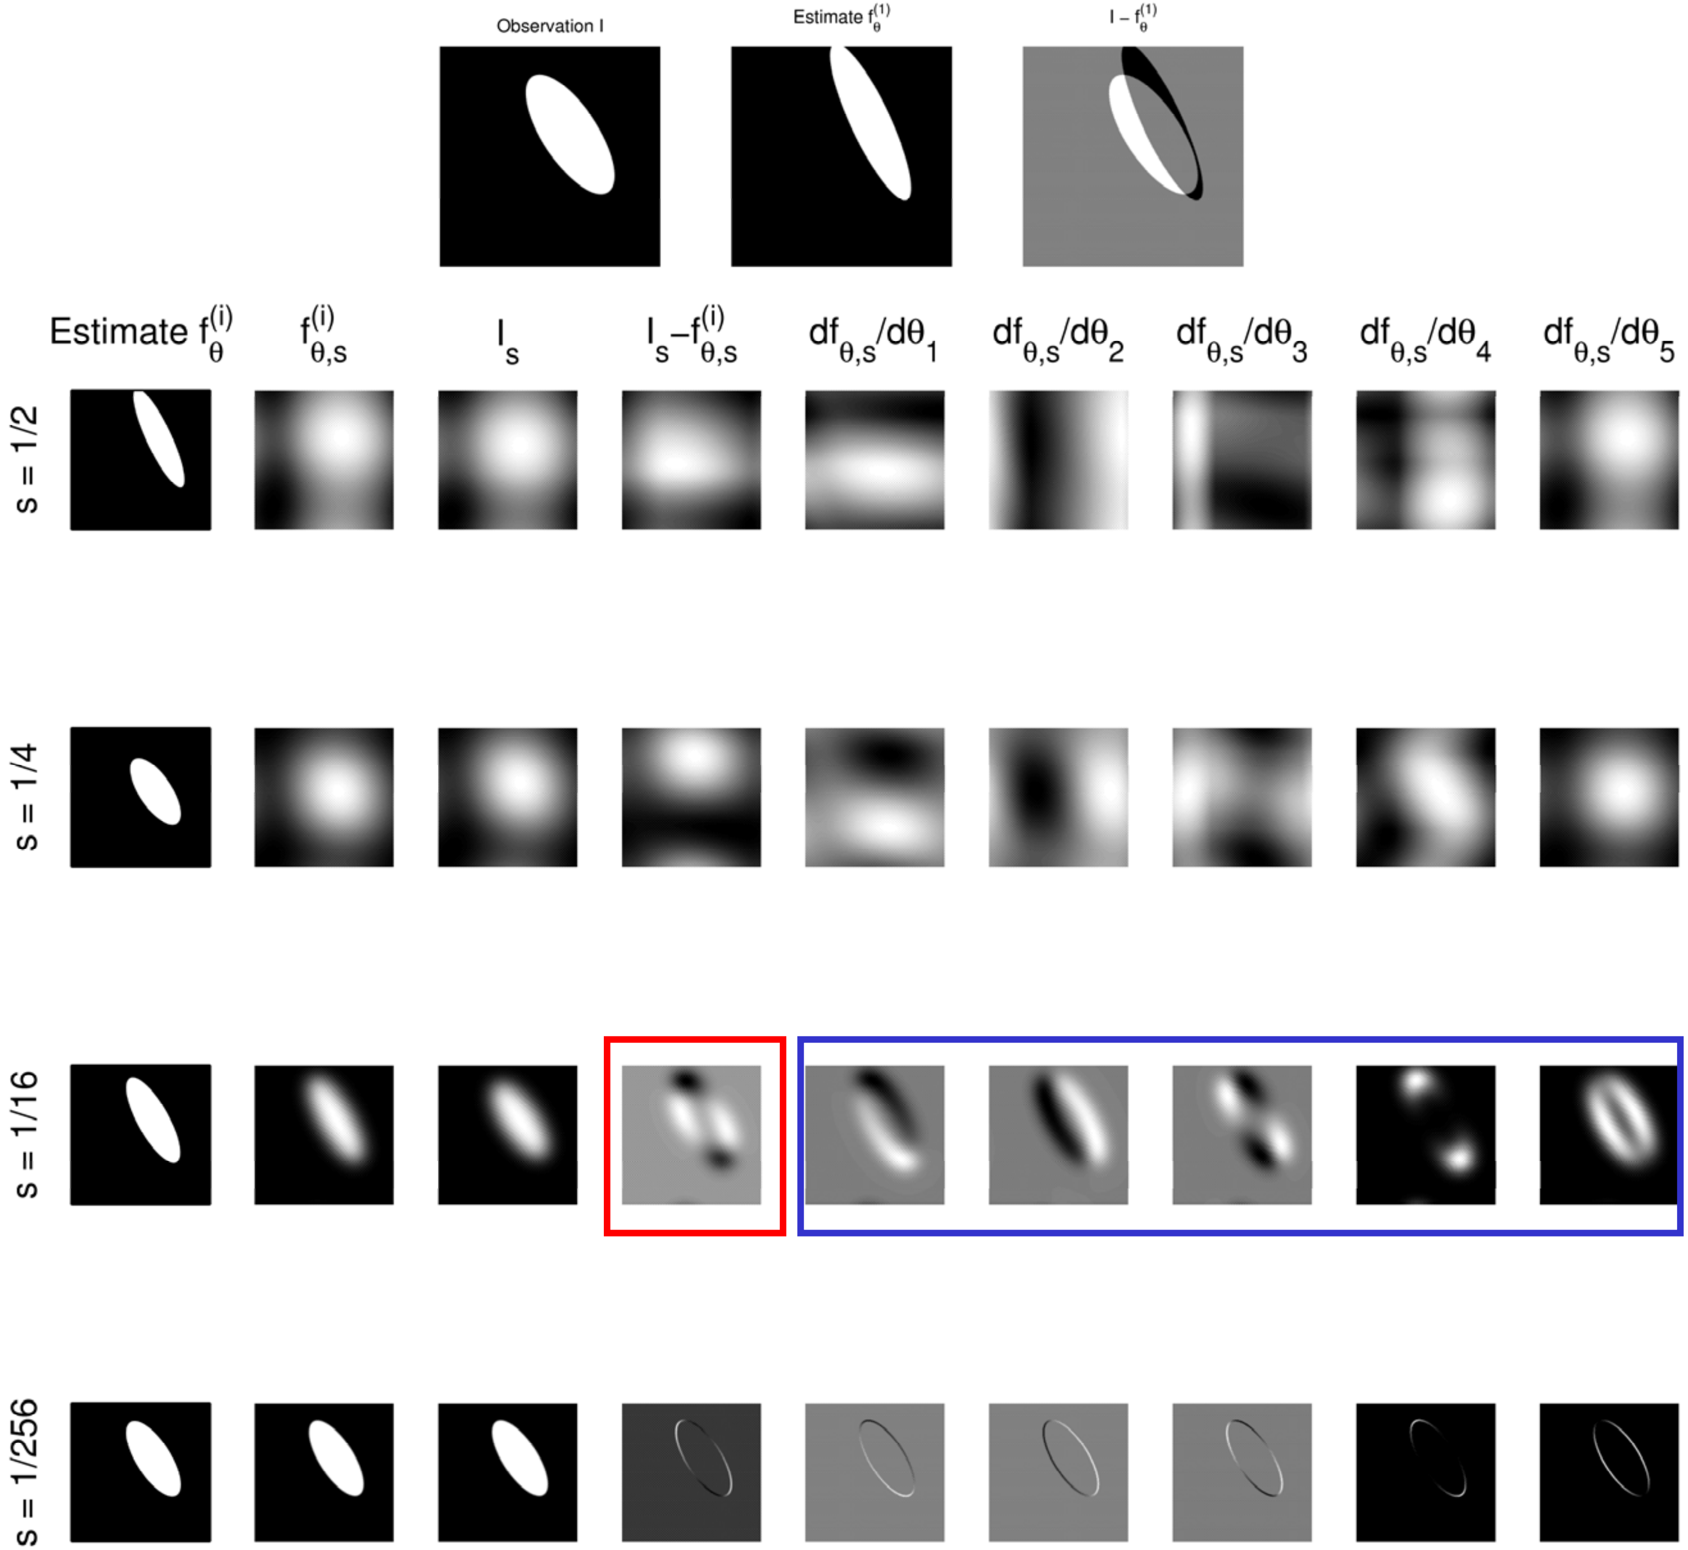

Observation I
Estimate $f_\theta^{(1)}$
$I - f_\theta^{(1)}$
Estimate $f_\theta^{(i)}$
$f_{\theta,s}^{(i)}$
I_s
$I_s - f_{\theta,s}^{(i)}$
$df_{\theta,s}/d\theta_1$
$df_{\theta,s}/d\theta_2$
$df_{\theta,s}/d\theta_3$
$df_{\theta,s}/d\theta_4$
$df_{\theta,s}/d\theta_5$
s = 1/2
s = 1/4
s = 1/16
s = 1/256

Overview

- Motivating application: parameter estimation
- Non-differentiability from edge migration
- Parameter estimation (revisited)
- Non-differentiability from edge occlusion
- Manifolds in Compressive Sensing

Occlusion-based Non-differentiability

- Sudden appearance/disappearance of edges
- Tangent spaces changing dimension
  - different “left”, “right” tangents
- Occurs at every scale

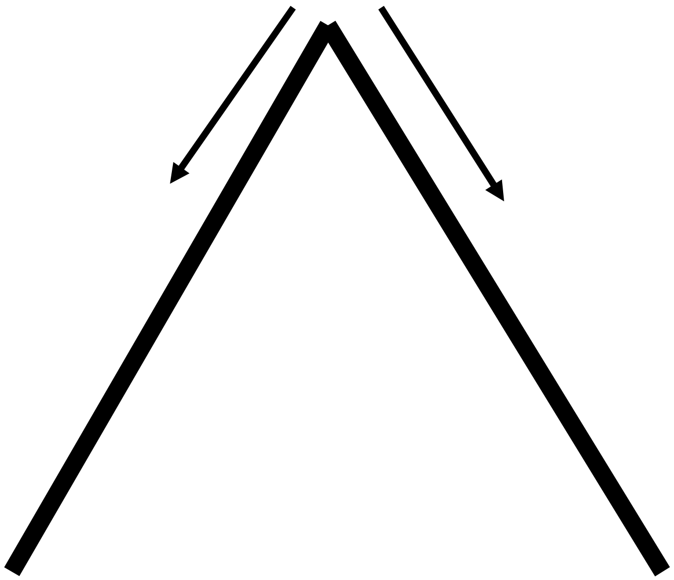

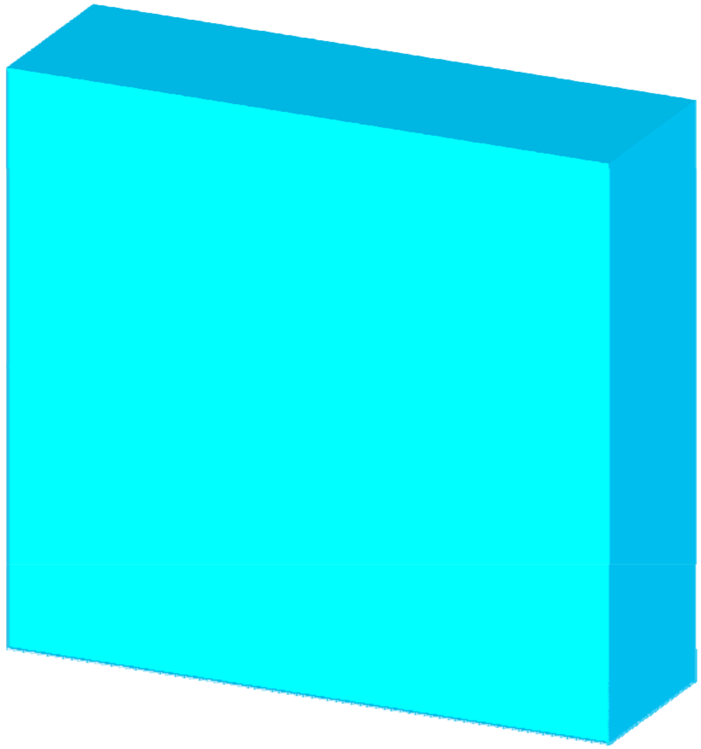

θ: *pitch*, roll, yaw

Occlusion-based Non-differentiability

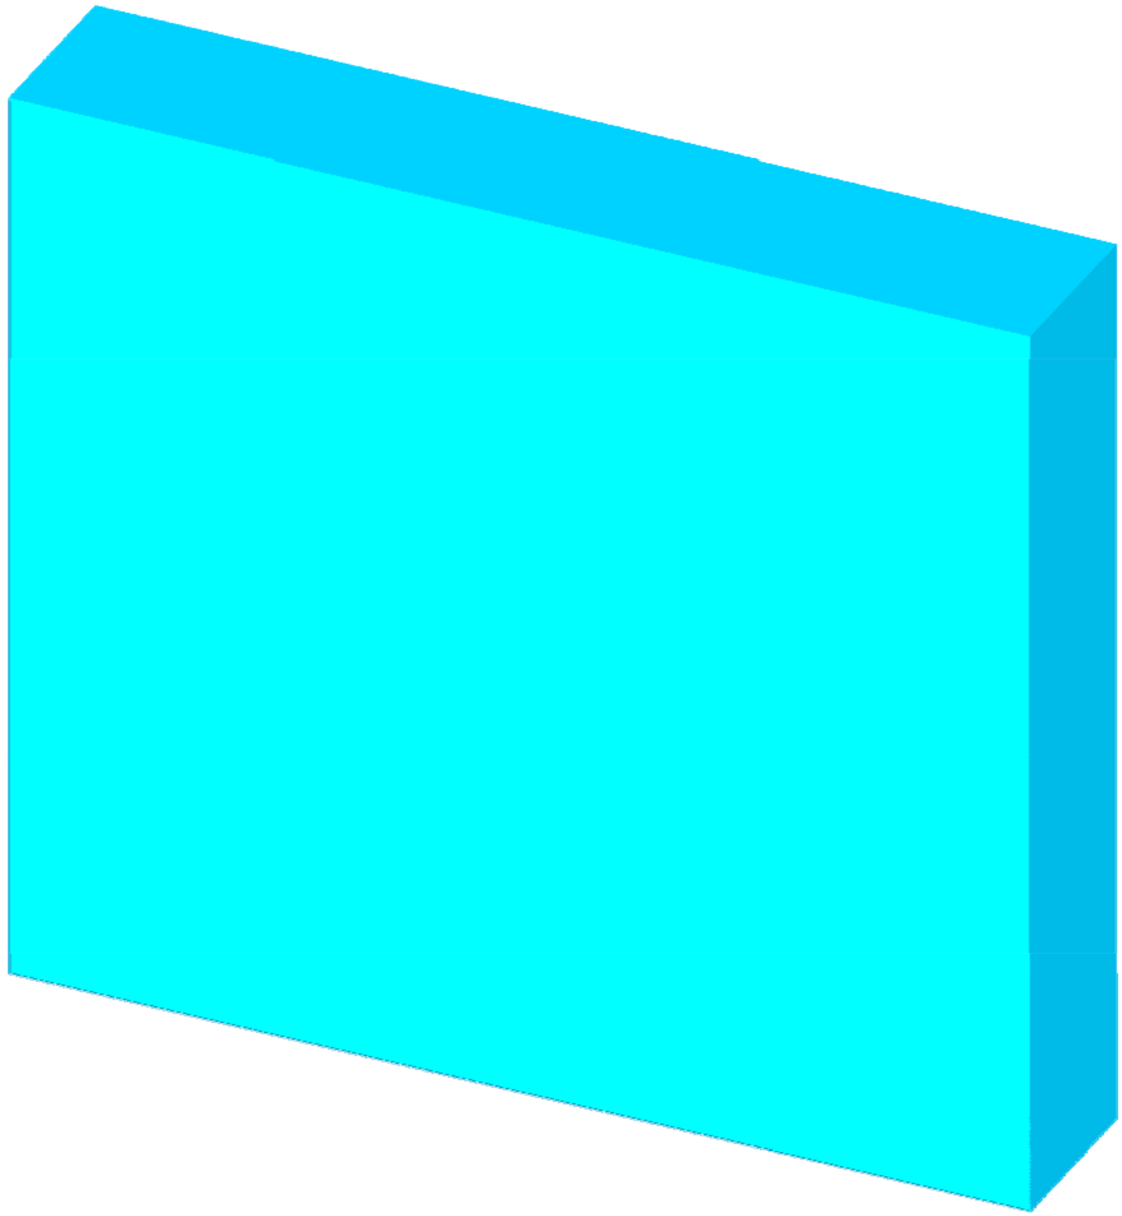

Occlusion-based Non-differentiability

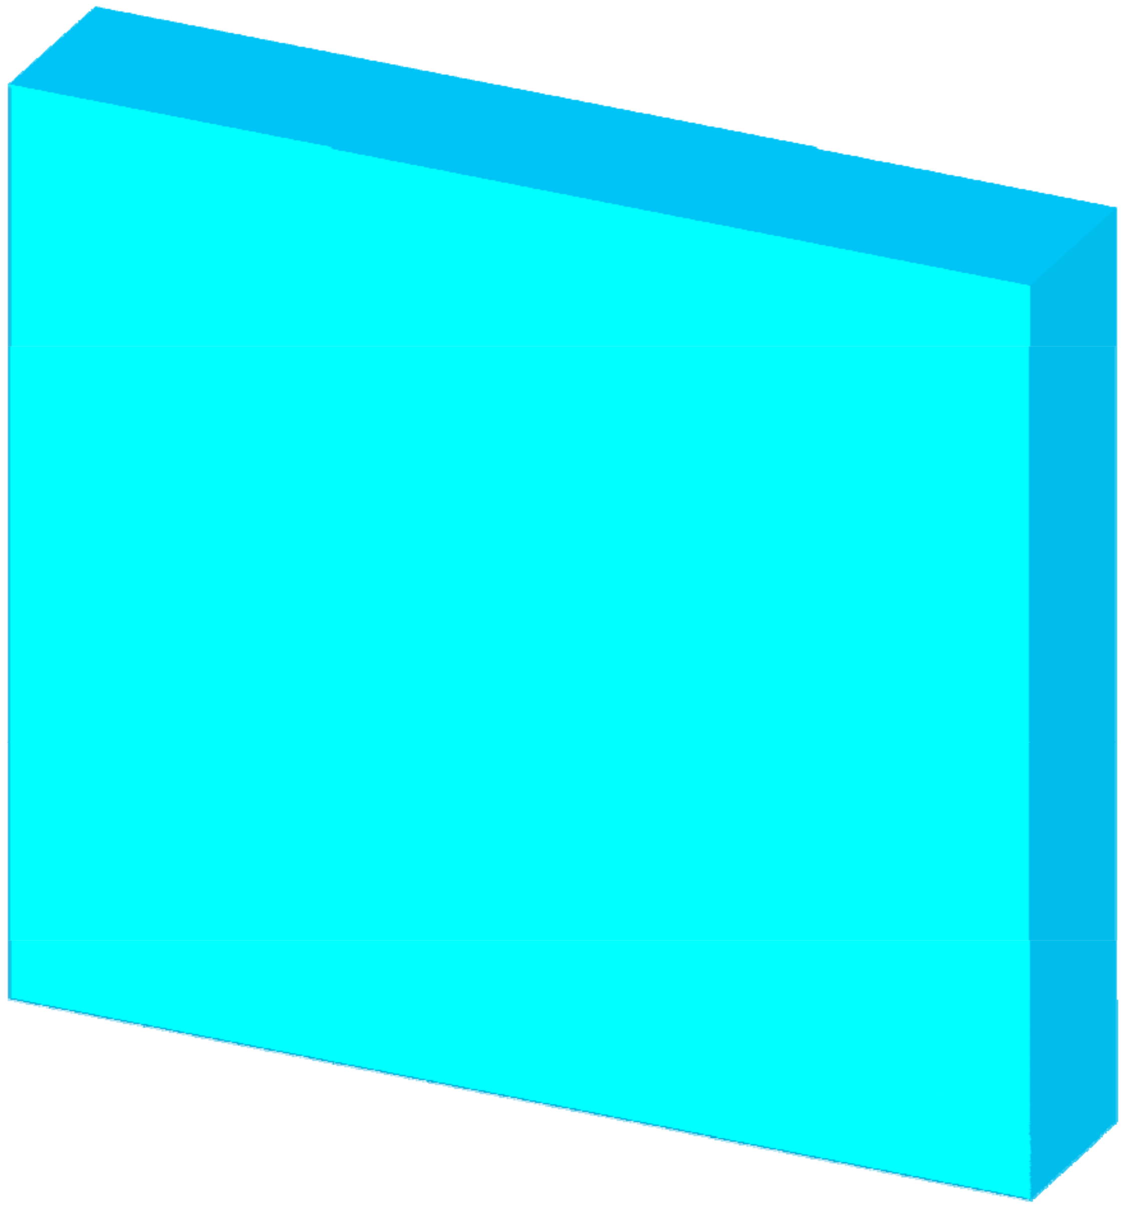

Occlusion-based Non-differentiability

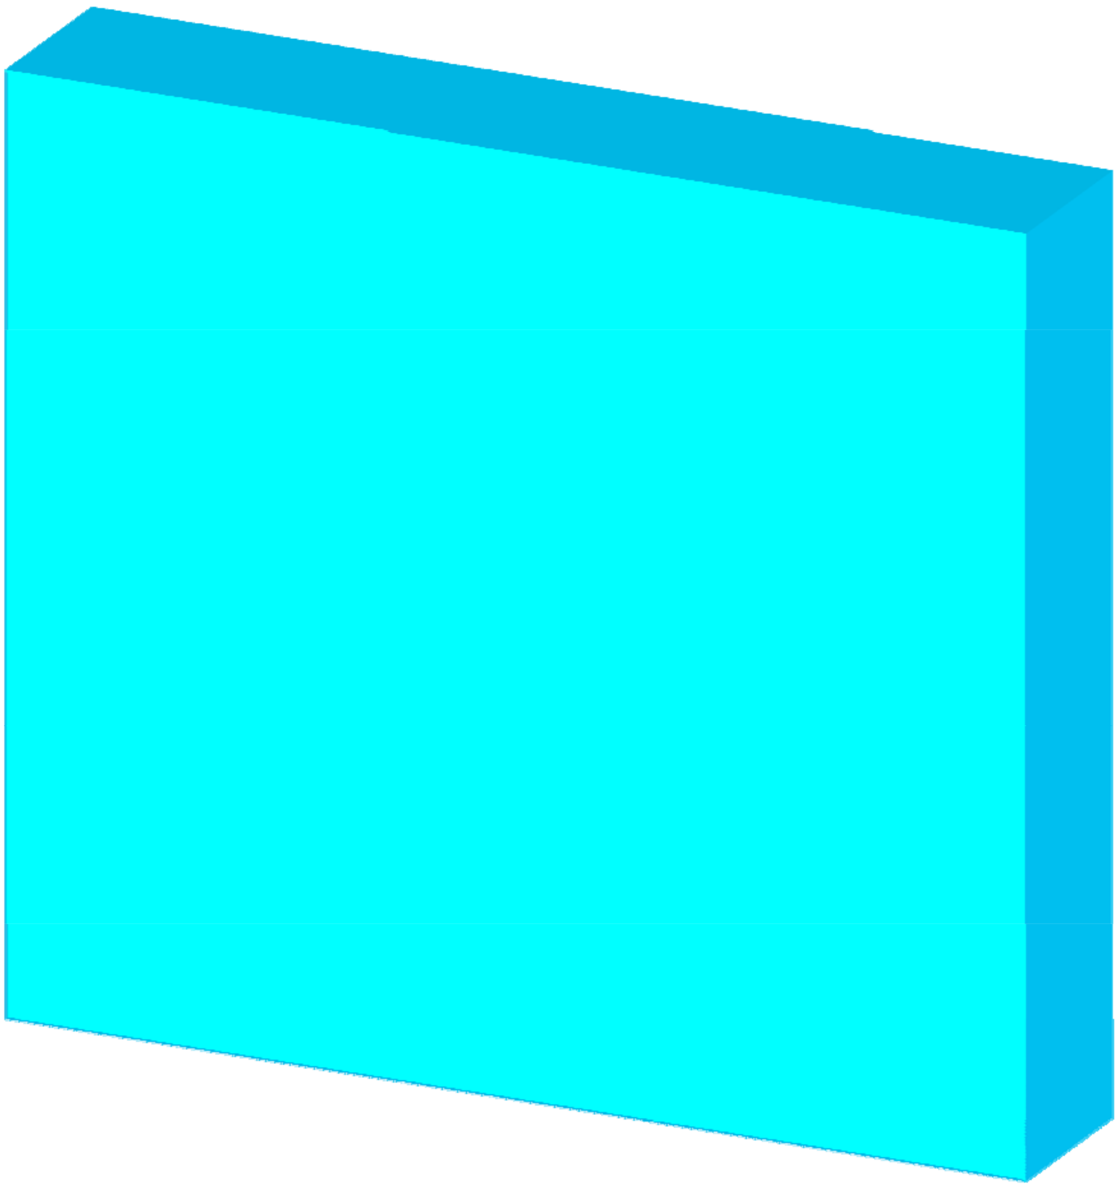

Occlusion-based Non-differentiability

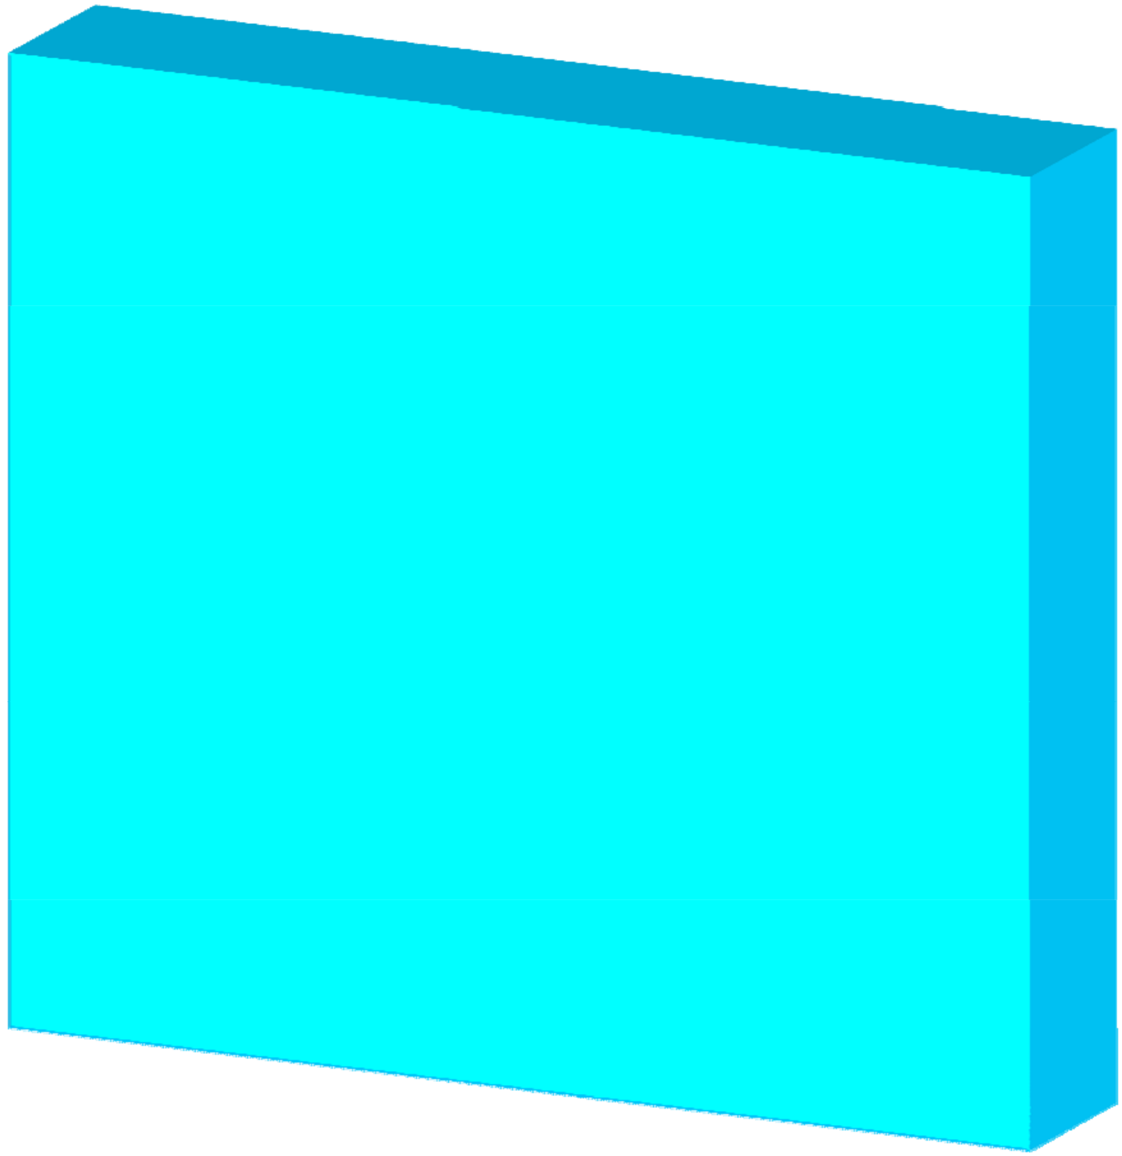

Occlusion-based Non-differentiability

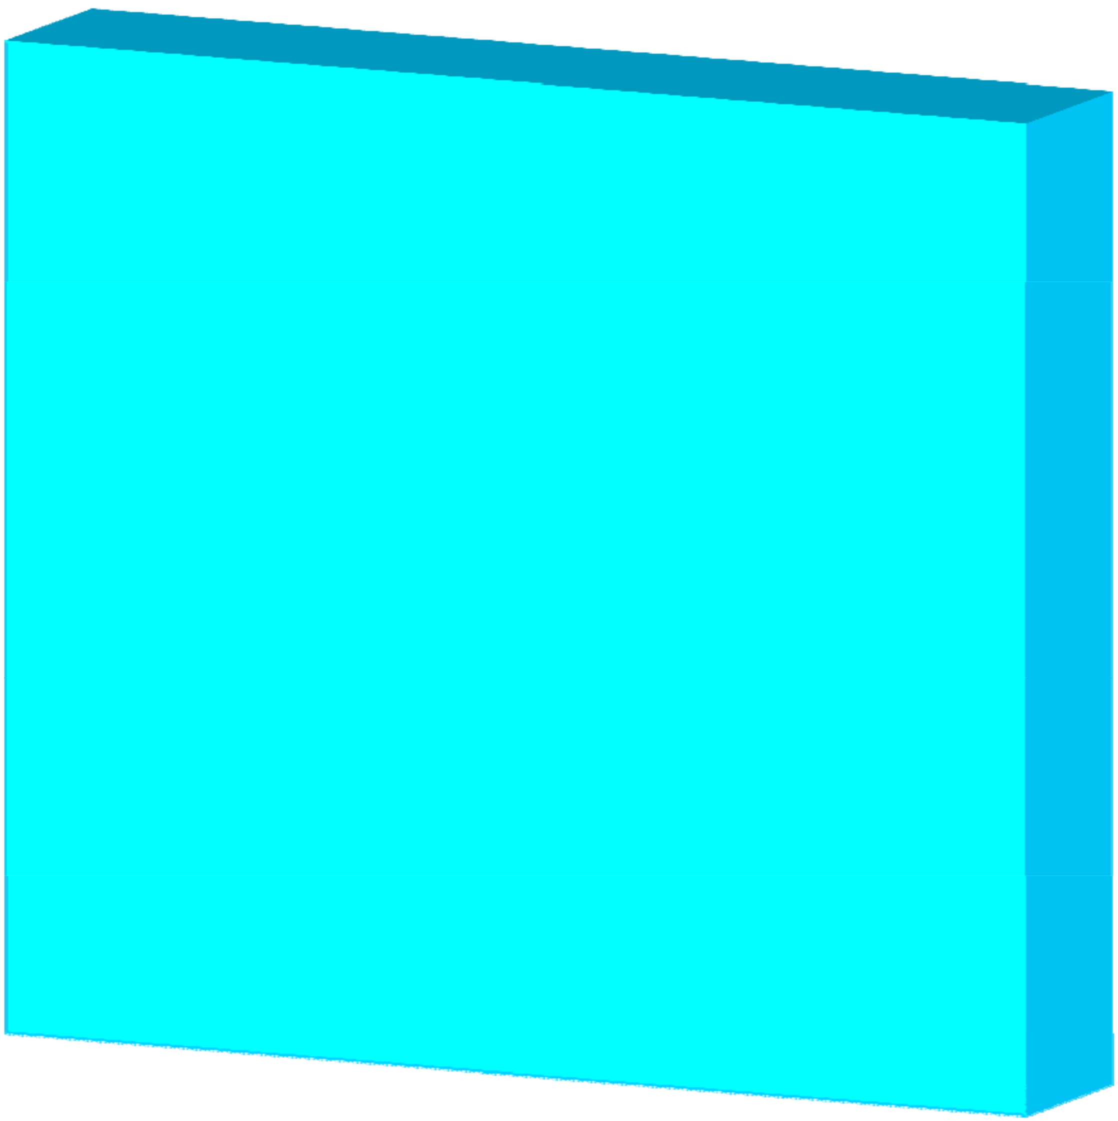

Occlusion-based Non-differentiability

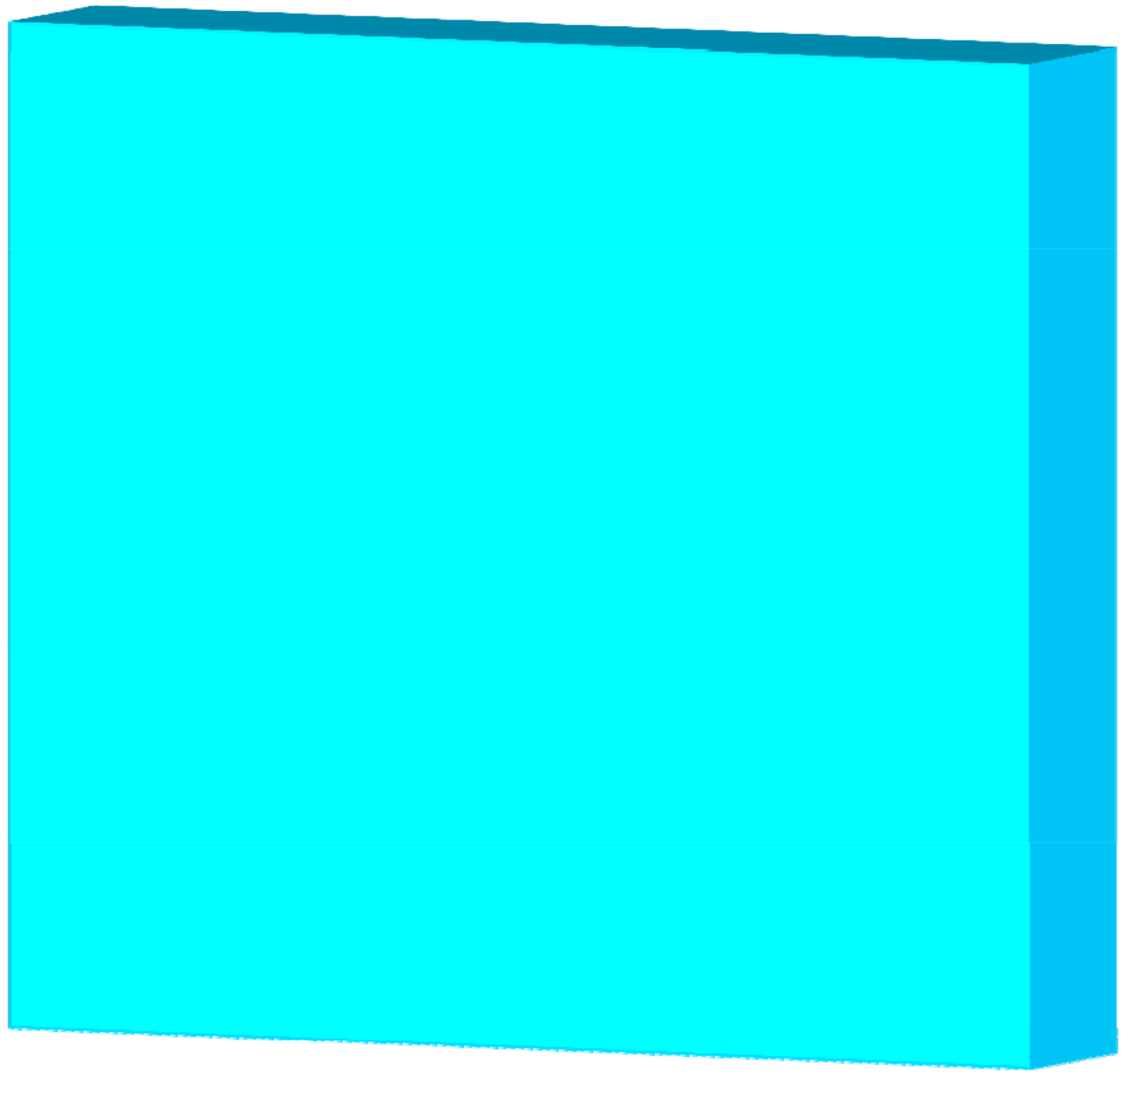

Occlusion-based Non-differentiability

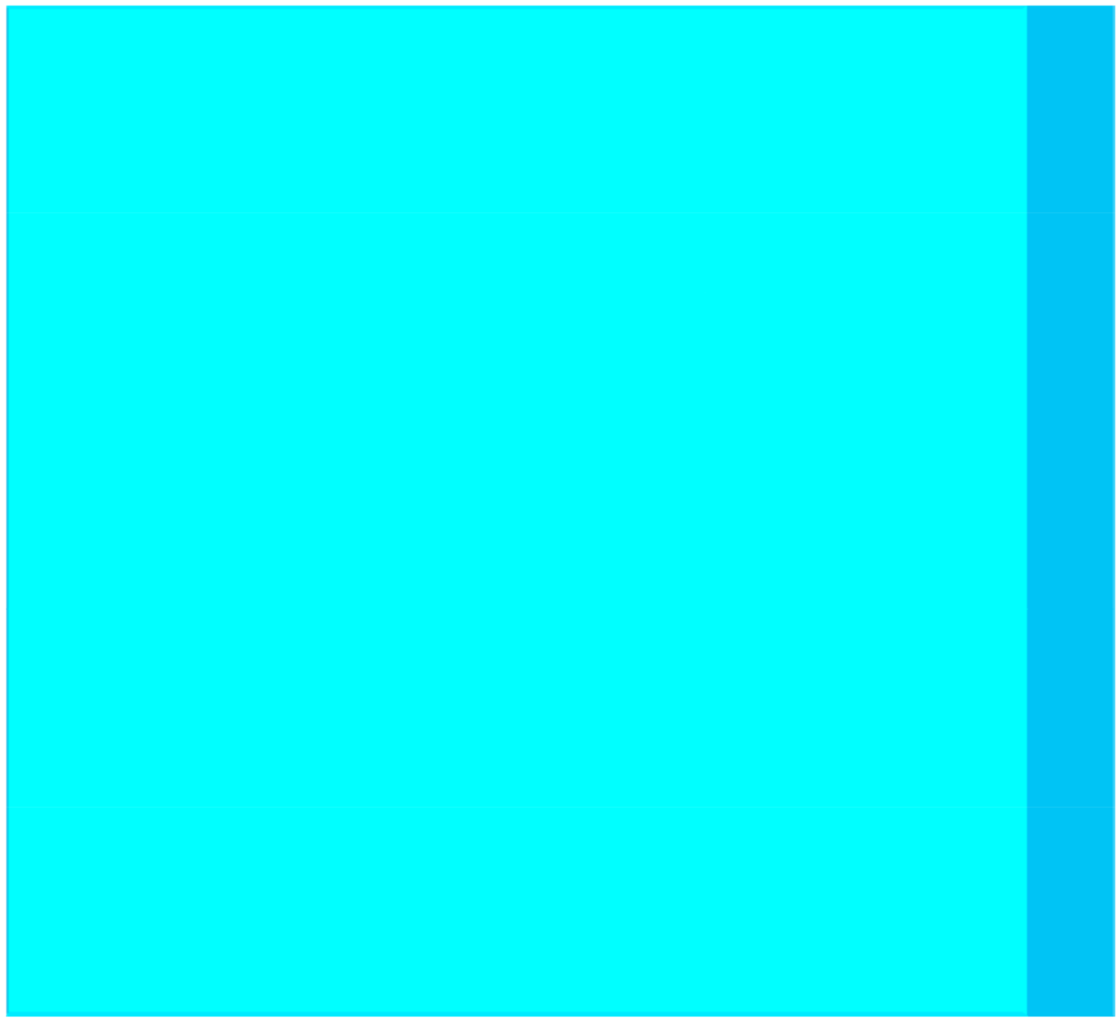

Occlusion-based Non-differentiability

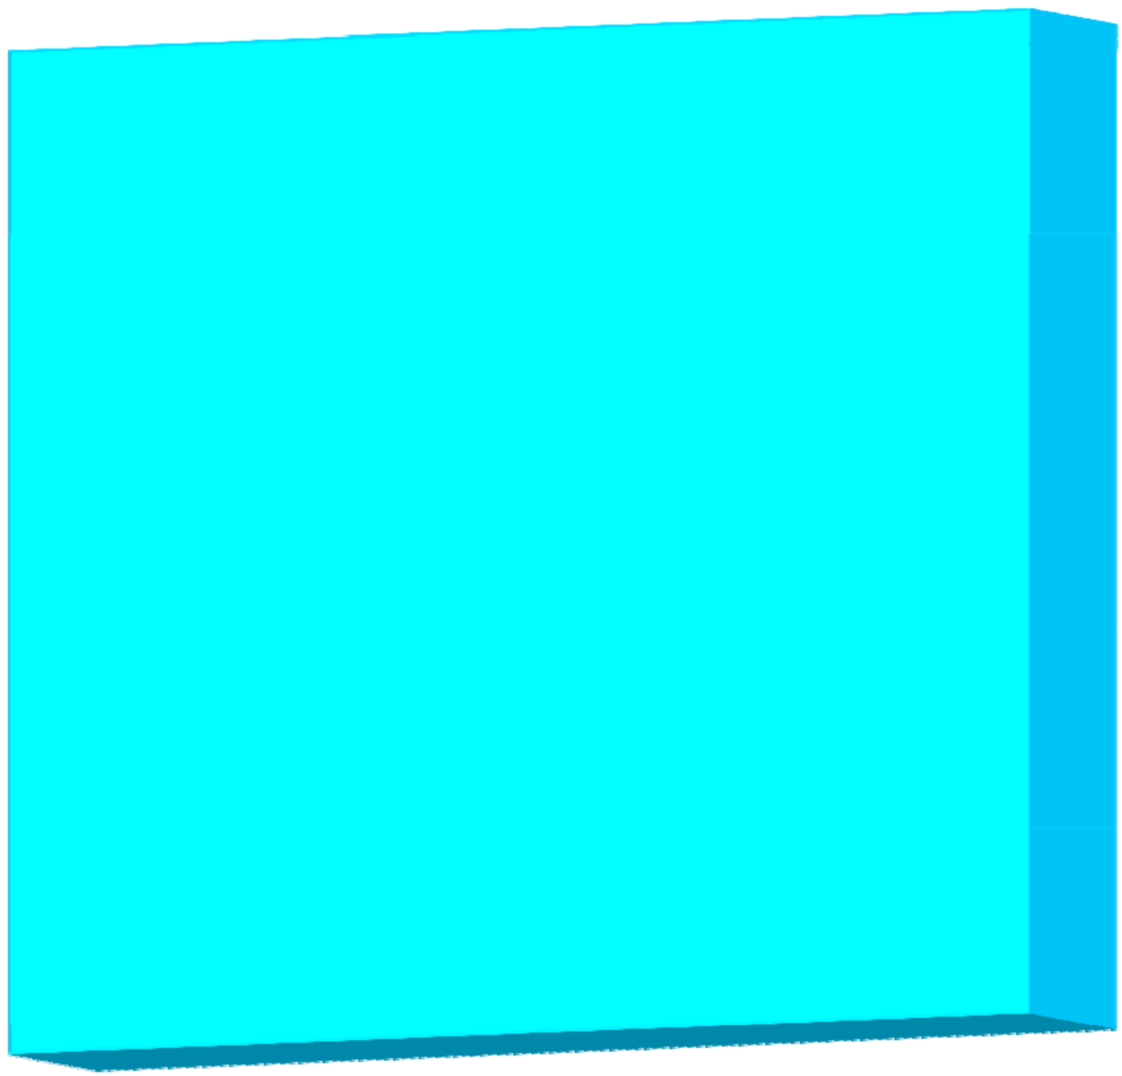

Occlusion-based Non-differentiability

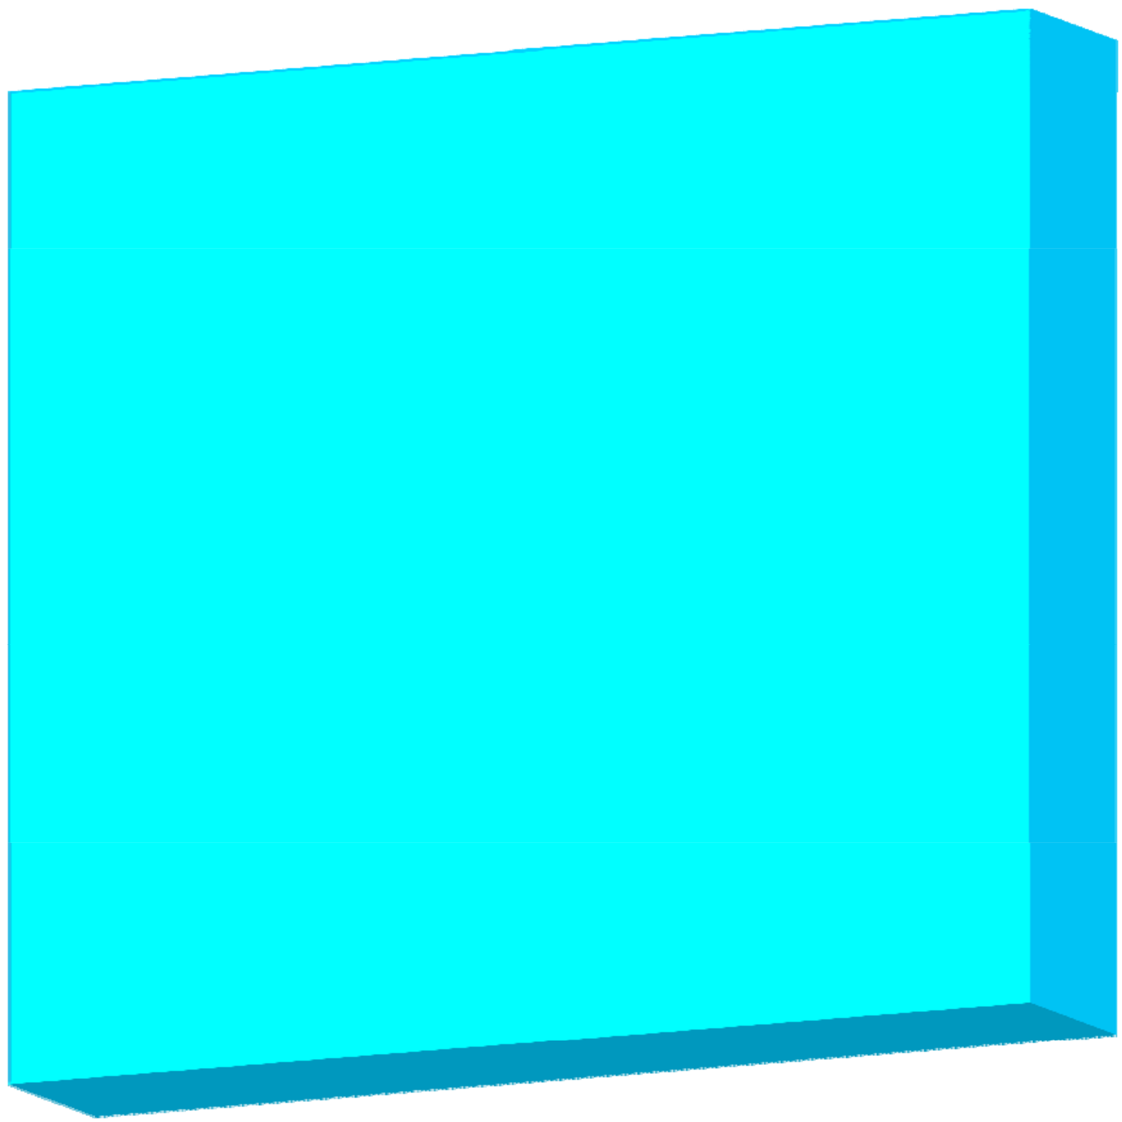

Occlusion-based Non-differentiability

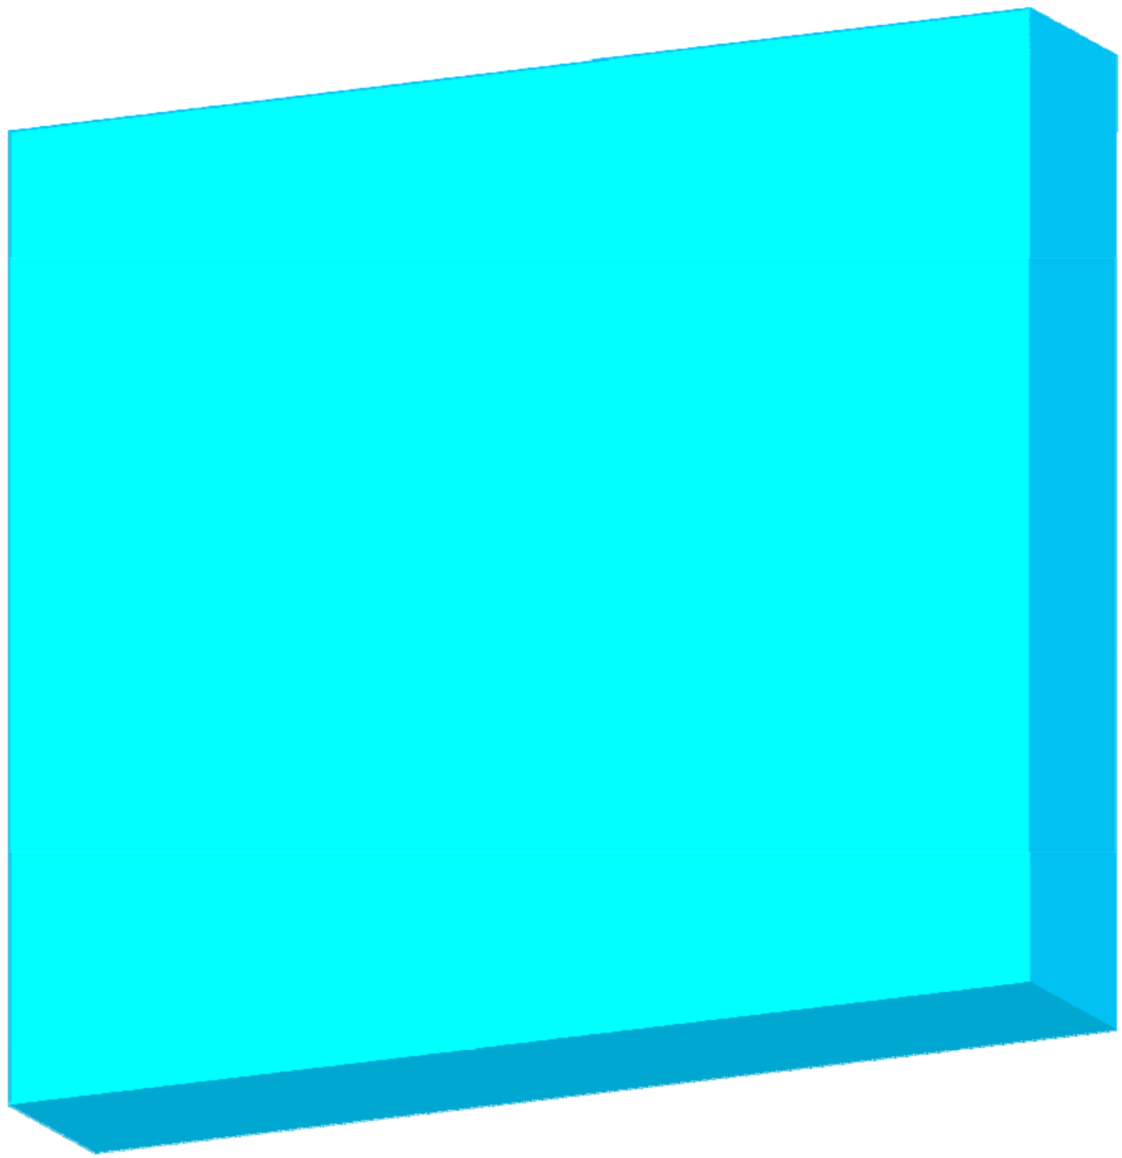

Occlusion-based Non-differentiability

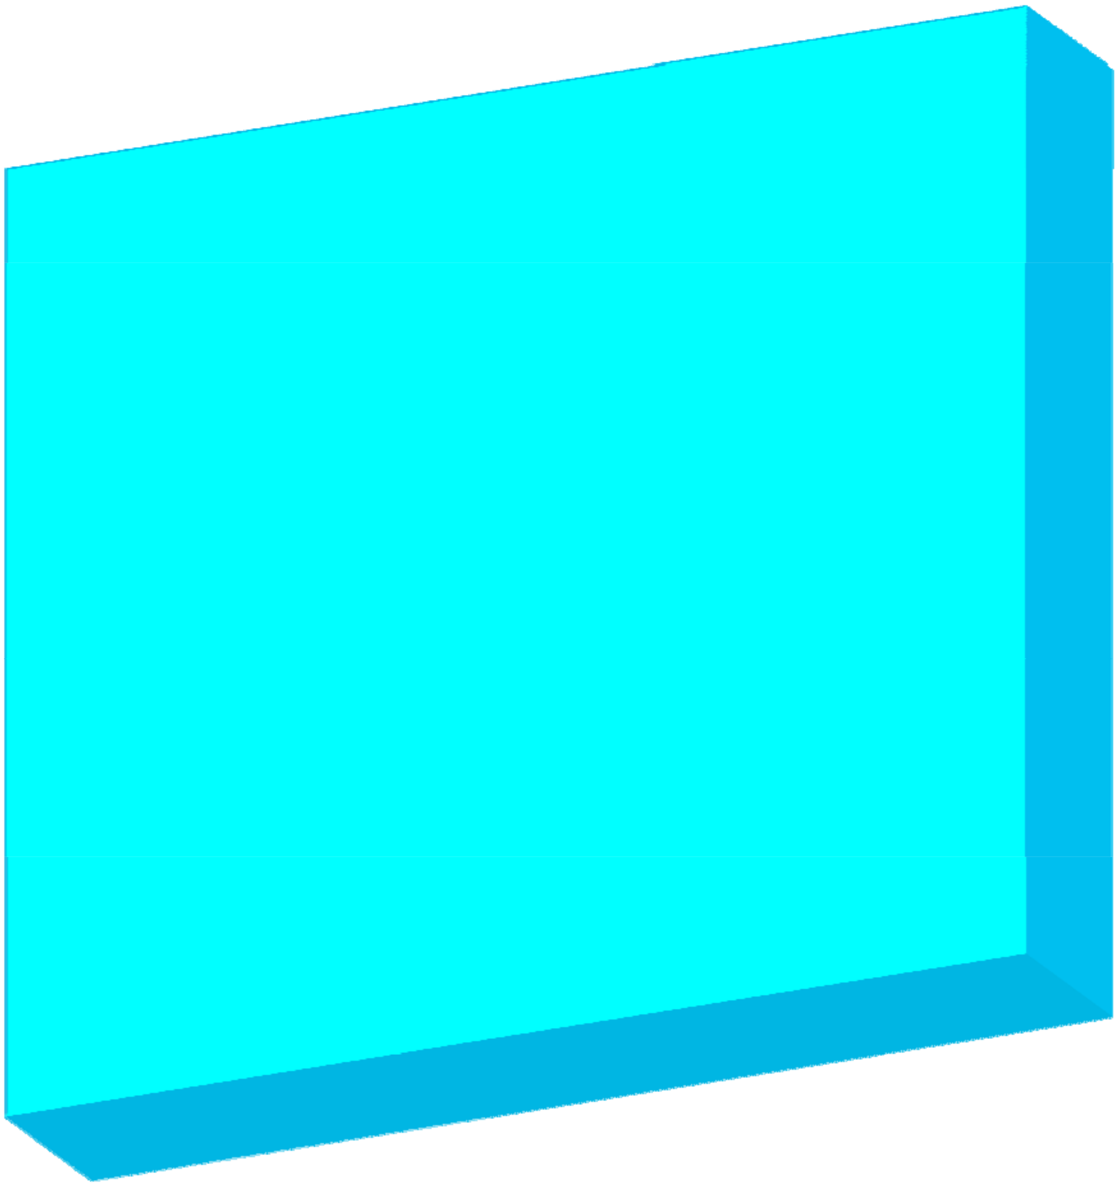

Five Relevant Tangent Vectors

“head-on” view

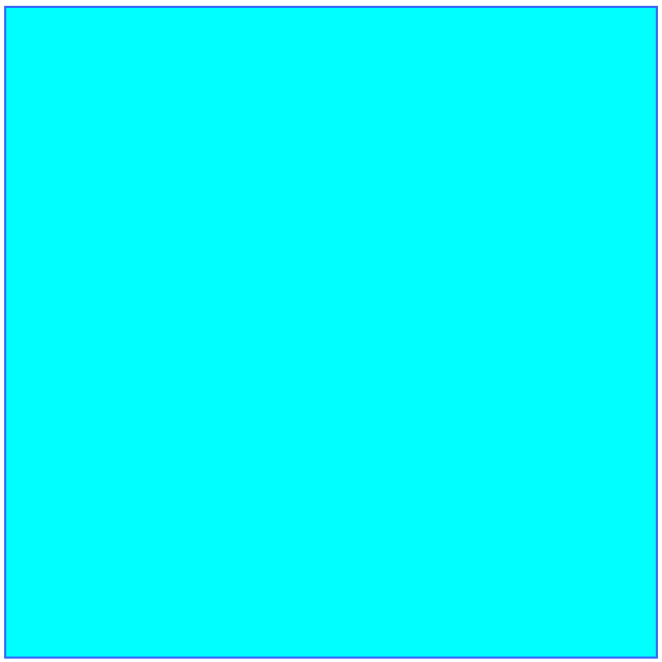

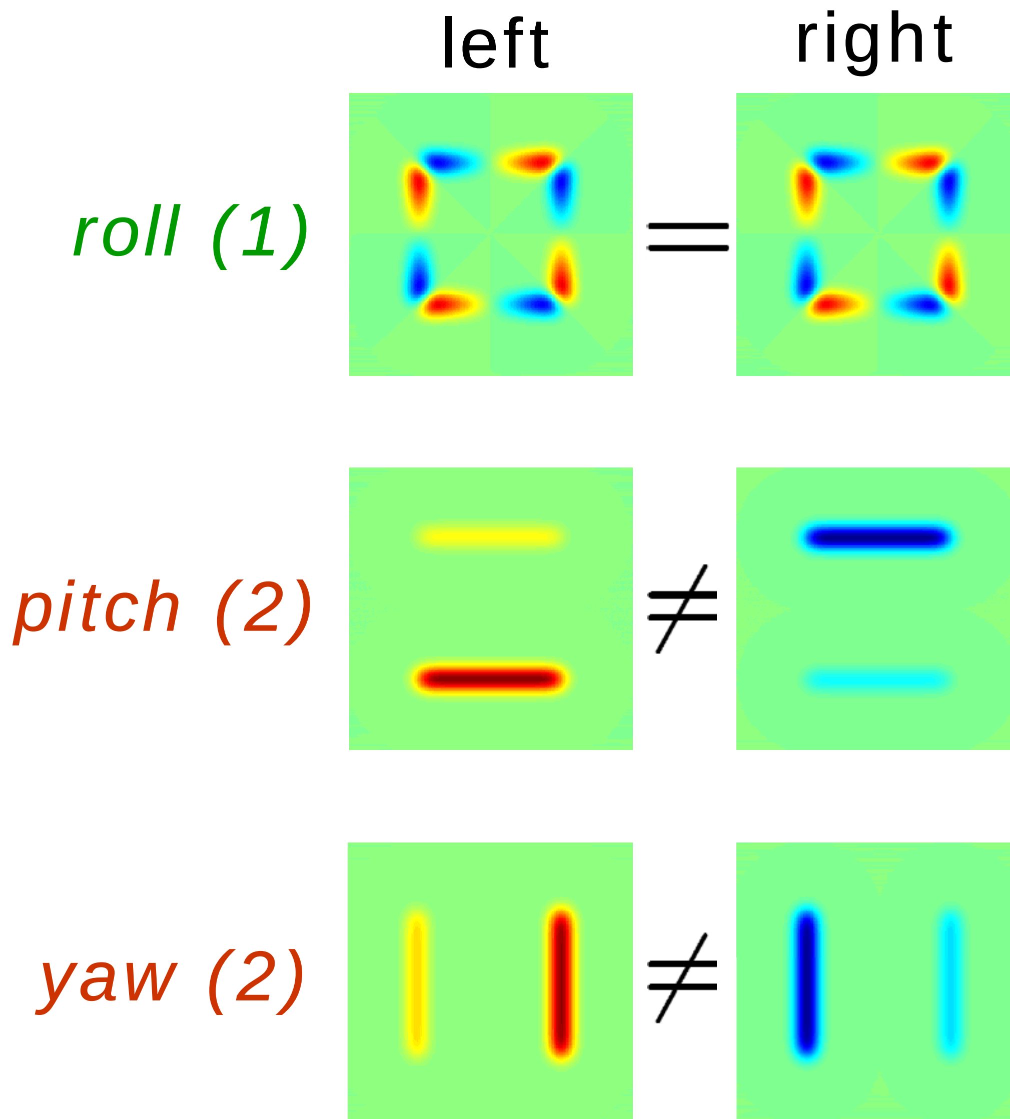

- Can explicitly consider such points in parameter estimation

Overview

- Motivating application: parameter estimation
- Non-differentiability from edge migration
- Parameter estimation (revisited)
- Non-differentiability from edge occlusion
- Manifolds in Compressive Sensing

Compressive Sensing

- Signal x is K-*sparse* in basis/dictionary Ψ
- Collect linear measurements $y = \Phi x$
  - measurement operator Φ *incoherent* with elements from Ψ
  - not adapted to signal x – *random Φ will work*

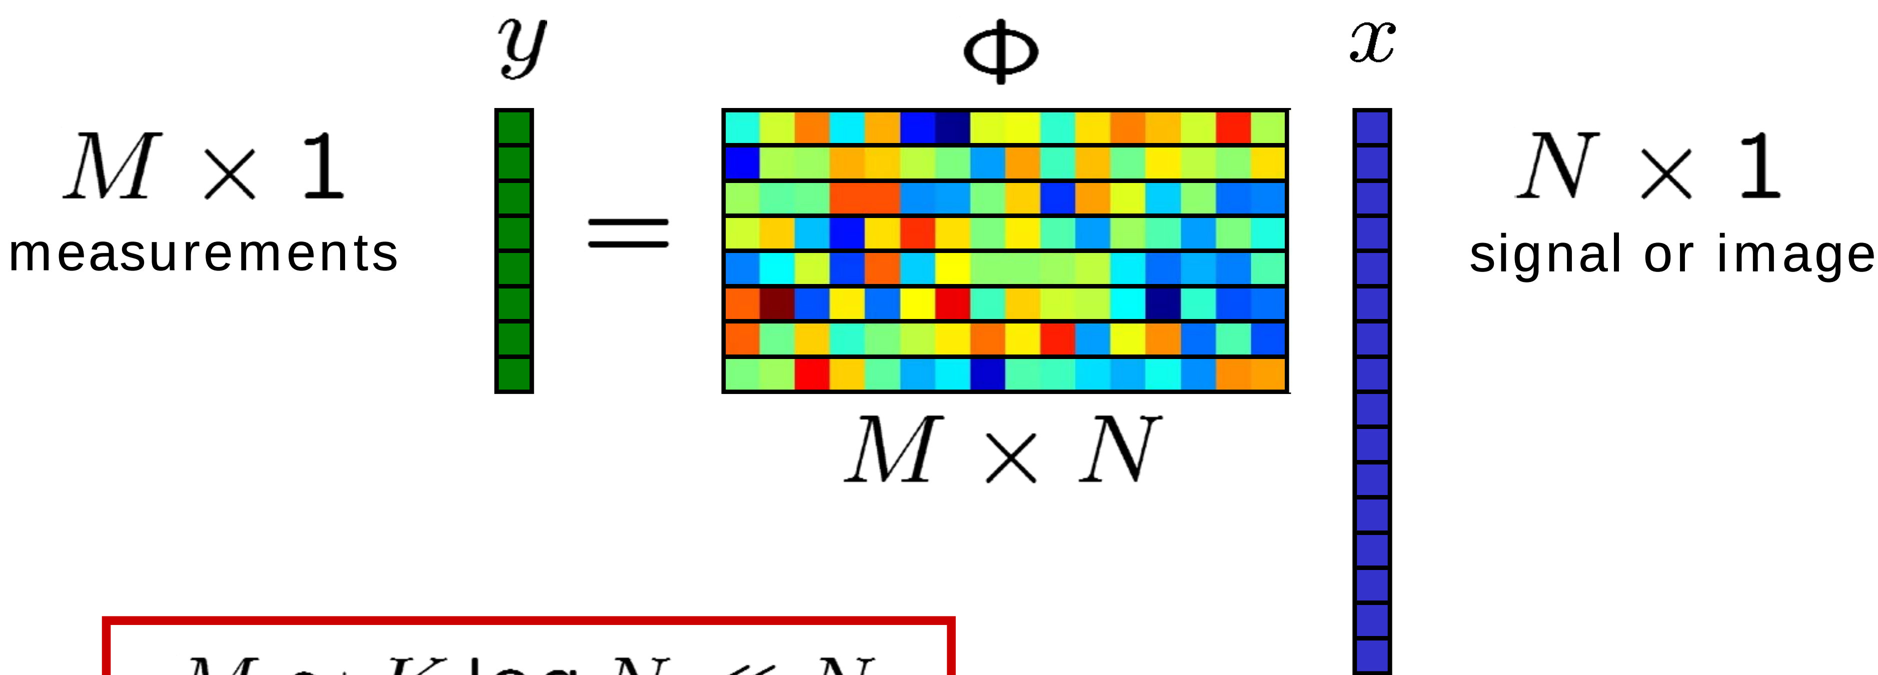

$$M \approx K \log N \ll N$$

[Candès, Romberg, Tao; Donoho]

"Single Pixel" CS Camera

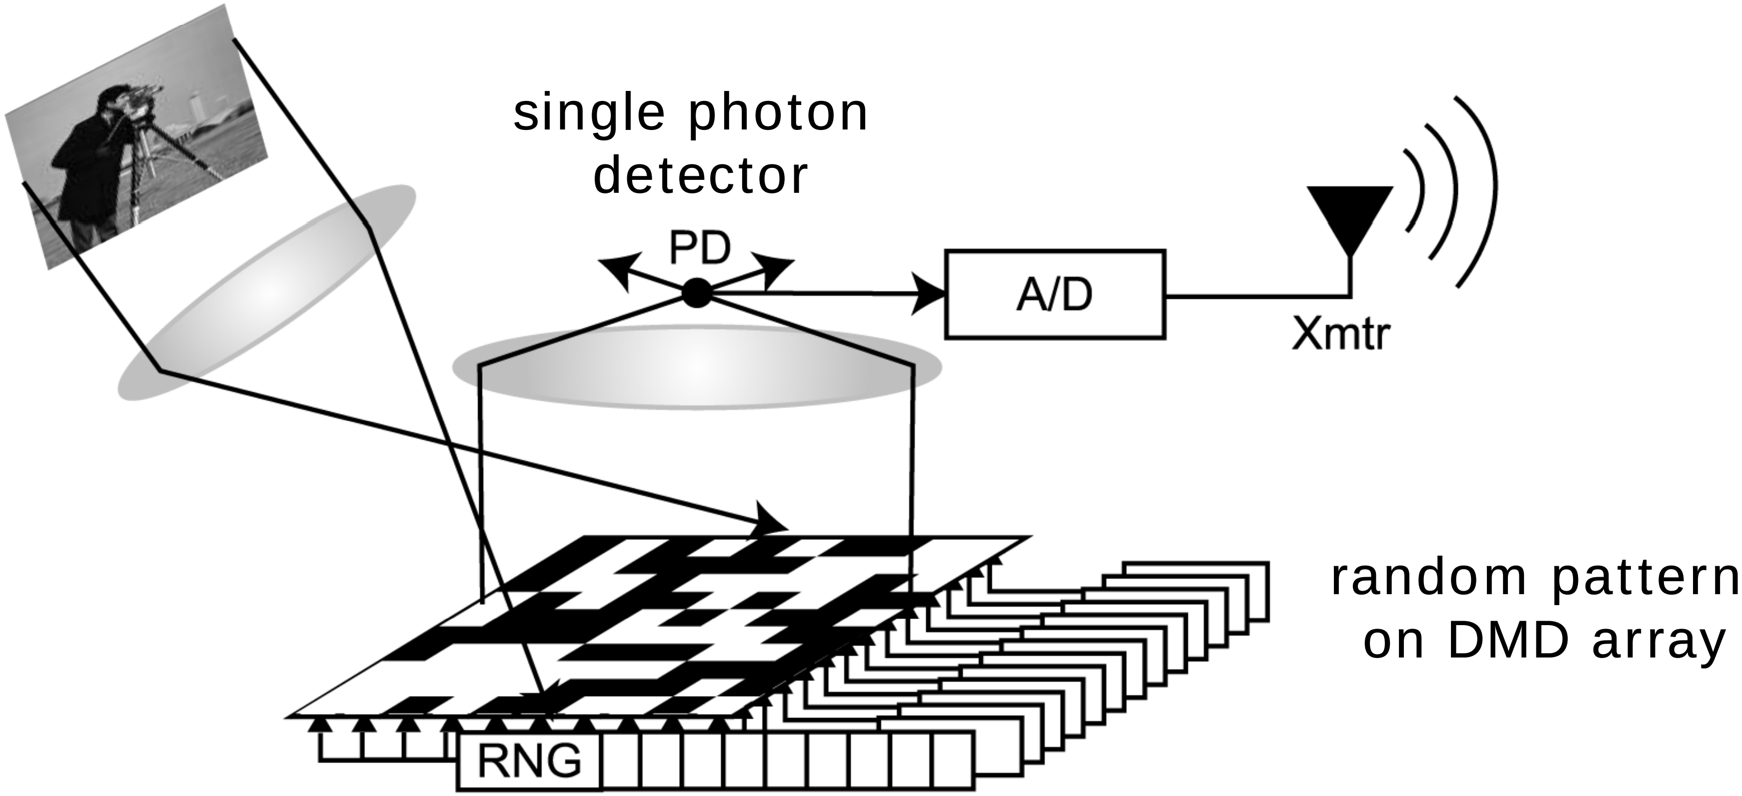

4096 pixels
1600 measurements
(40%)

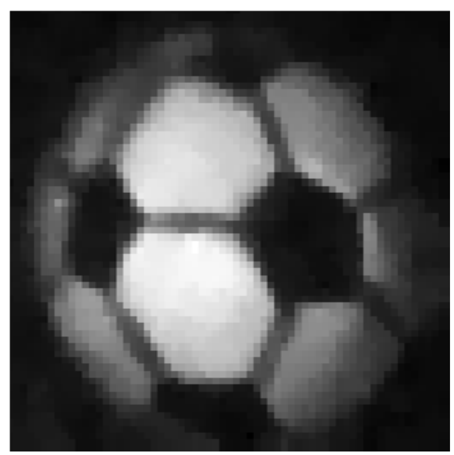

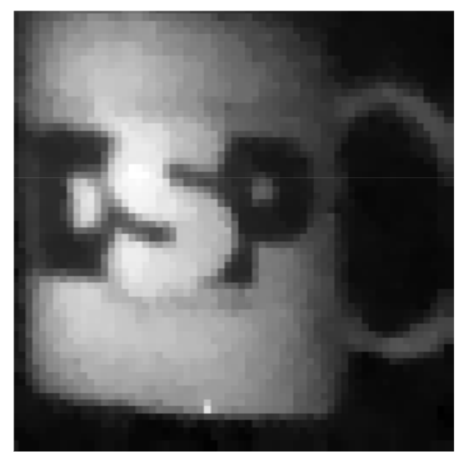

[with R. Baraniuk + Rice CS Team]

Why CS Works: Stable Embeddings

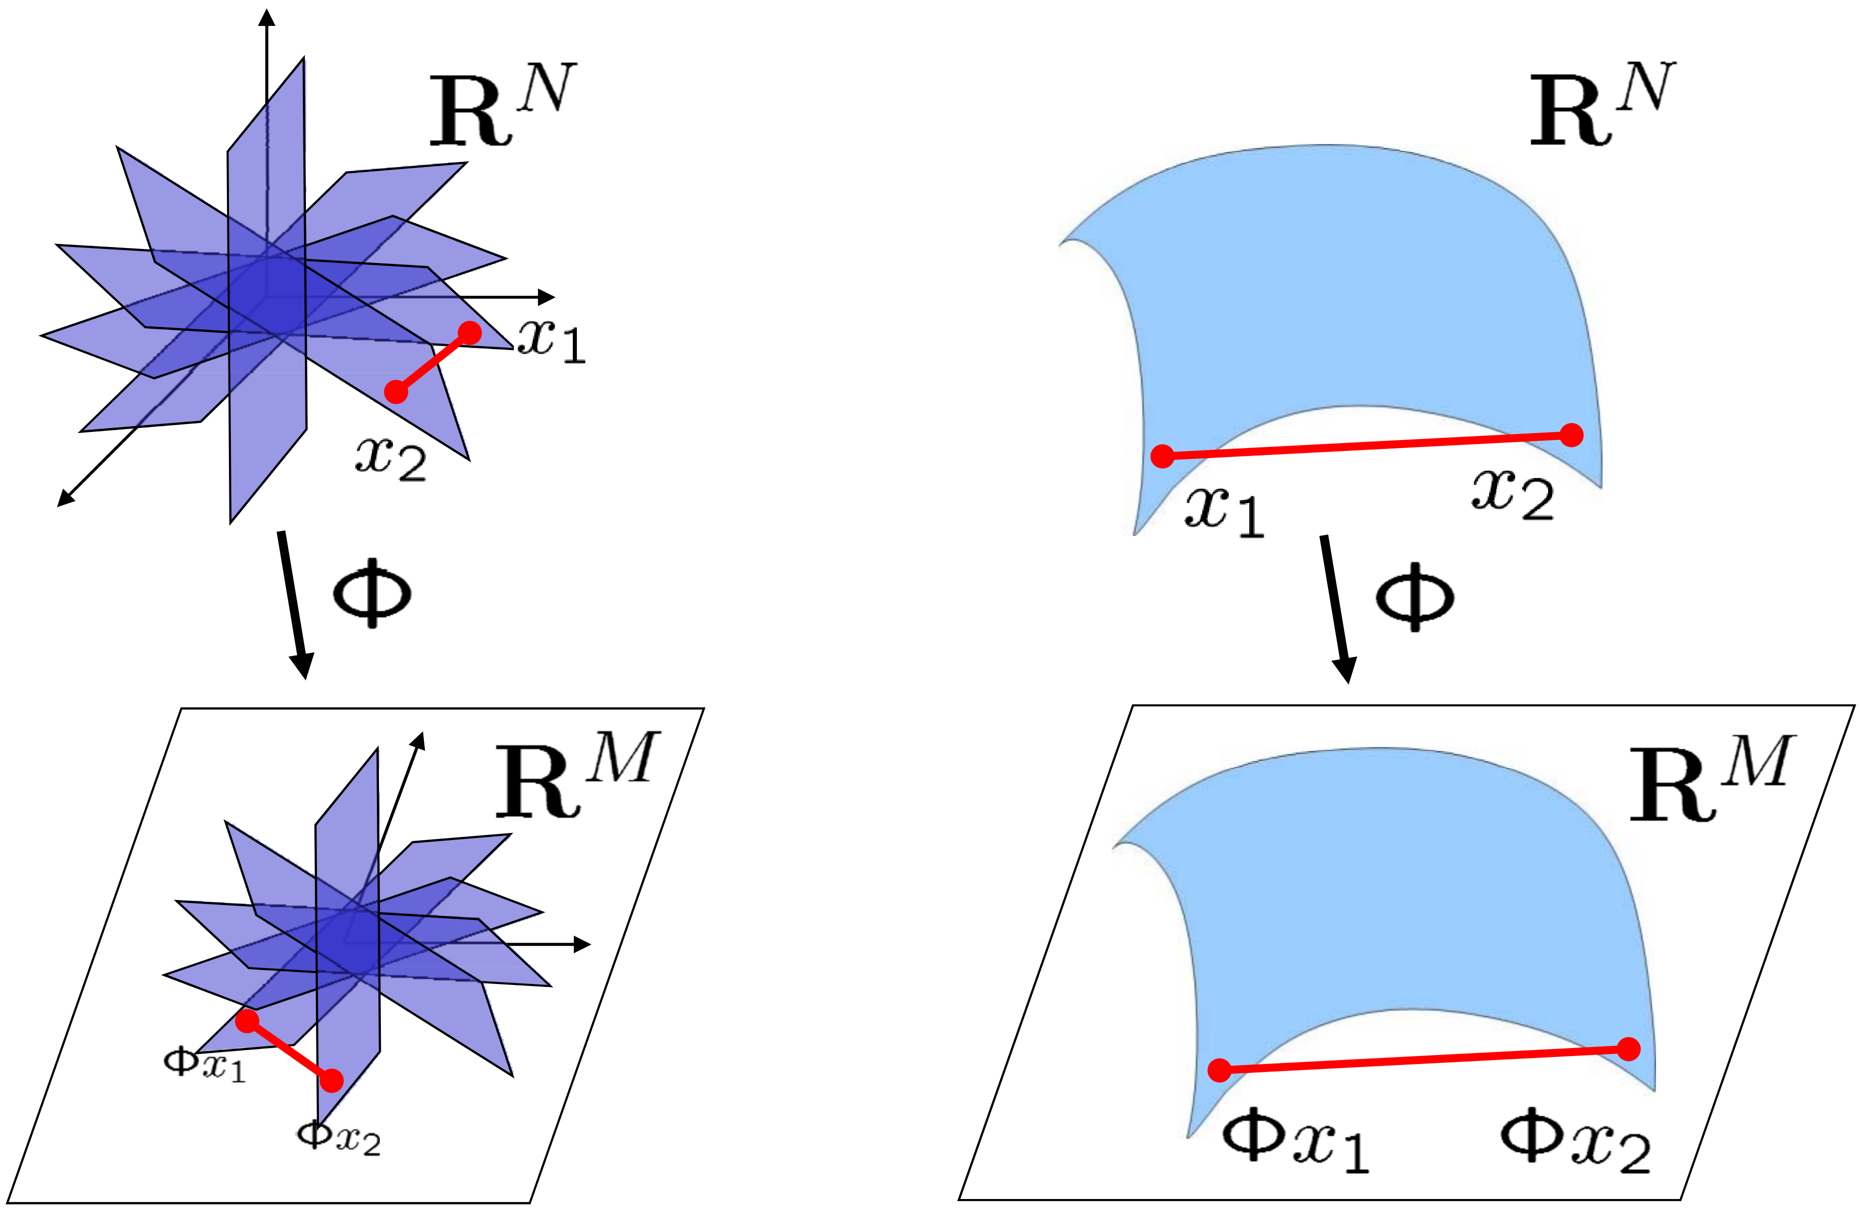

One Challenge: Non-Differentiability

- Many image manifolds are non-differentiable
  - no embedding guarantee
  - difficult to navigate

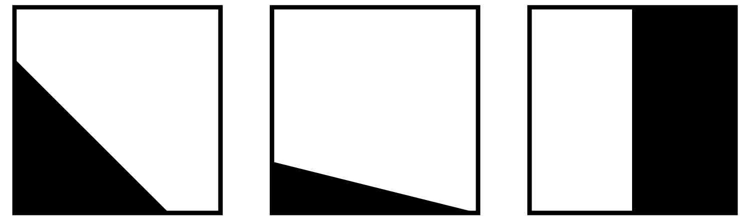

- Solution: *multiscale random projections*

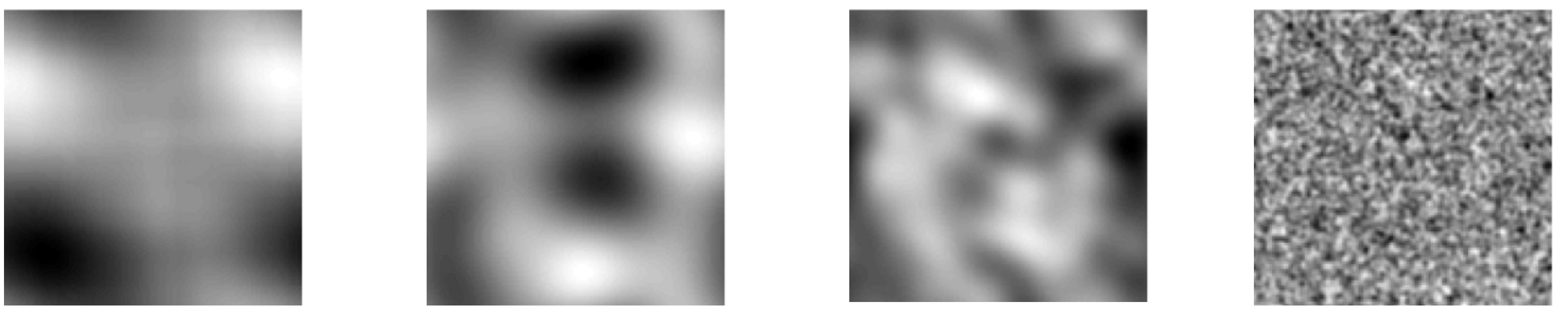

- Noiselets [Coifman et al.]

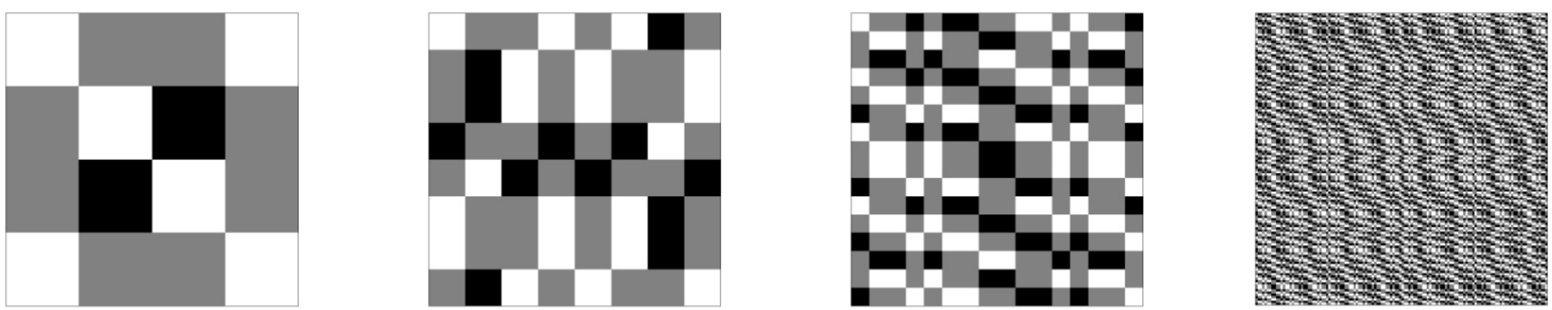

Example: Ellipse Parameters

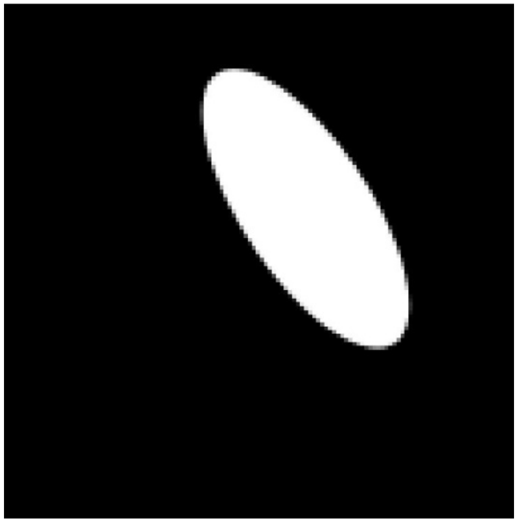

original

initial guess

initial error

N = 128x128 = 16384

K = 5 (major & minor axes; rotation; up & down)

M = 6 per scale (30 total): 57% success

M = 20 per scale (100 total): 99% success

Multi-Signal Recovery: "Manifold Lifting"

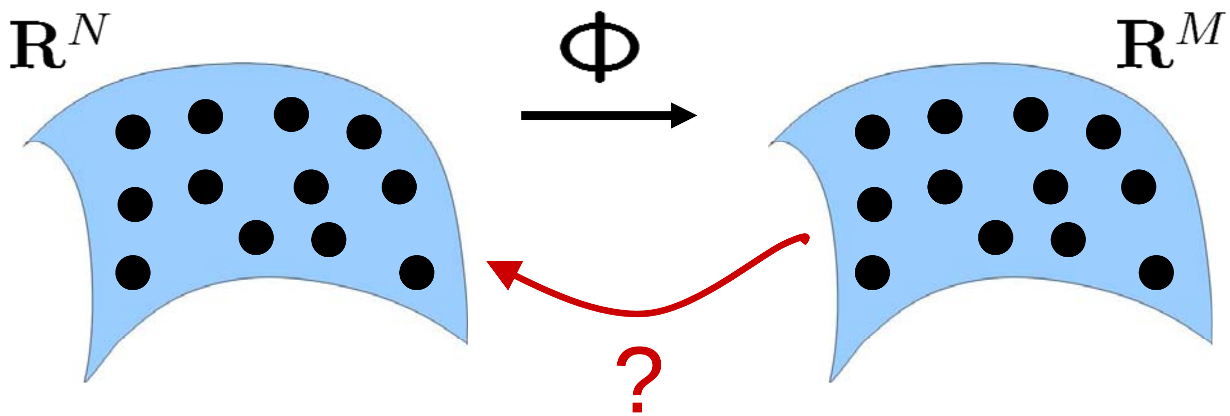

200 images

$N = 64^2 = 4096$

192

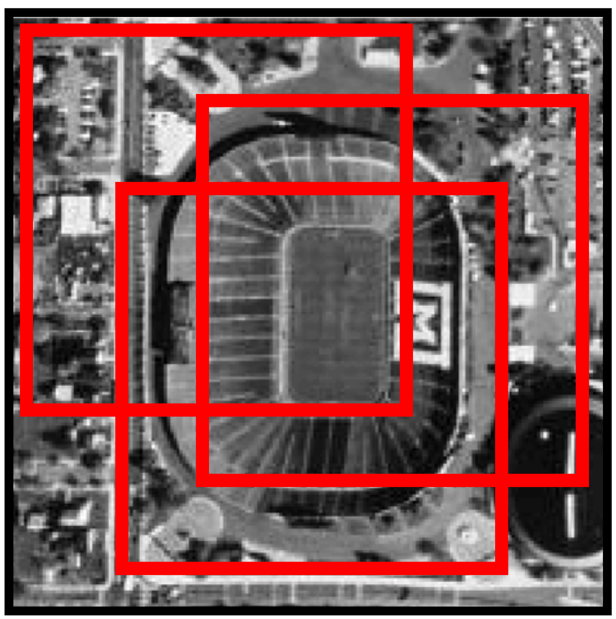

192

$M = 96$

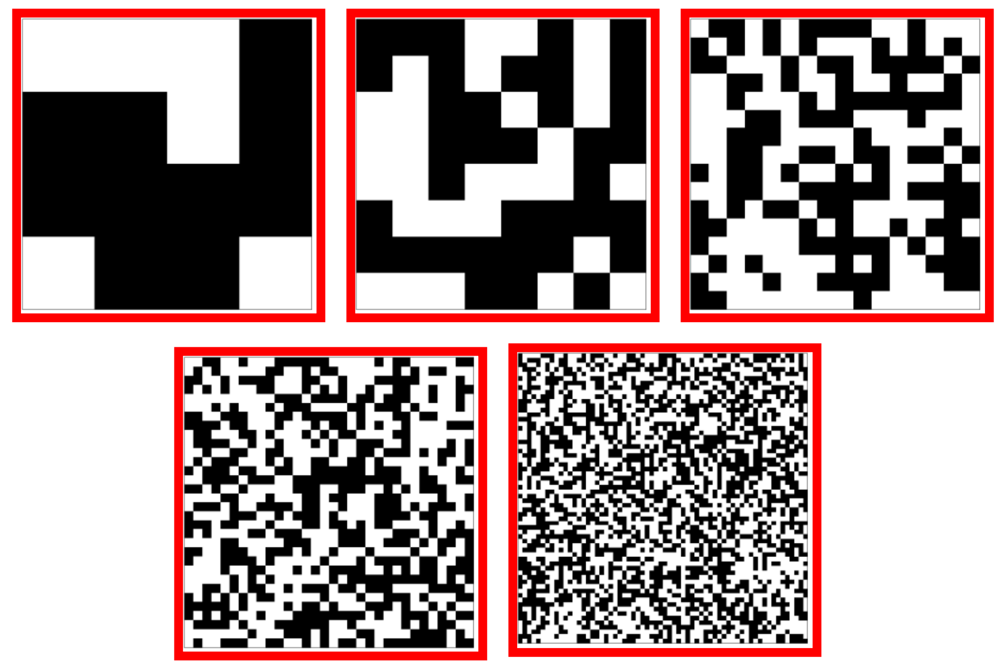

Final Reconstruction

image-by-image reconstruction without using manifold structure

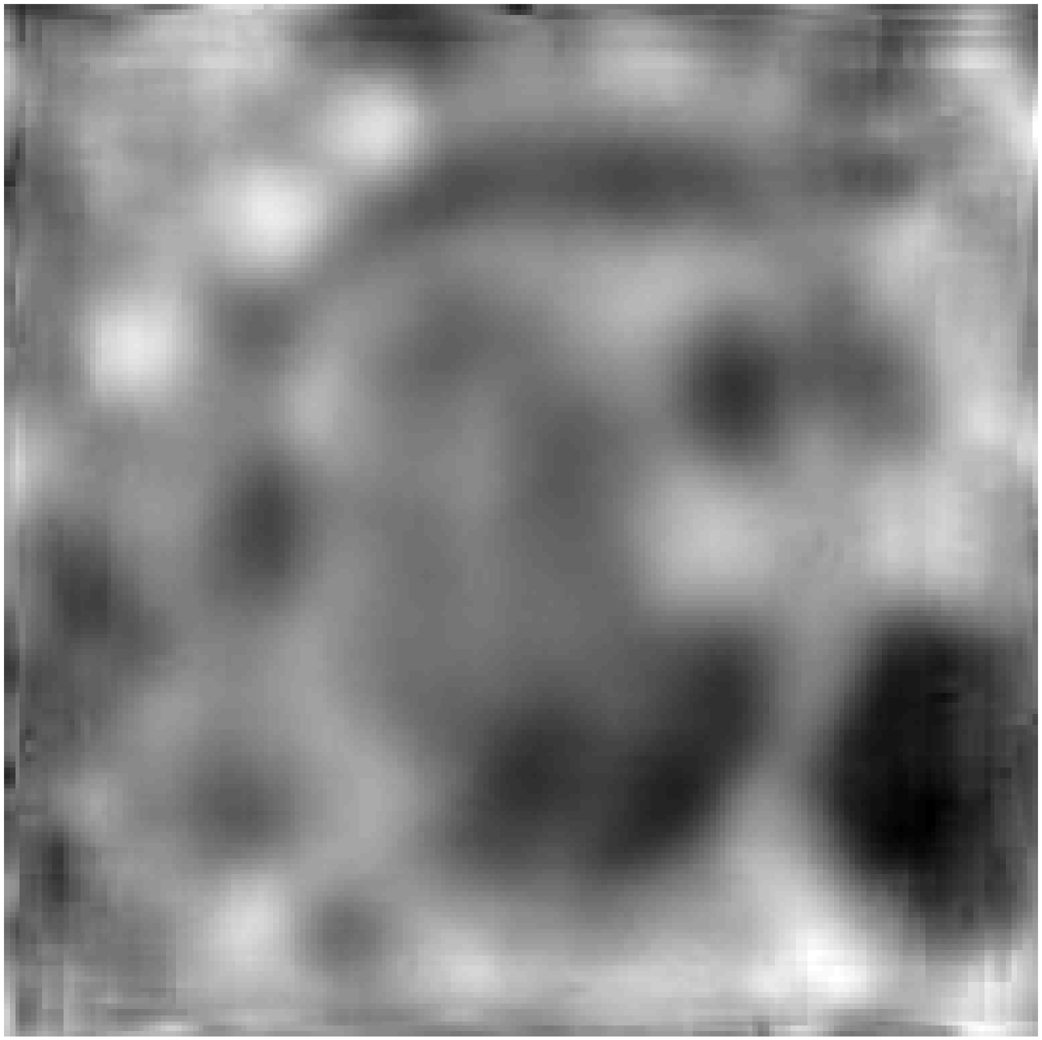

PSNR 15.4dB

joint reconstruction using manifold structure

PSNR 23.8dB

Conclusions

- Image manifolds contain rich geometric structure
- Image appearance manifolds
  - non-differentiable, due to sharp edges
  - edge migration $\rightarrow$ global non-differentiability
    - wavelet-like multiscale structure
    - accessible by regularizing each image
  - edge occlusion $\rightarrow$ local non-differentiability
- Can exploit multiscale structure in algorithms
  - proxy for standard calculus
  - new interpretation for image registration, etc.
  - applications in Compressive Sensing