Multiscale Random Projections for Compressive Classification

Marco F. Duarte

Joint work with

Mark Davenport, Michael Wakin, Jason Laska, Dharmpal Takhar, Kevin Kelly and Rich Baraniuk

dsp.rice.edu/cs

Data Explosion

- DSP revolution:
  sample first and ask questions later

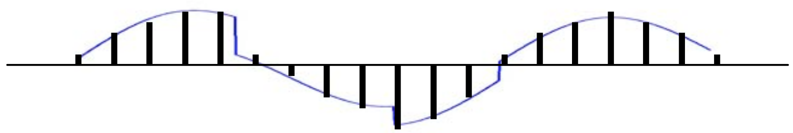

- Increasing ***pressure*** on classification algorithms
  - ever faster training and classification rates
  - ever larger, higher-dimensional data
  - ever lower energy consumption
  - radically new sensing modalities

- How can we acquire and process high-dimensional data quickly and efficiently?

Compressive Classification

- ***Random projections*** *preserve information*
  - Compressive Sensing (CS) (Candès, Donoho – 2004)
  - Johnson-Lindenstrauss Lemma (point clouds – 1984)

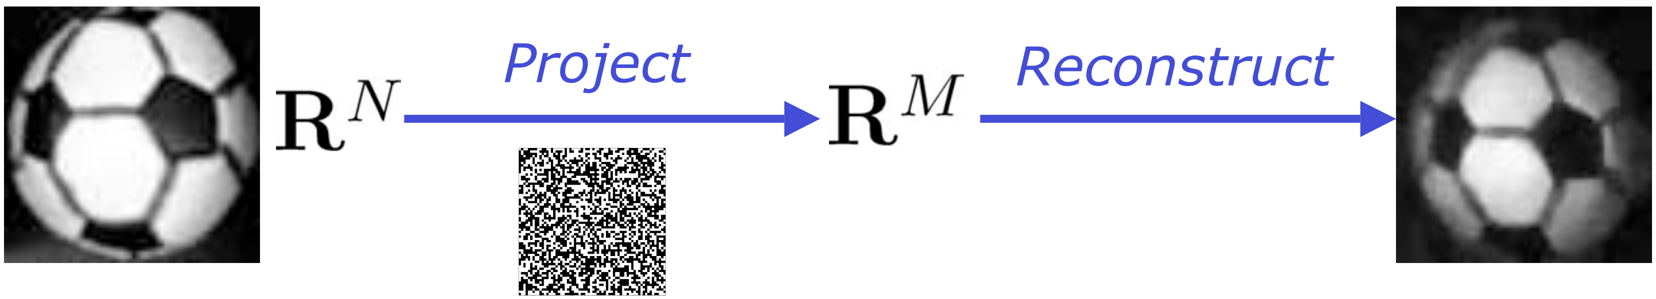

- If we can reconstruct a signal from compressive measurements, we should be able to perform
  - detection
  - classification
  - estimation
  - …

Matched Filter

- Signal x belongs to one of J classes
- Observed with some parameterized transformation
  - translation, rotation, scaling, lighting conditions, etc.
  - observation parameter unknown

$$
\begin{aligned}
\mathcal{H}_1 : &\quad x = \mathcal{T}_{\theta_1} s_1 + n \\
\mathcal{H}_2 : &\quad x = \mathcal{T}_{\theta_2} s_2 + n \\
\vdots \\
\mathcal{H}_J : &\quad x = \mathcal{T}_{\theta_J} s_J + n
\end{aligned}
$$

- *Maximum likelihood* classifier with AWGN

$$
\min_{j, \widehat{\theta}_j} \| x - \mathcal{T}_{\widehat{\theta}_j} s_j \|_2
$$

- Solve via convolution when parameter = translation

Manifold Models

- K-dimensional *parameter* $\theta \in \Theta$

  captures degrees of freedom in signal $x_\theta \in R^N$

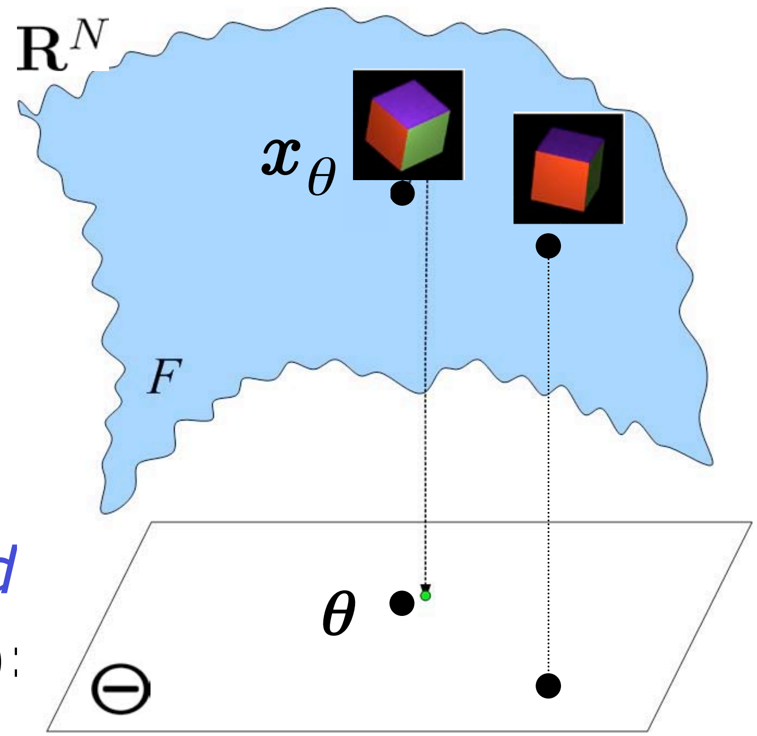

- Signal class $F = \{x_\theta : \theta \in \Theta\}$

  forms a K-dimensional *manifold*
  - Image appereance manifolds (IAM): shifts, rotations, etc.

- Dimensionality reduction and manifold learning
  - embeddings [ISOMAP; LLE; HLLE; …]
  - harmonic analysis [Belkin; Coifman; …]

Matched Filter

- *Maximum likelihood* classifier with AWGN

$$\min_{j,\widehat{\theta}_j} \|x - \mathcal{T}_{\widehat{\theta}_j} s_j\|_2$$

reduces to *nearest neighbor* classification when signal classes form manifolds

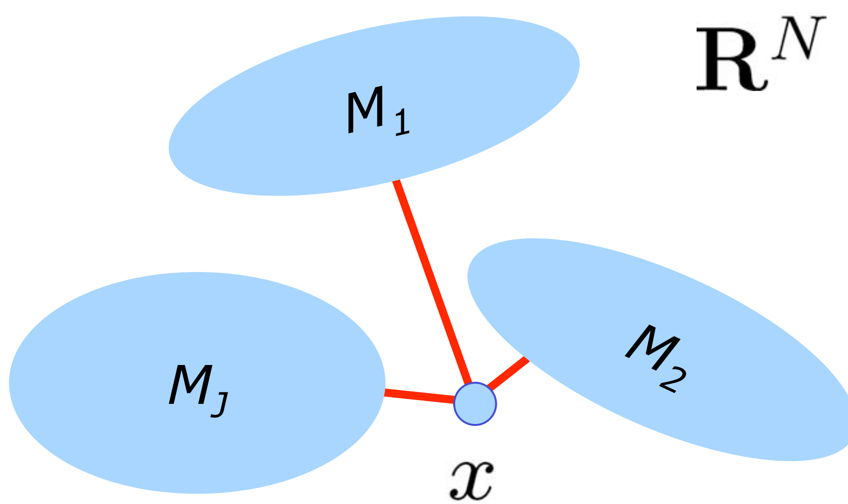

"Smashed Filter"

- Solve "nearest manifold" problem using random projections

$$\min_{j,\widehat{\theta}_j} \|\Phi x - \Phi \mathcal{T}_{\widehat{\theta}_j} s_j\|_2$$

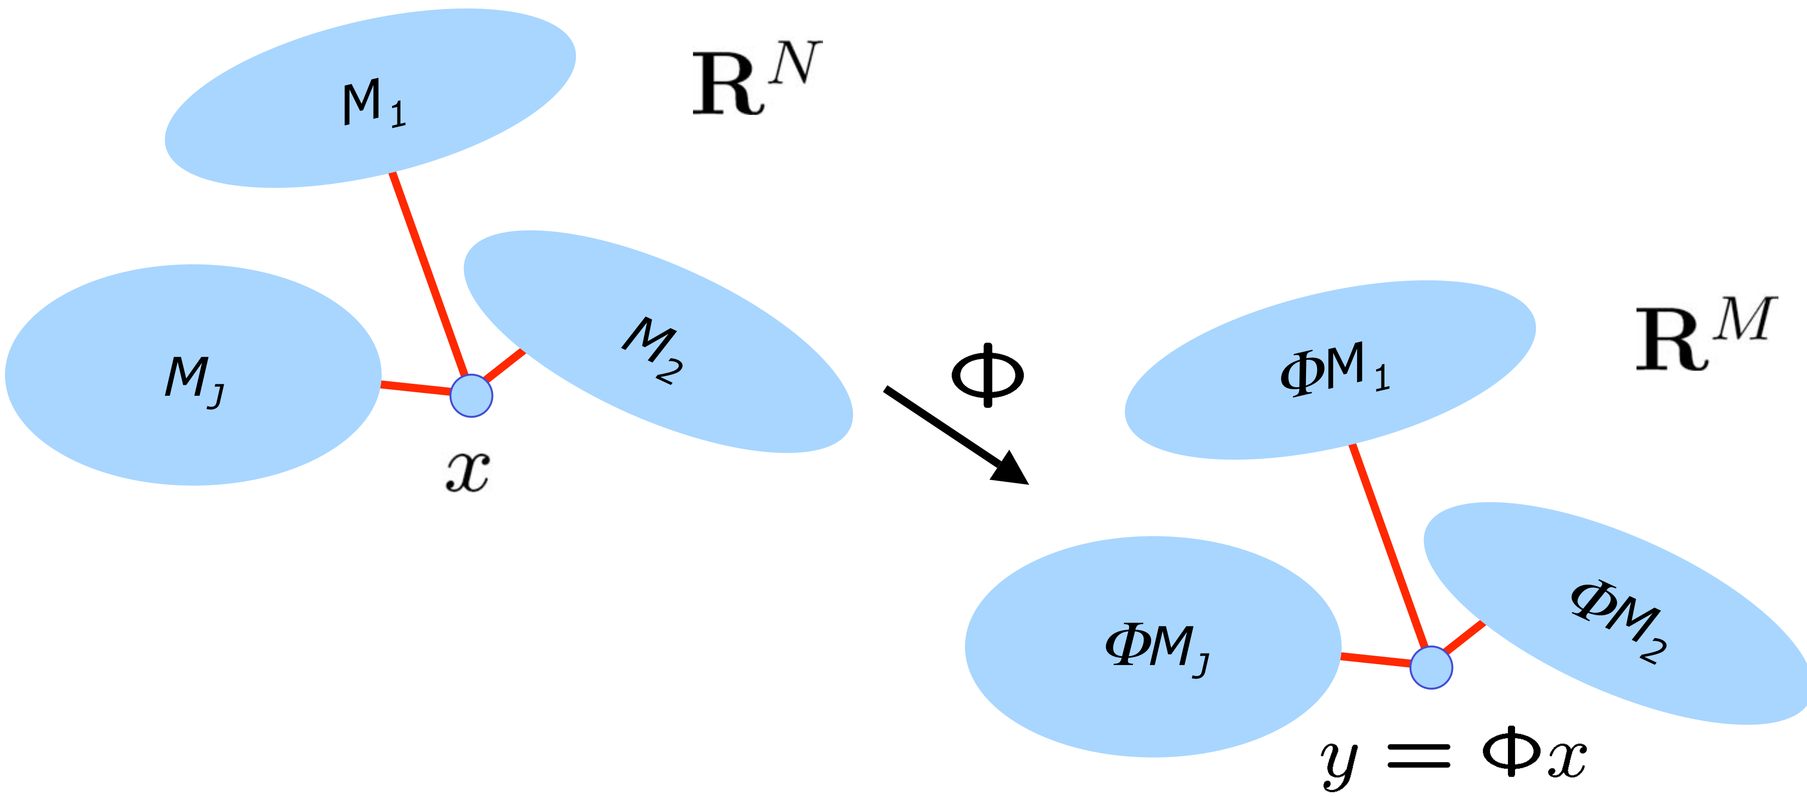

[M. Davenport et al., SPIE Electronic Imaging 07]

Stable Manifold Embedding

Theorem:
Let $F \subset R^N$ be a compact *K-dimensional manifold* with

- condition number $1/\tau$ (curvature, self-avoiding)
- volume V

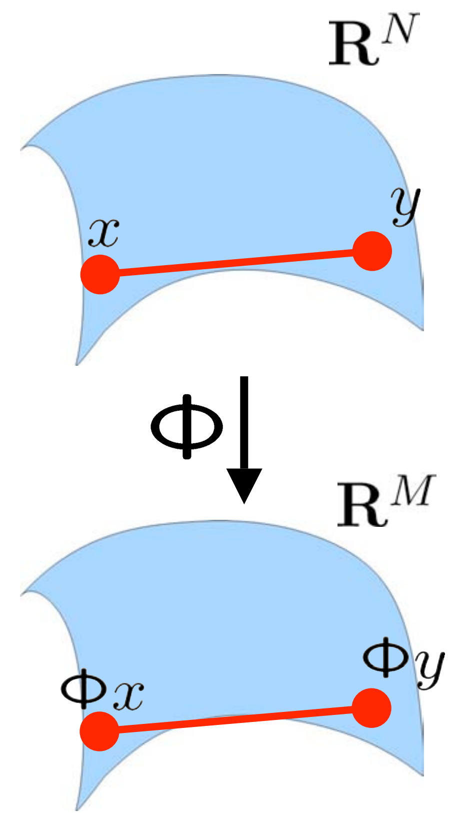

Let Φ be a random MxN orthoprojector with

$$M = O\left(\frac{K \log(NV\tau^{-1}\epsilon^{-1}) \log(1/\rho)}{\epsilon^2}\right).$$

Then with probability at least $1-\rho$, the following statement holds:
For every pair $x,y \in F$,

$$(1-\epsilon)\, \|x-y\|_2 \leq \|\Phi x - \Phi y\|_2 \leq (1+\epsilon)\, \|x-y\|_2 .$$

[R. Baraniuk, M. Wakin, FOCM, in press]

Multiple Manifold Embedding

Corollary:

Let $M_1,\ldots,M_P \subset \mathbb{R}^N$ be compact K-dimensional manifolds with

- condition number $1/\tau$ (curvature, self-avoiding)
- volume V
- min dist$(M_j, M_k) > \tau$

Let Φ be a random MxN orthoprojector with

$$M = O\left(\frac{K \log(NPV\tau^{-1}\epsilon^{-1}) \log(1/\rho)}{\epsilon^2}\right).$$

Then with probability at least 1-ρ, the following statement holds:

For every pair $x,y \in \cup M_j$,

$$(1-\epsilon)\,\|x-y\|_2 \le \|\Phi x - \Phi y\|_2 \le (1+\epsilon)\,\|x-y\|_2 .$$

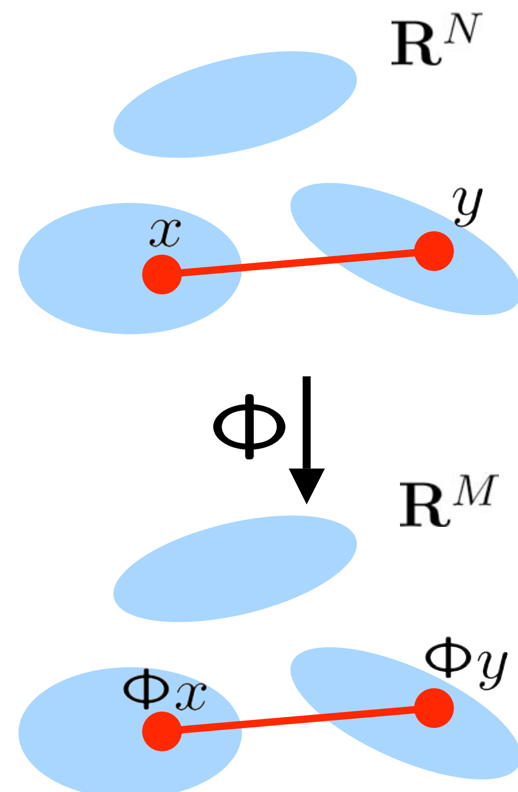

[M. Davenport et al., SPIE Electronic Imaging 07]

The *Smashed Filter*

- *Compressive* manifold classification with GLRT
  - nearest-manifold classifier
  - manifolds classified are now $\Phi M_j = \{\Phi f_j(\theta_j) \ : \ \theta_j \in \Theta_j\}$

$$H_j \ : \ y = \Phi(m_j + n), \quad m_j \in M_j$$

$$m_j = f_j(\theta_j)$$

$$\underset{j=1,\ldots,P}{\arg\min} \ \|y - \Phi f_j(\widehat{\theta}_j)\|_2$$

$$\widehat{\theta}_j = \underset{\theta \in \Theta_j}{\arg\min} \ \|y - \Phi f_j(\theta_j)\|_2$$

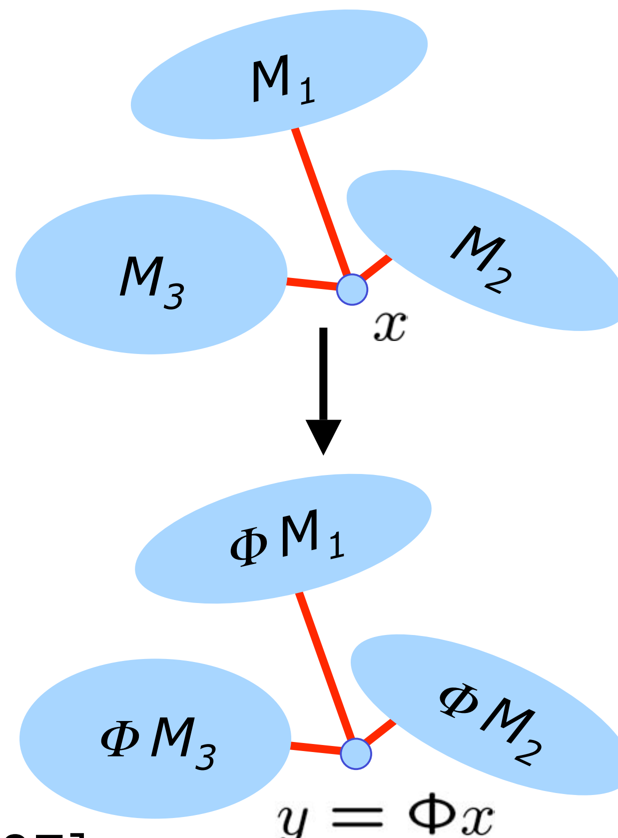

[M. Davenport et al., SPIE Electronic Imaging 07]

Newton's Method...

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

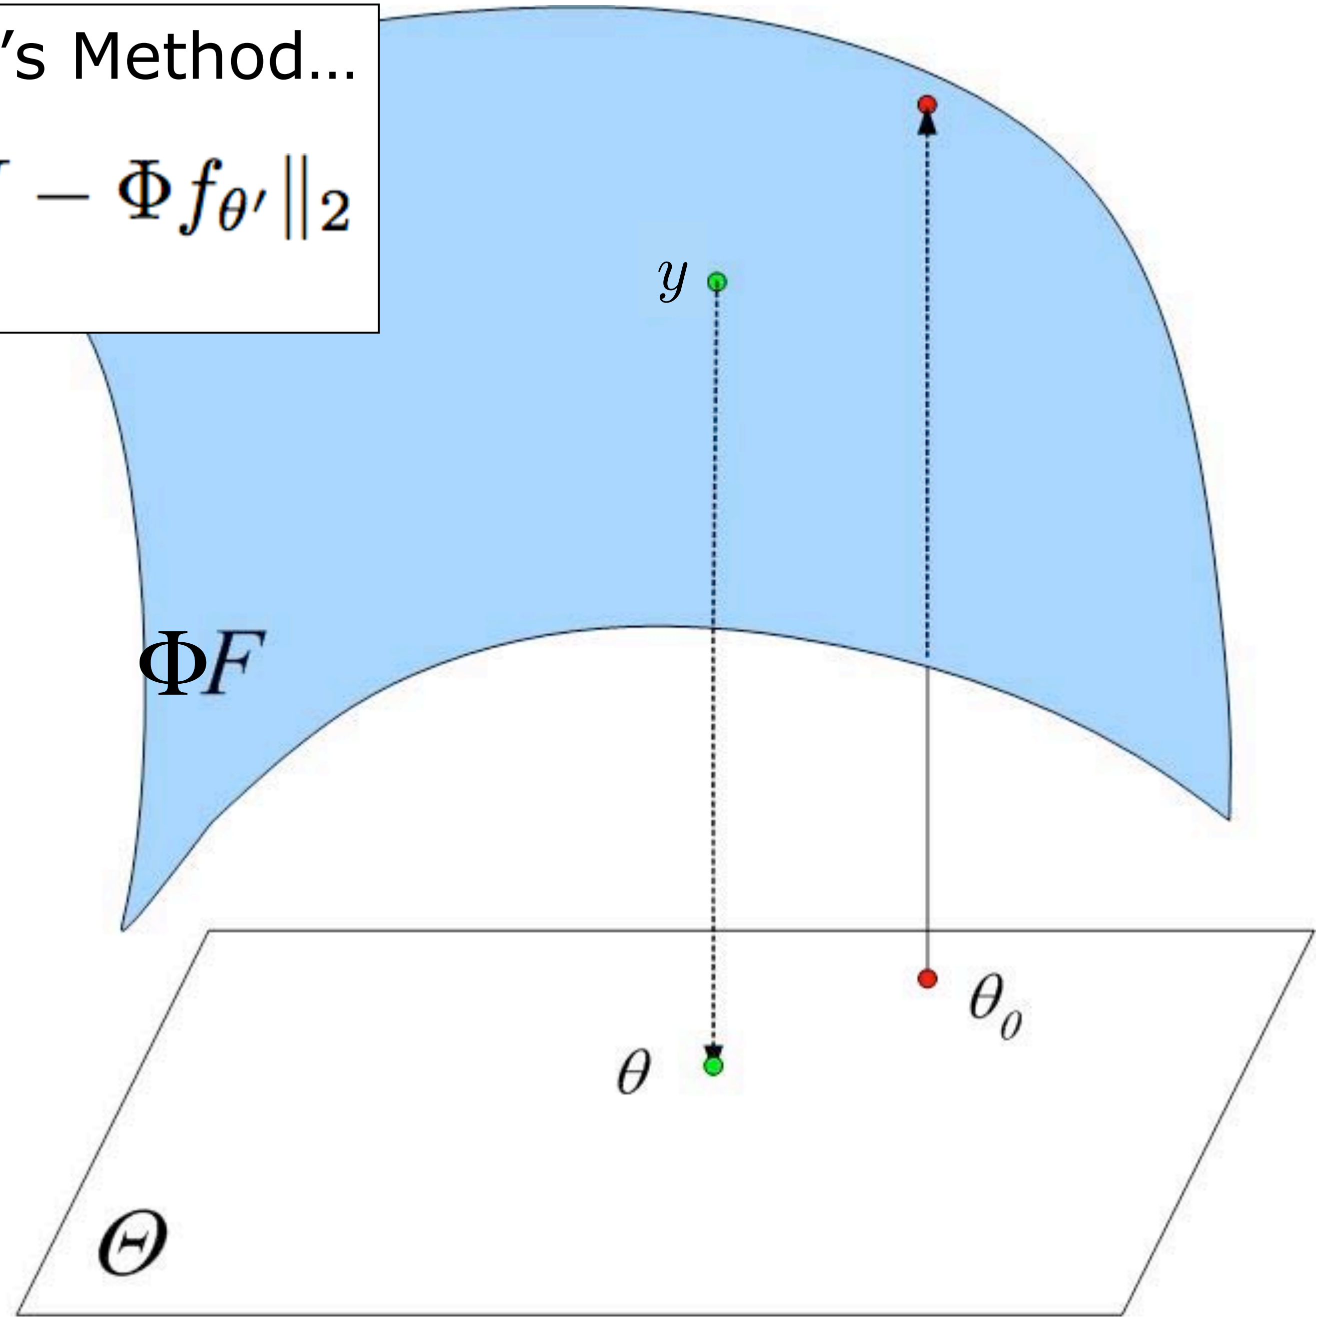

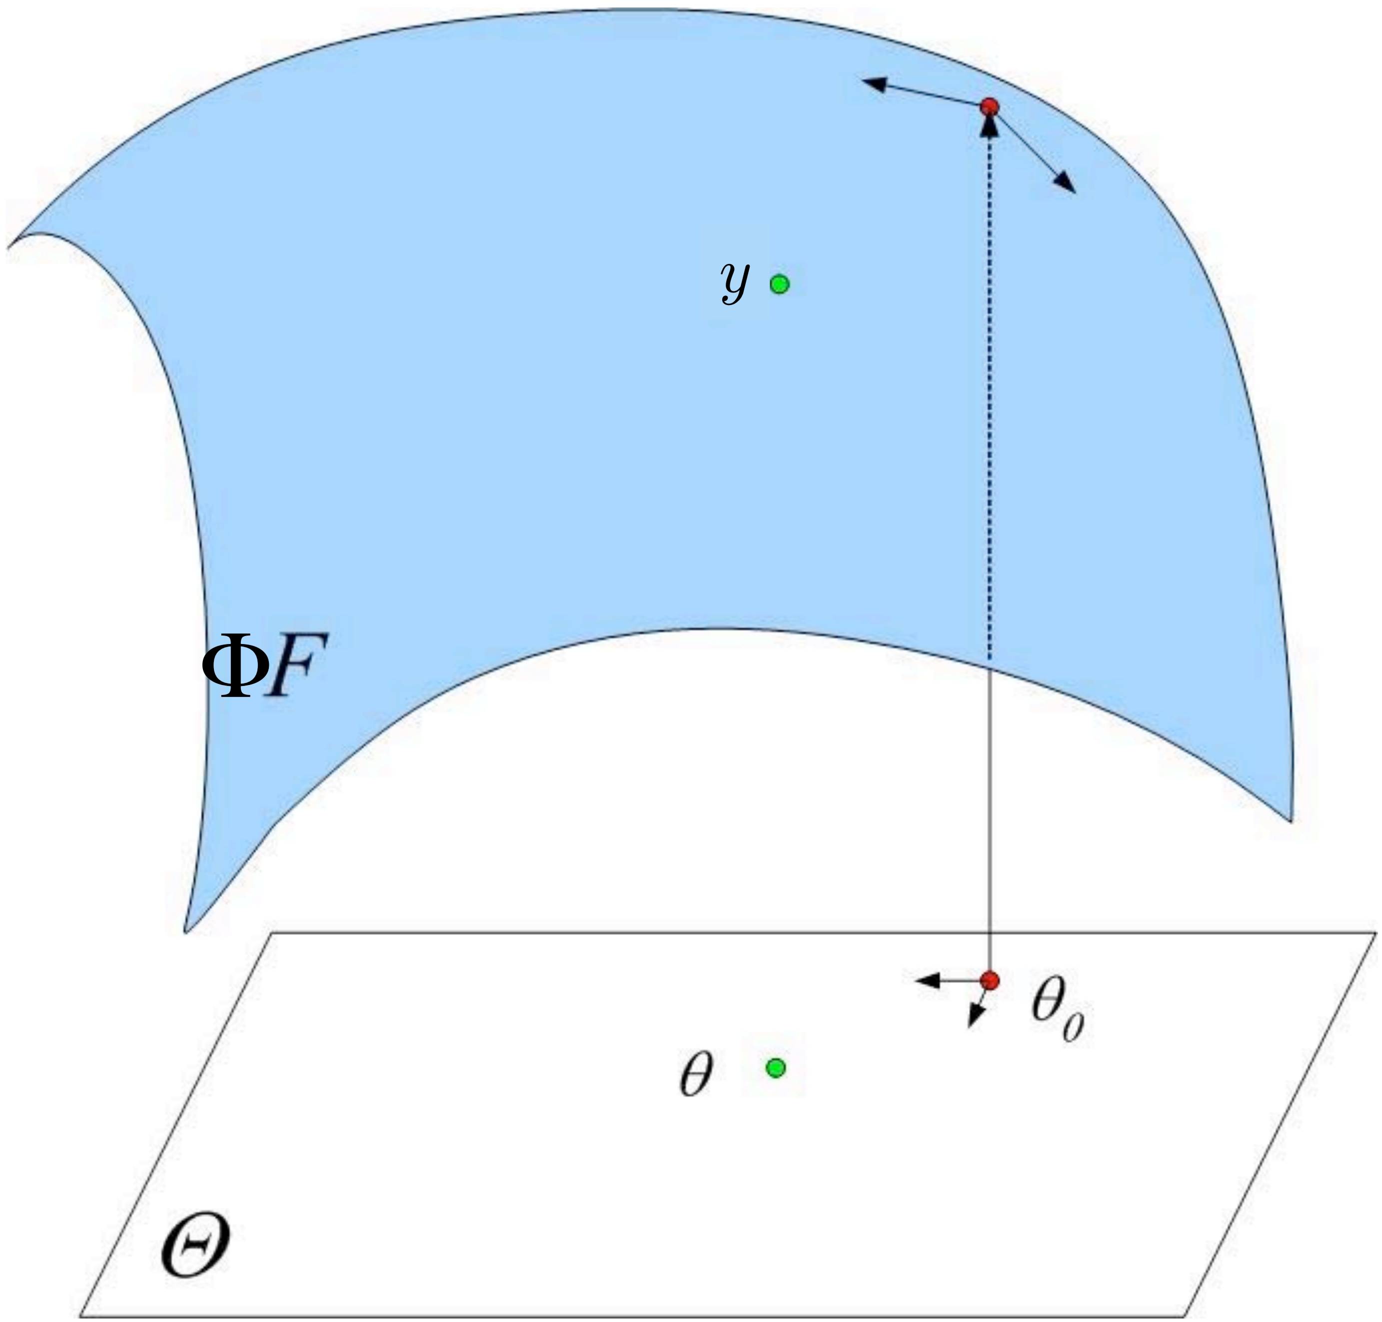
y
ΦF
θ_0
θ
Θ

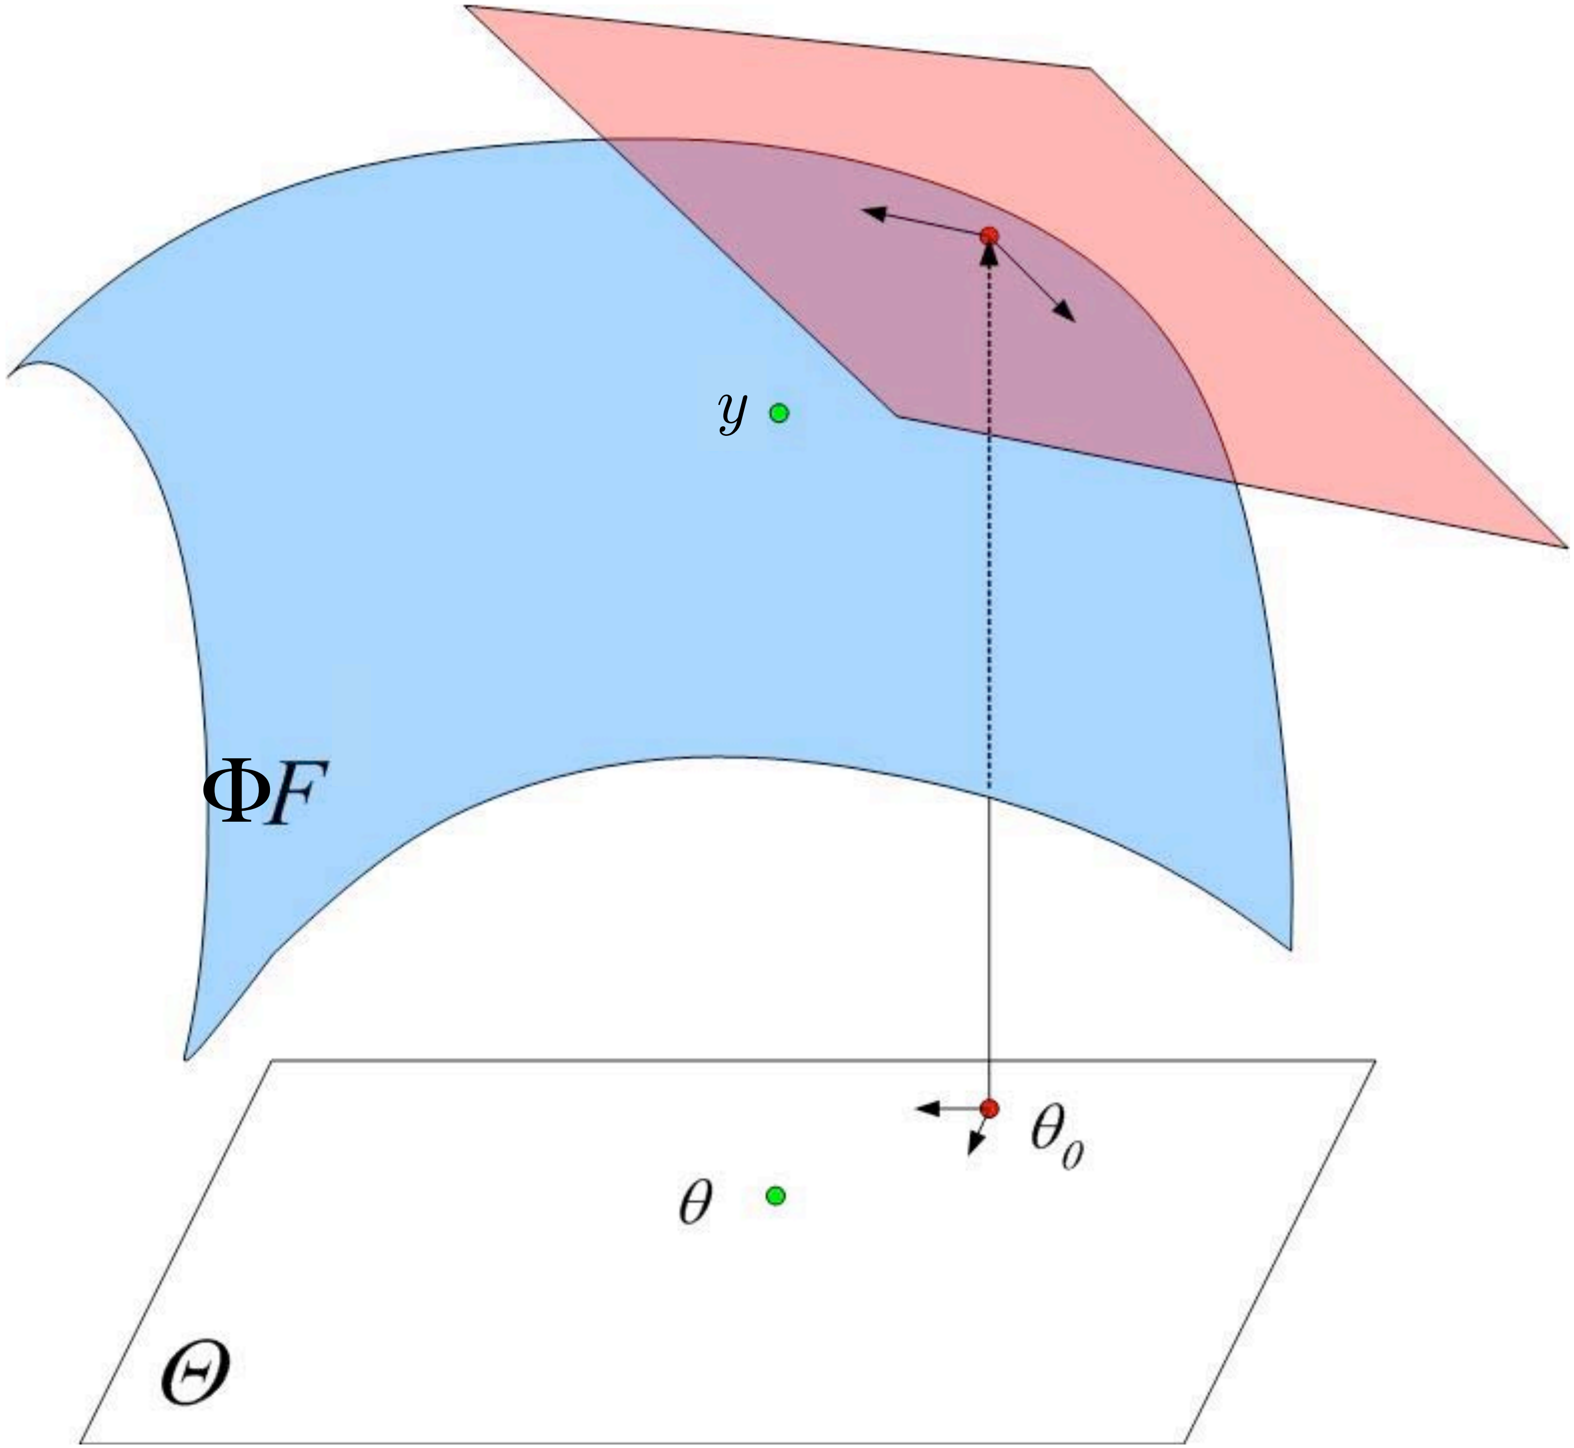

y
ΦF
θ₀
θ
Θ

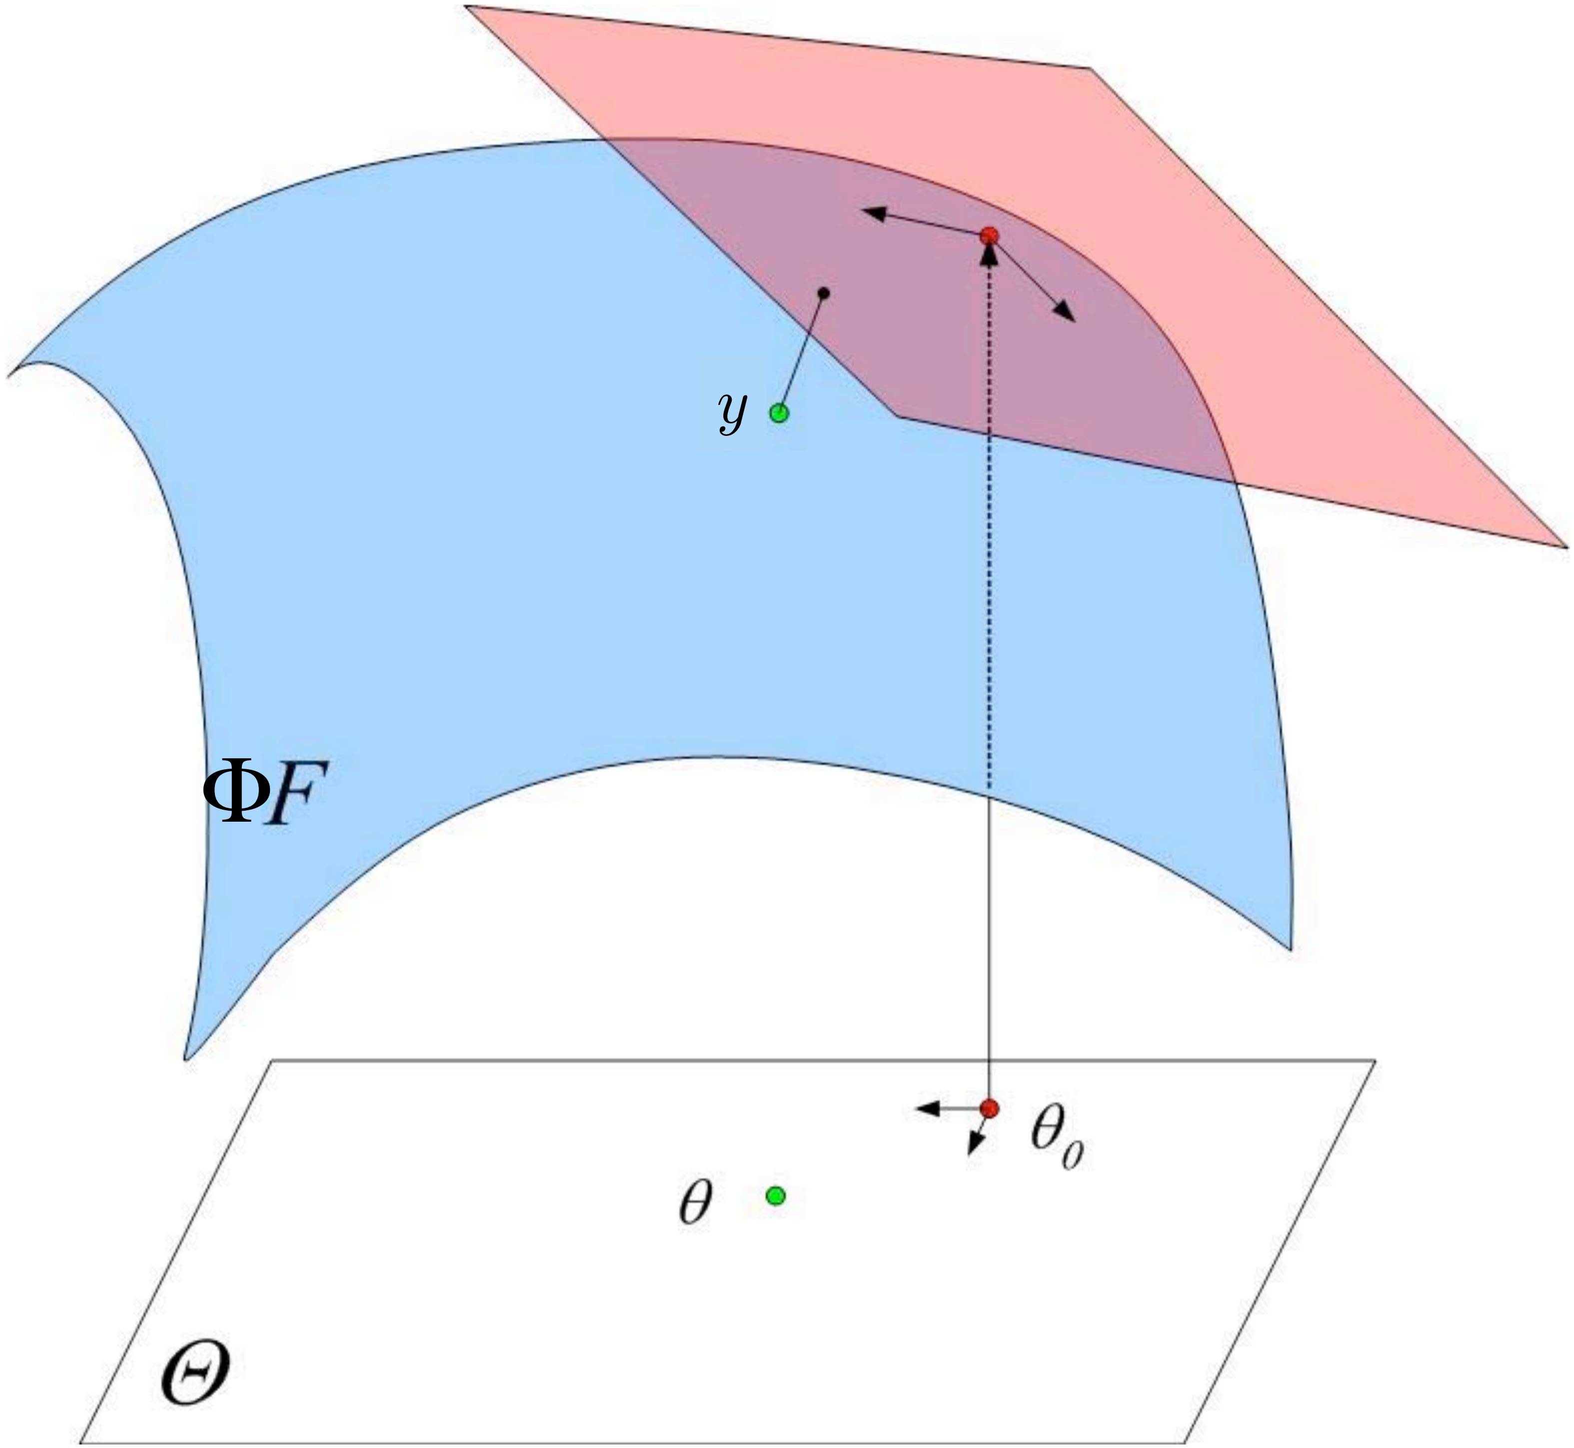

y
ΦF
θ0
θ
Θ

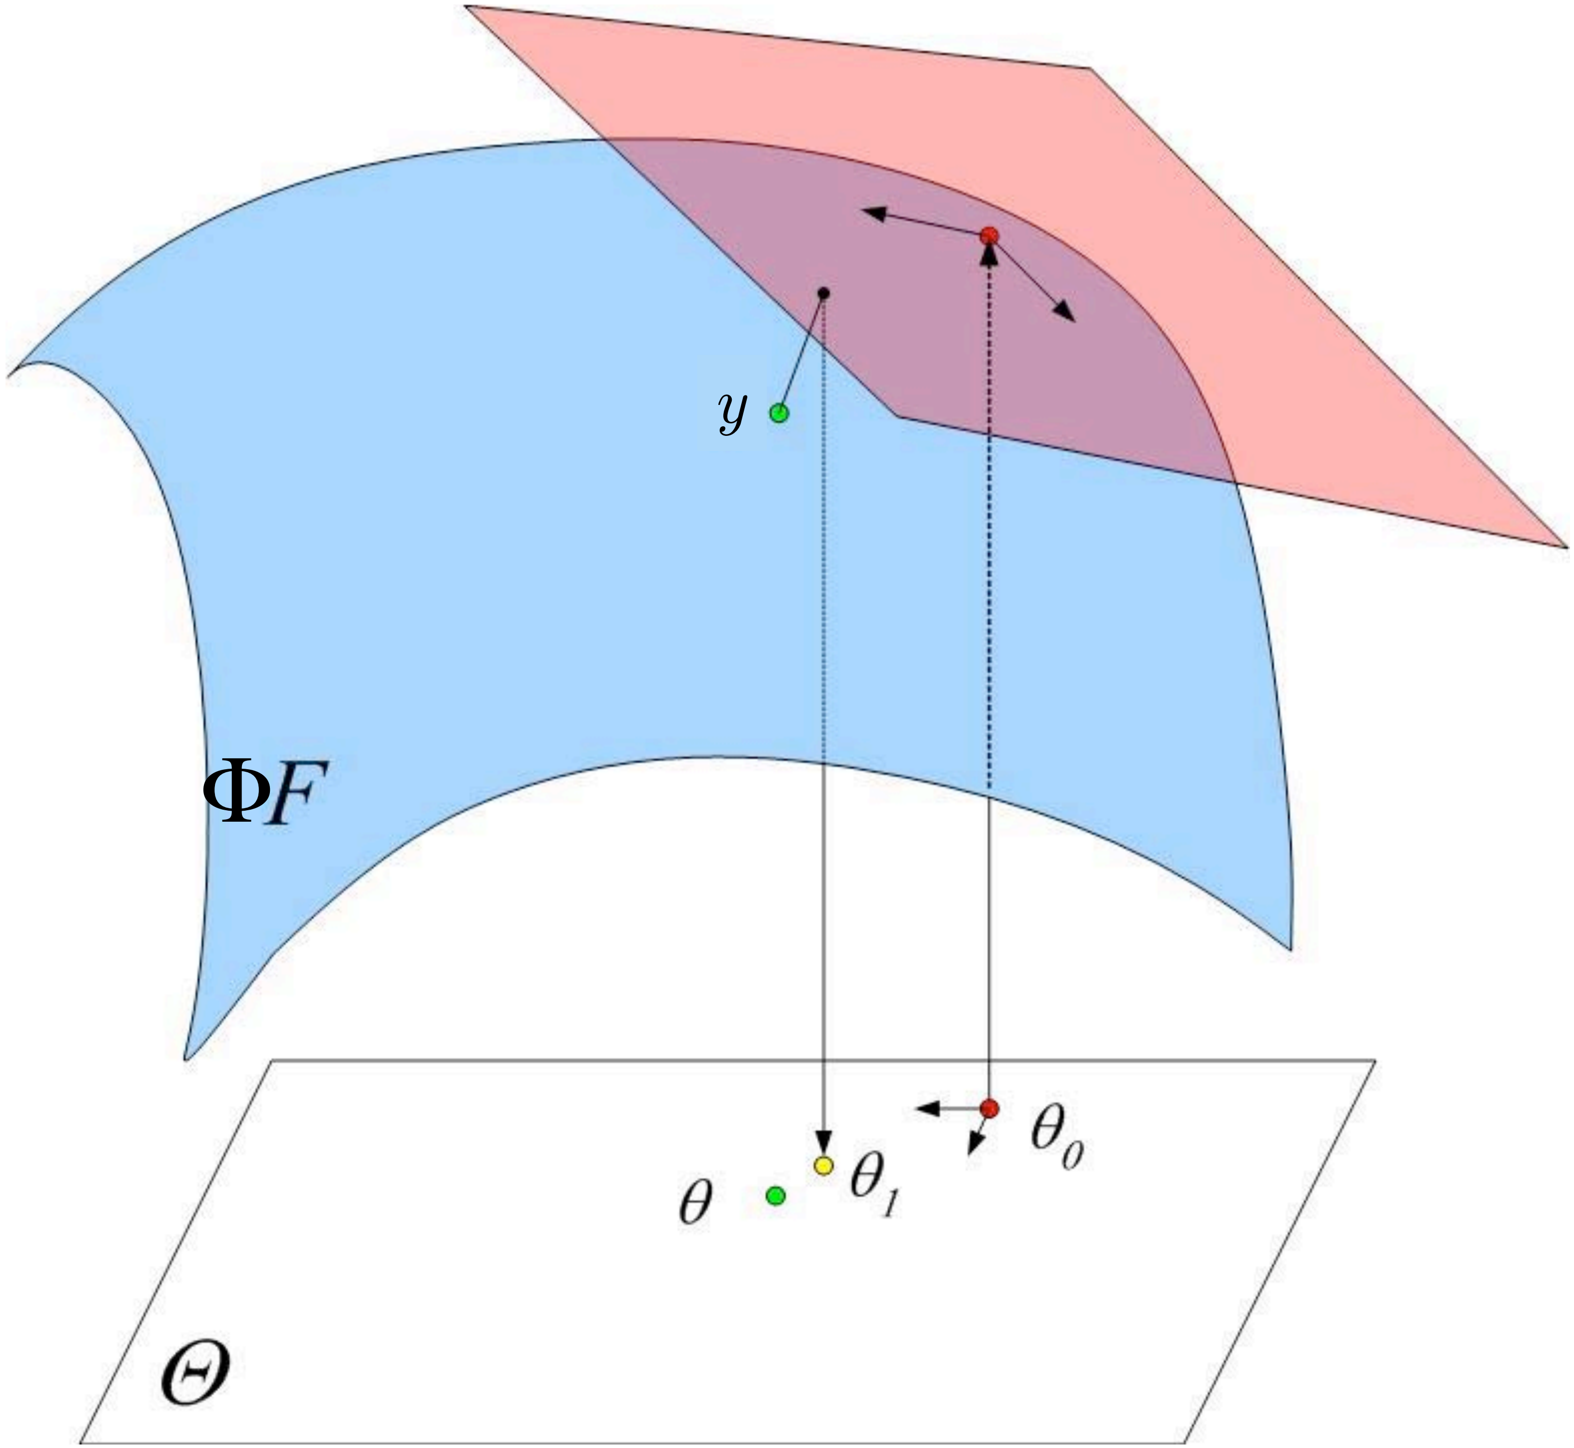

y
ΦF
θ_0
θ_1
θ
Θ

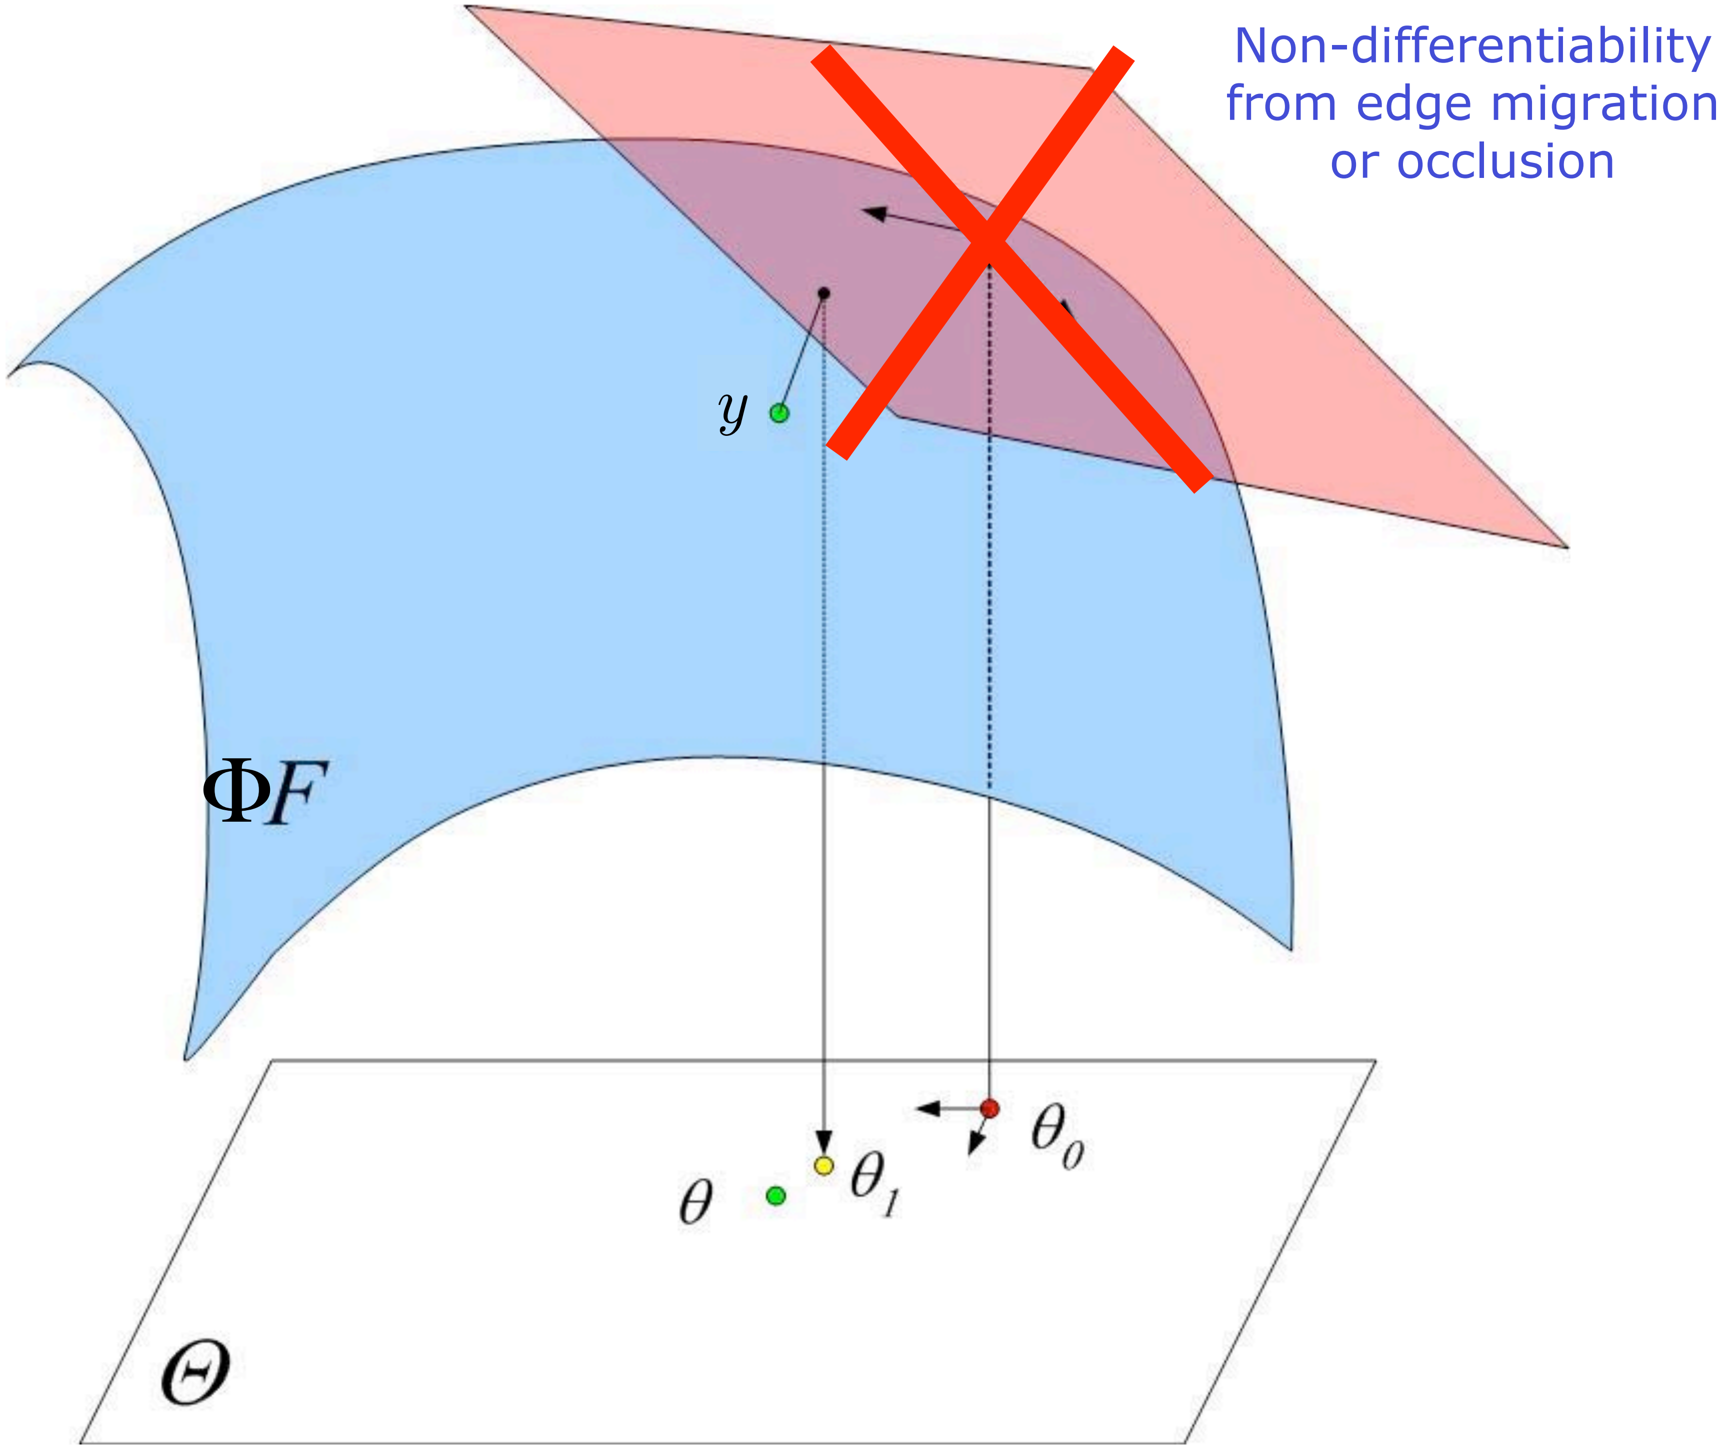
Non-differentiability
from edge migration
or occlusion
y
ΦF
θ
θ1
θ0
Θ

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

[M. Wakin, D. Donoho, H. Choi, and R. Baraniuk,
The Multiscale Structure of Non-Differentiable Image Manifolds, *SPIE Wavelets XI*, 2005.]

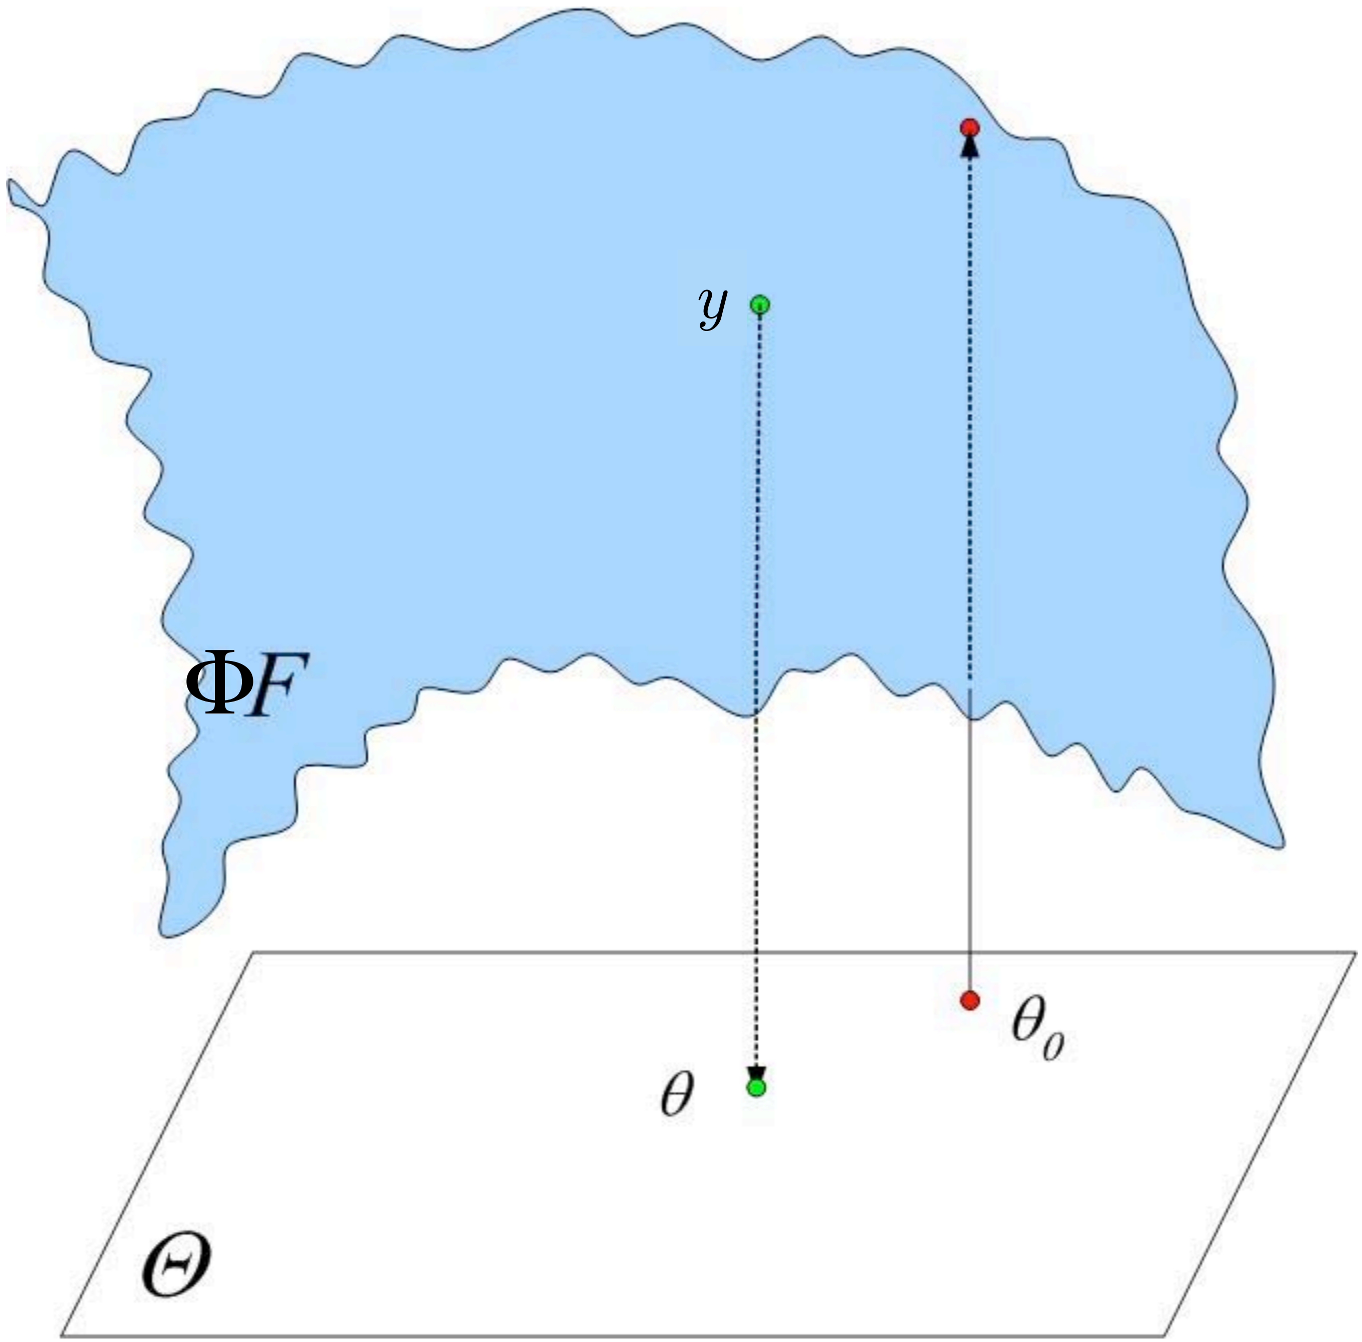

y
ΦF
θ
θ_0
Θ

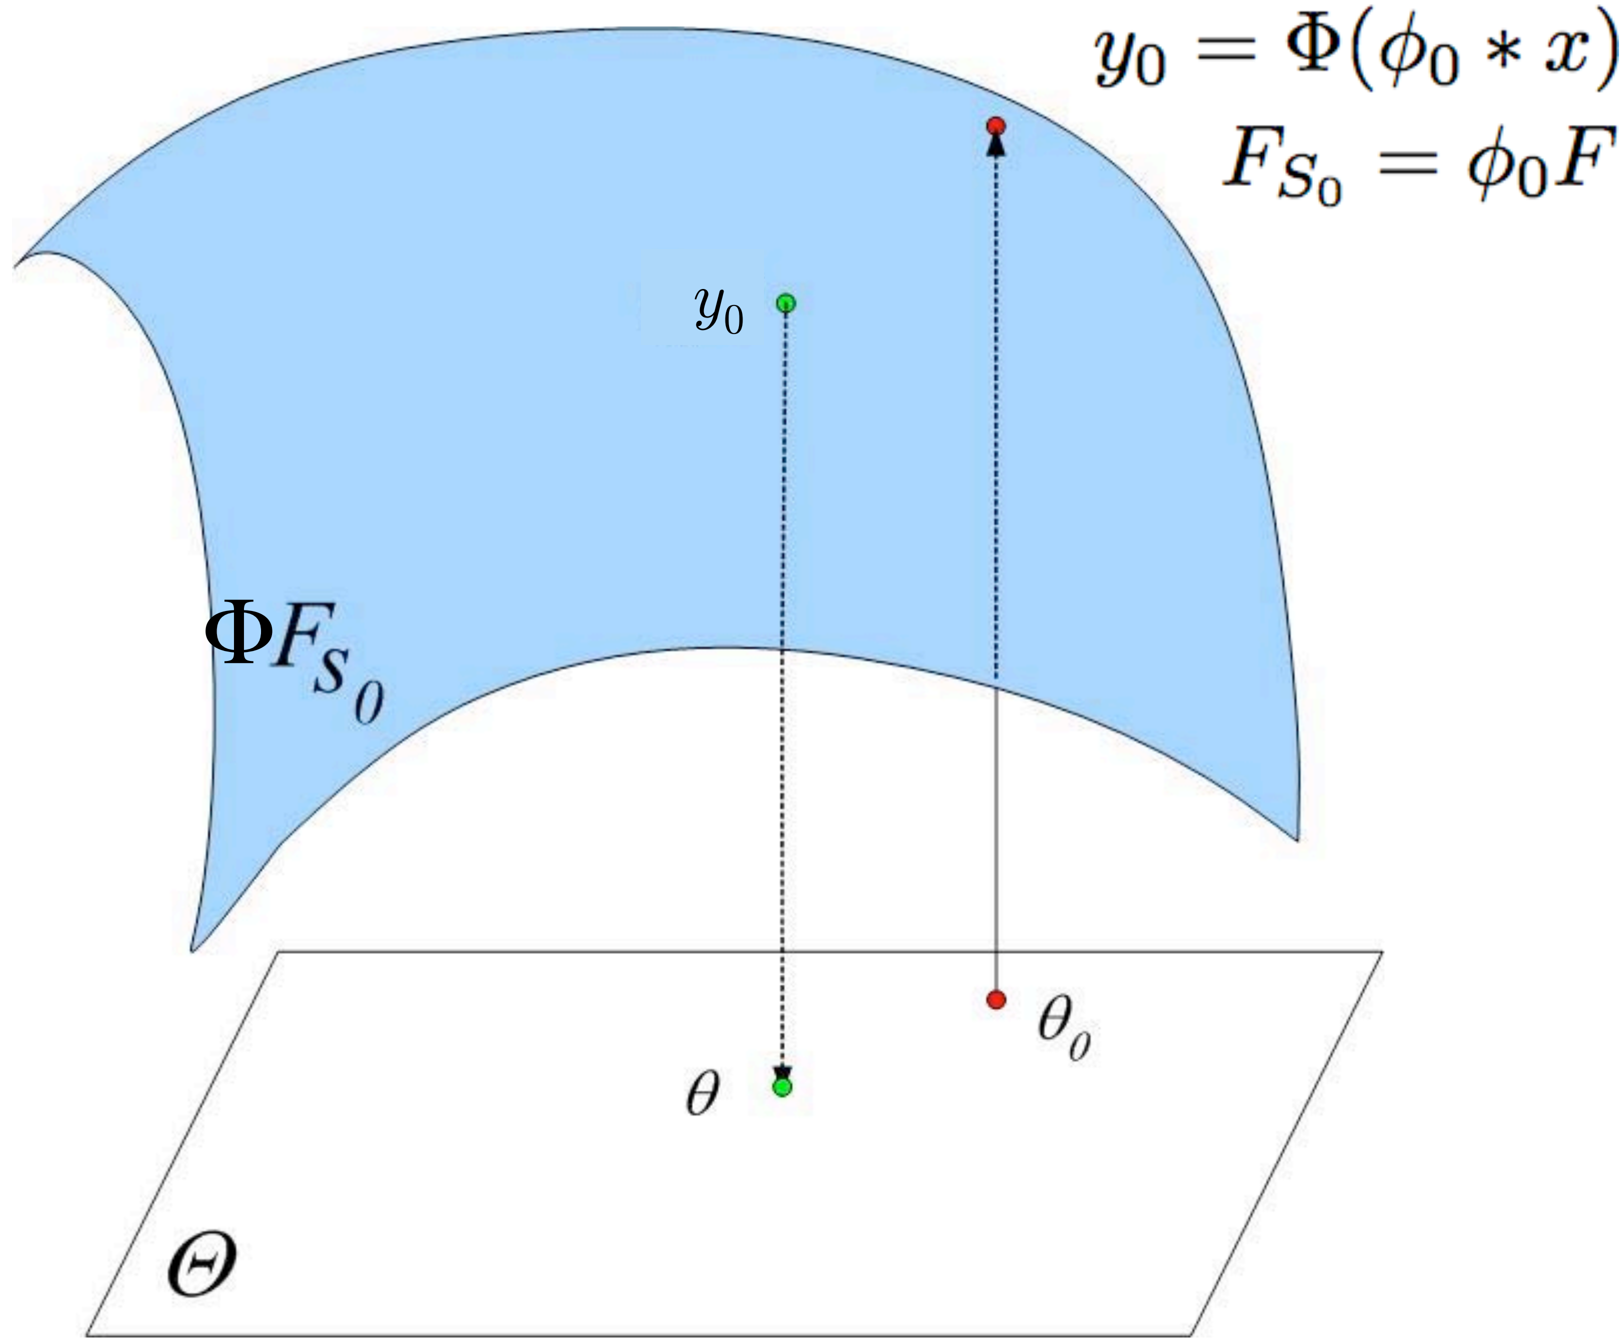

$y_0 = \Phi(\phi_0 * x)$
$F_{S_0} = \phi_0 F$
y_0
ΦF_{s_0}
θ_0
θ
Θ

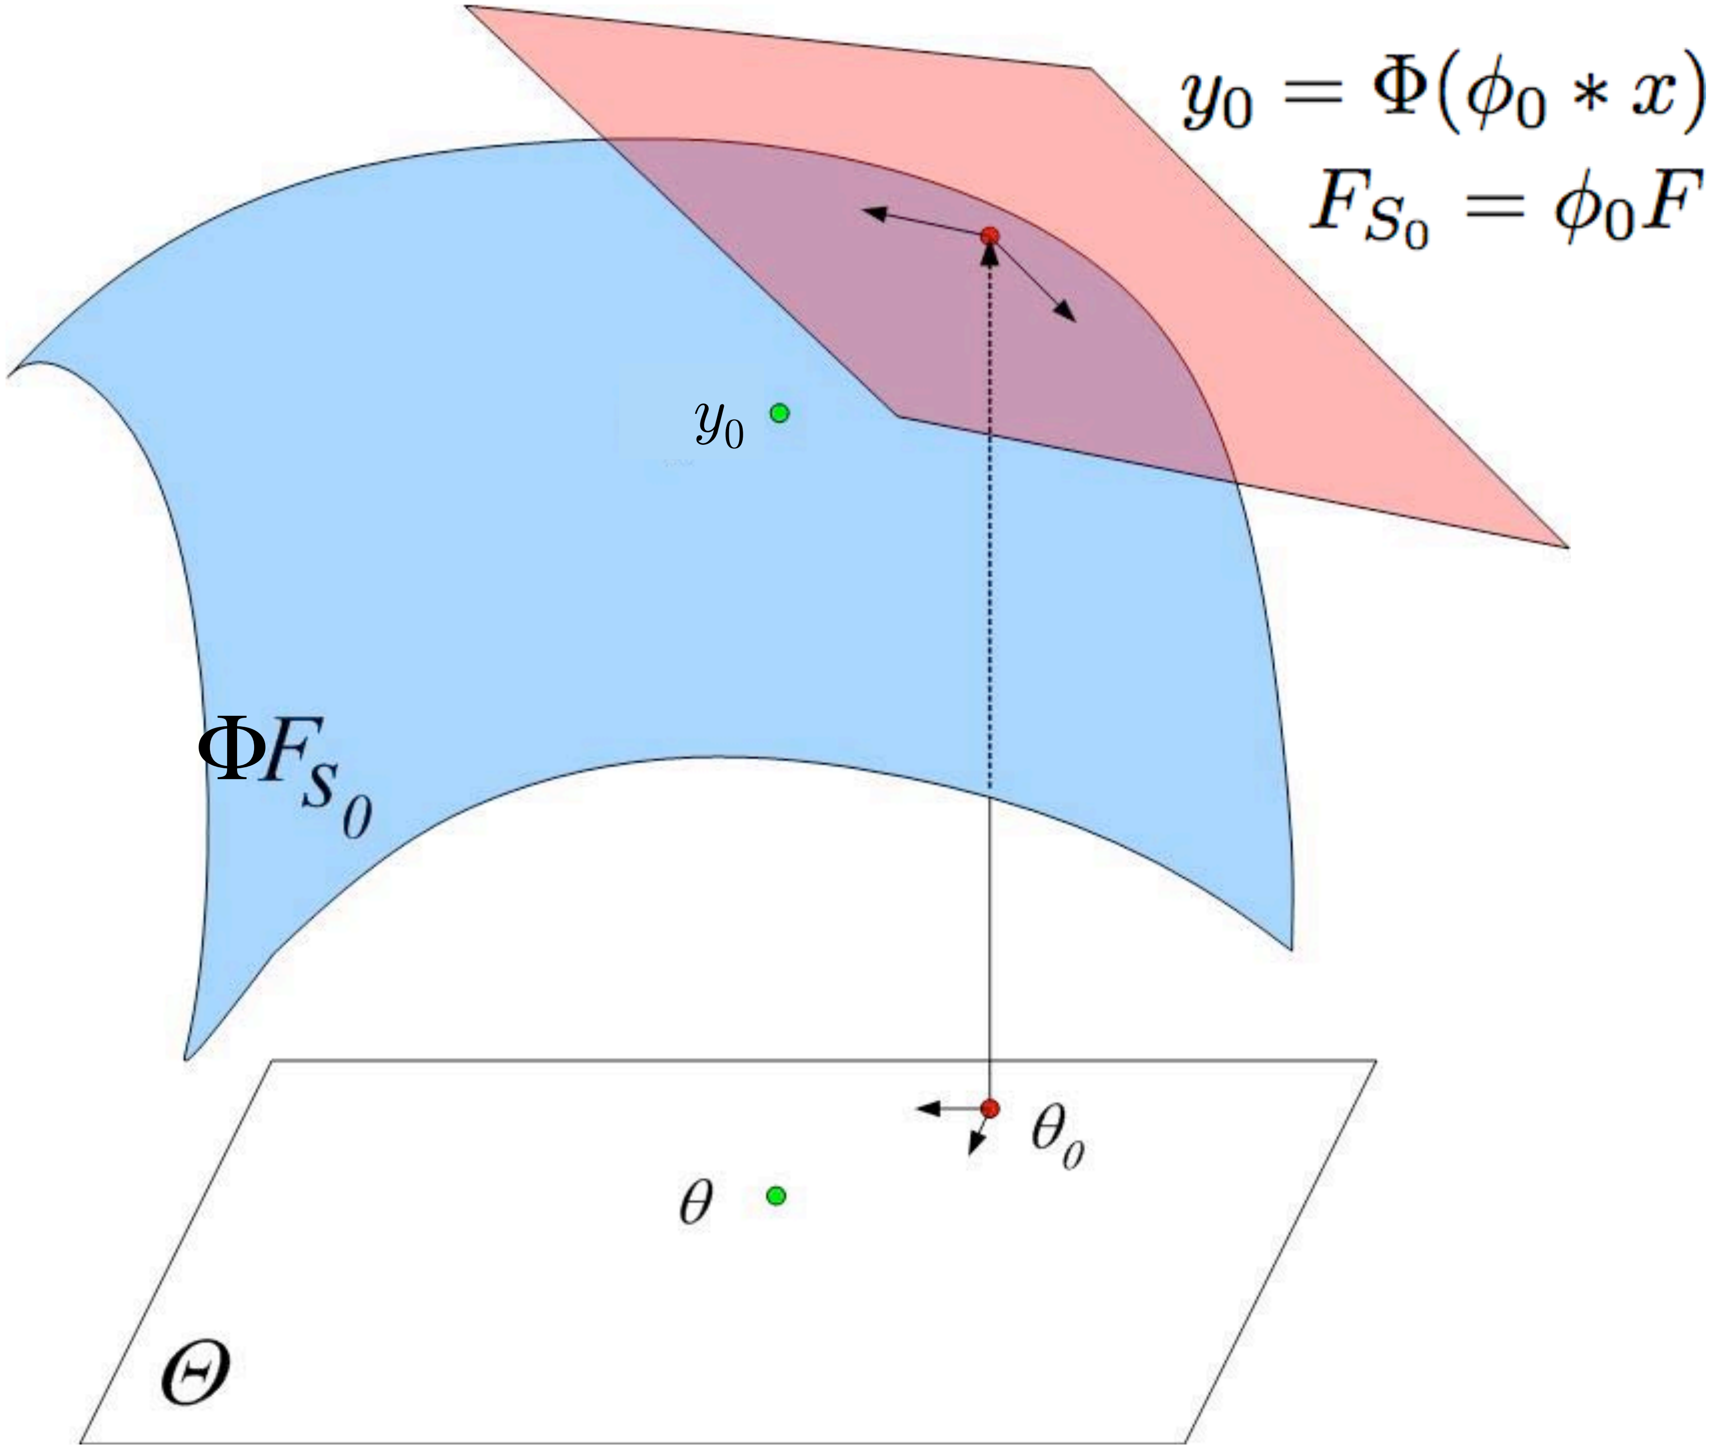

$y_0 = \Phi(\phi_0 * x)$
$F_{S_0} = \phi_0 F$
y_0
ΦF_{s_0}
θ_0
θ
Θ

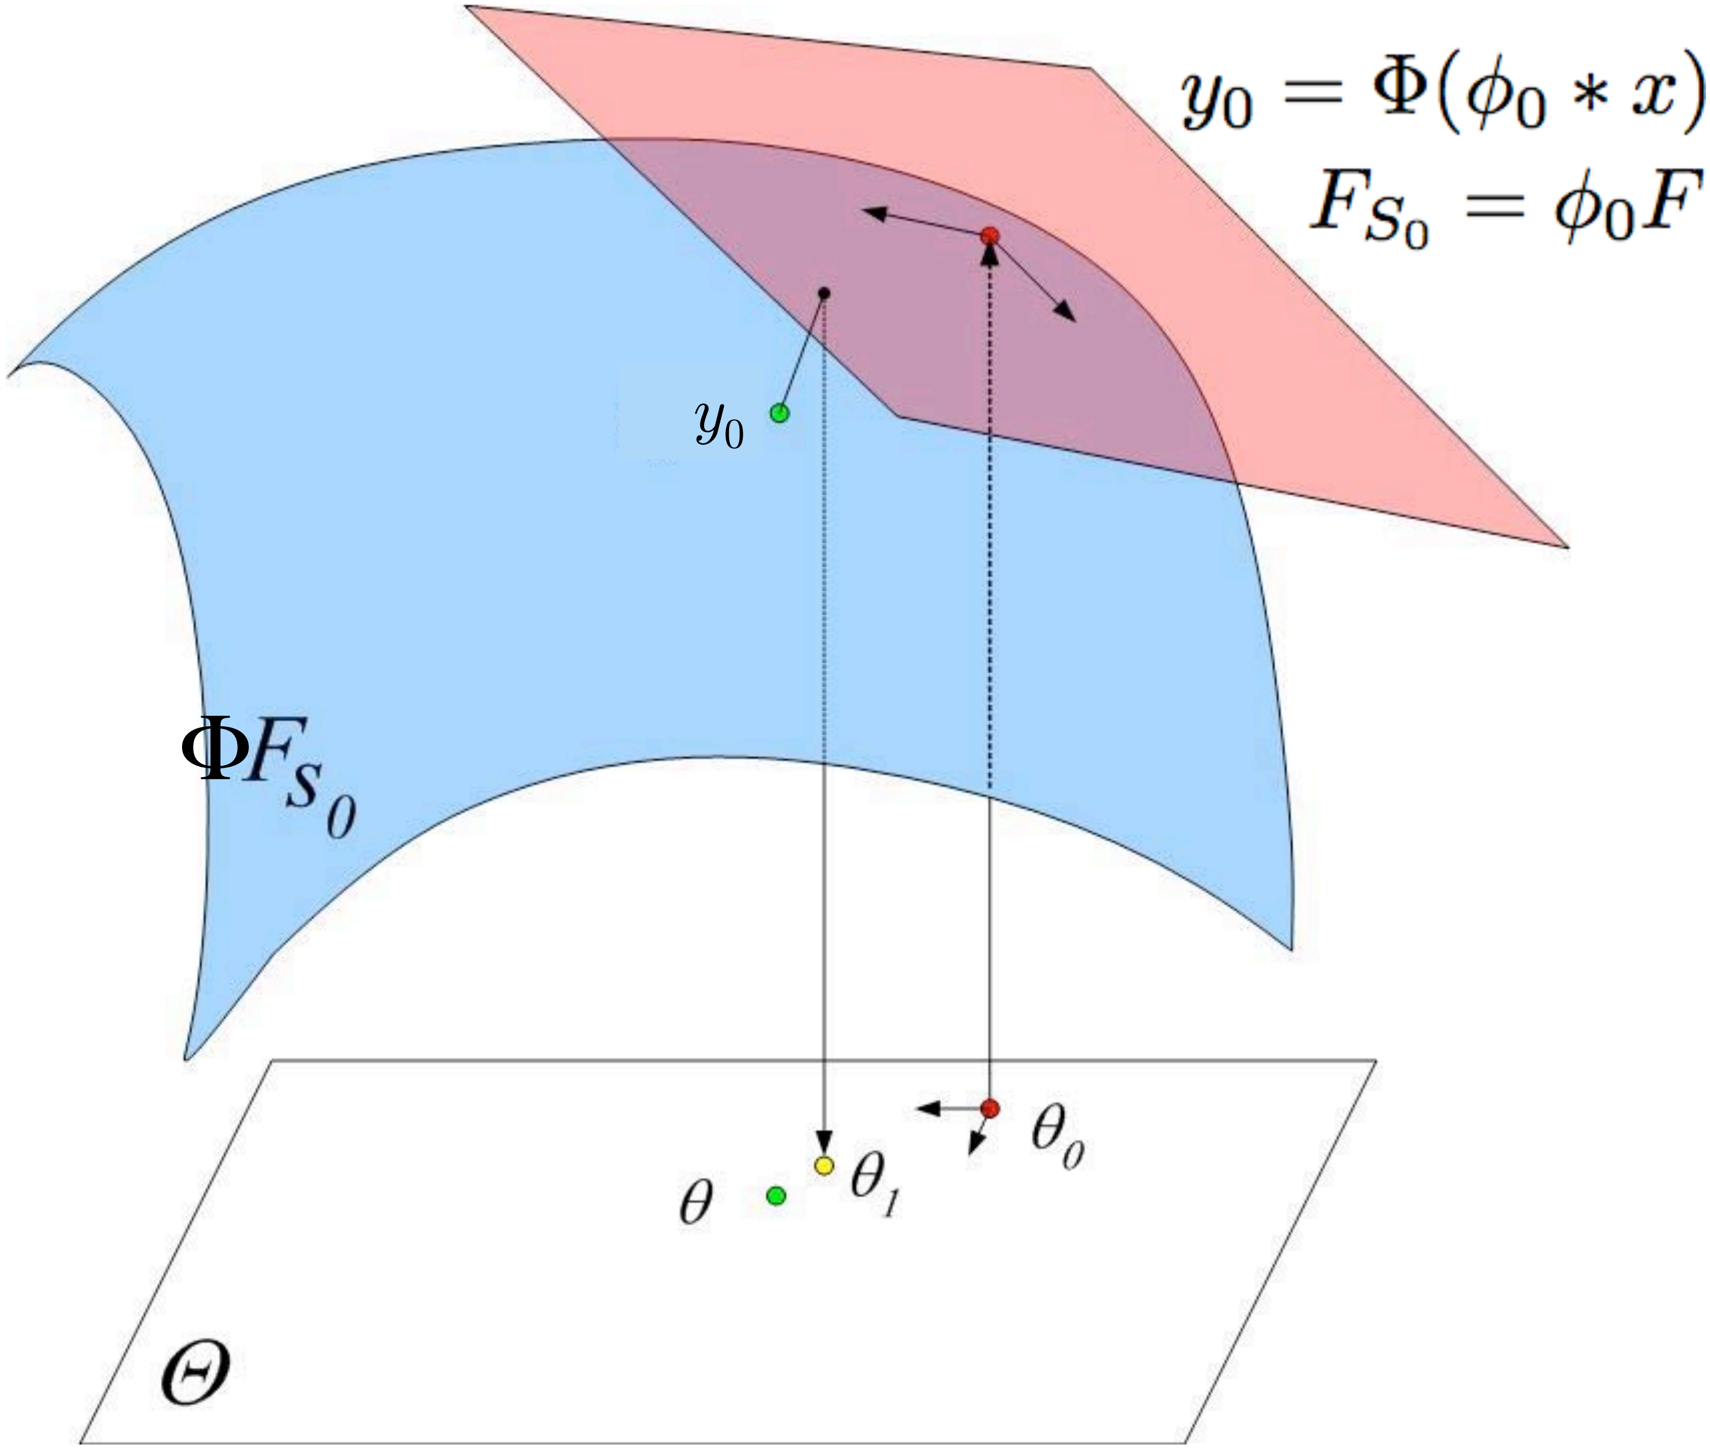

$y_0 = \Phi(\phi_0 * x)$
$F_{S_0} = \phi_0 F$
y_0
ΦF_{S_0}
θ_0
θ_1
θ
Θ

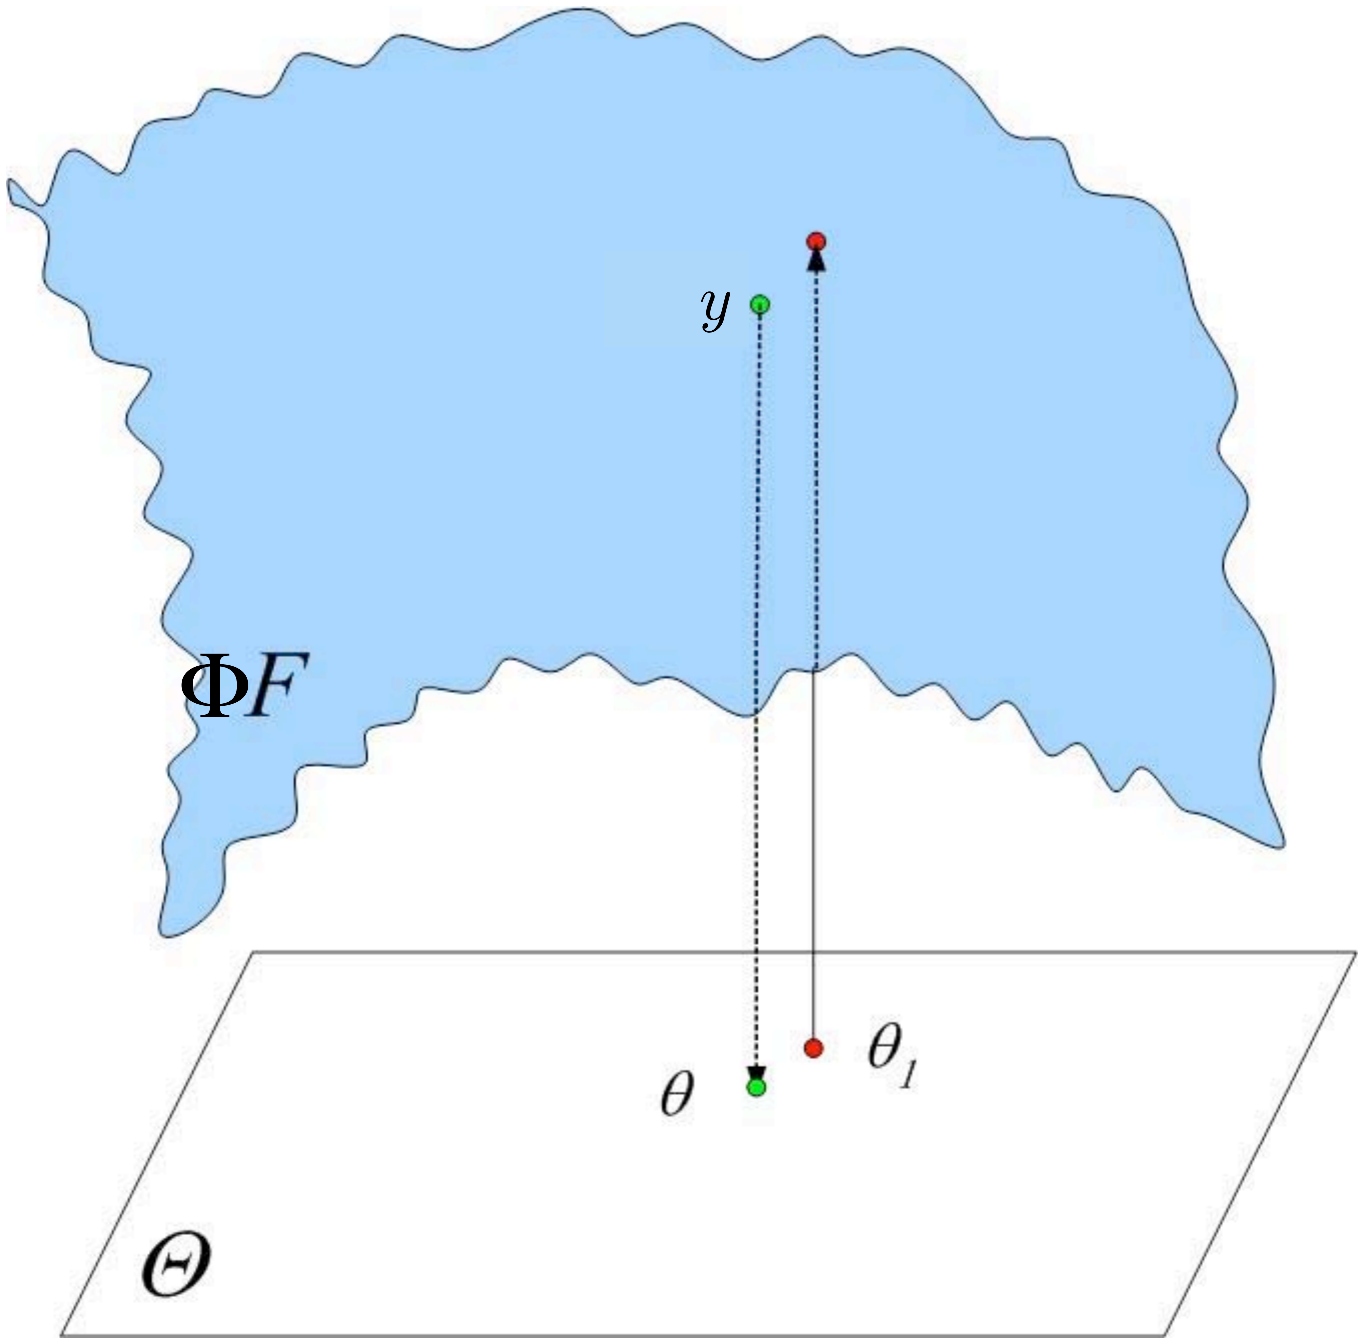

y
ΦF
θ
θ1
Θ

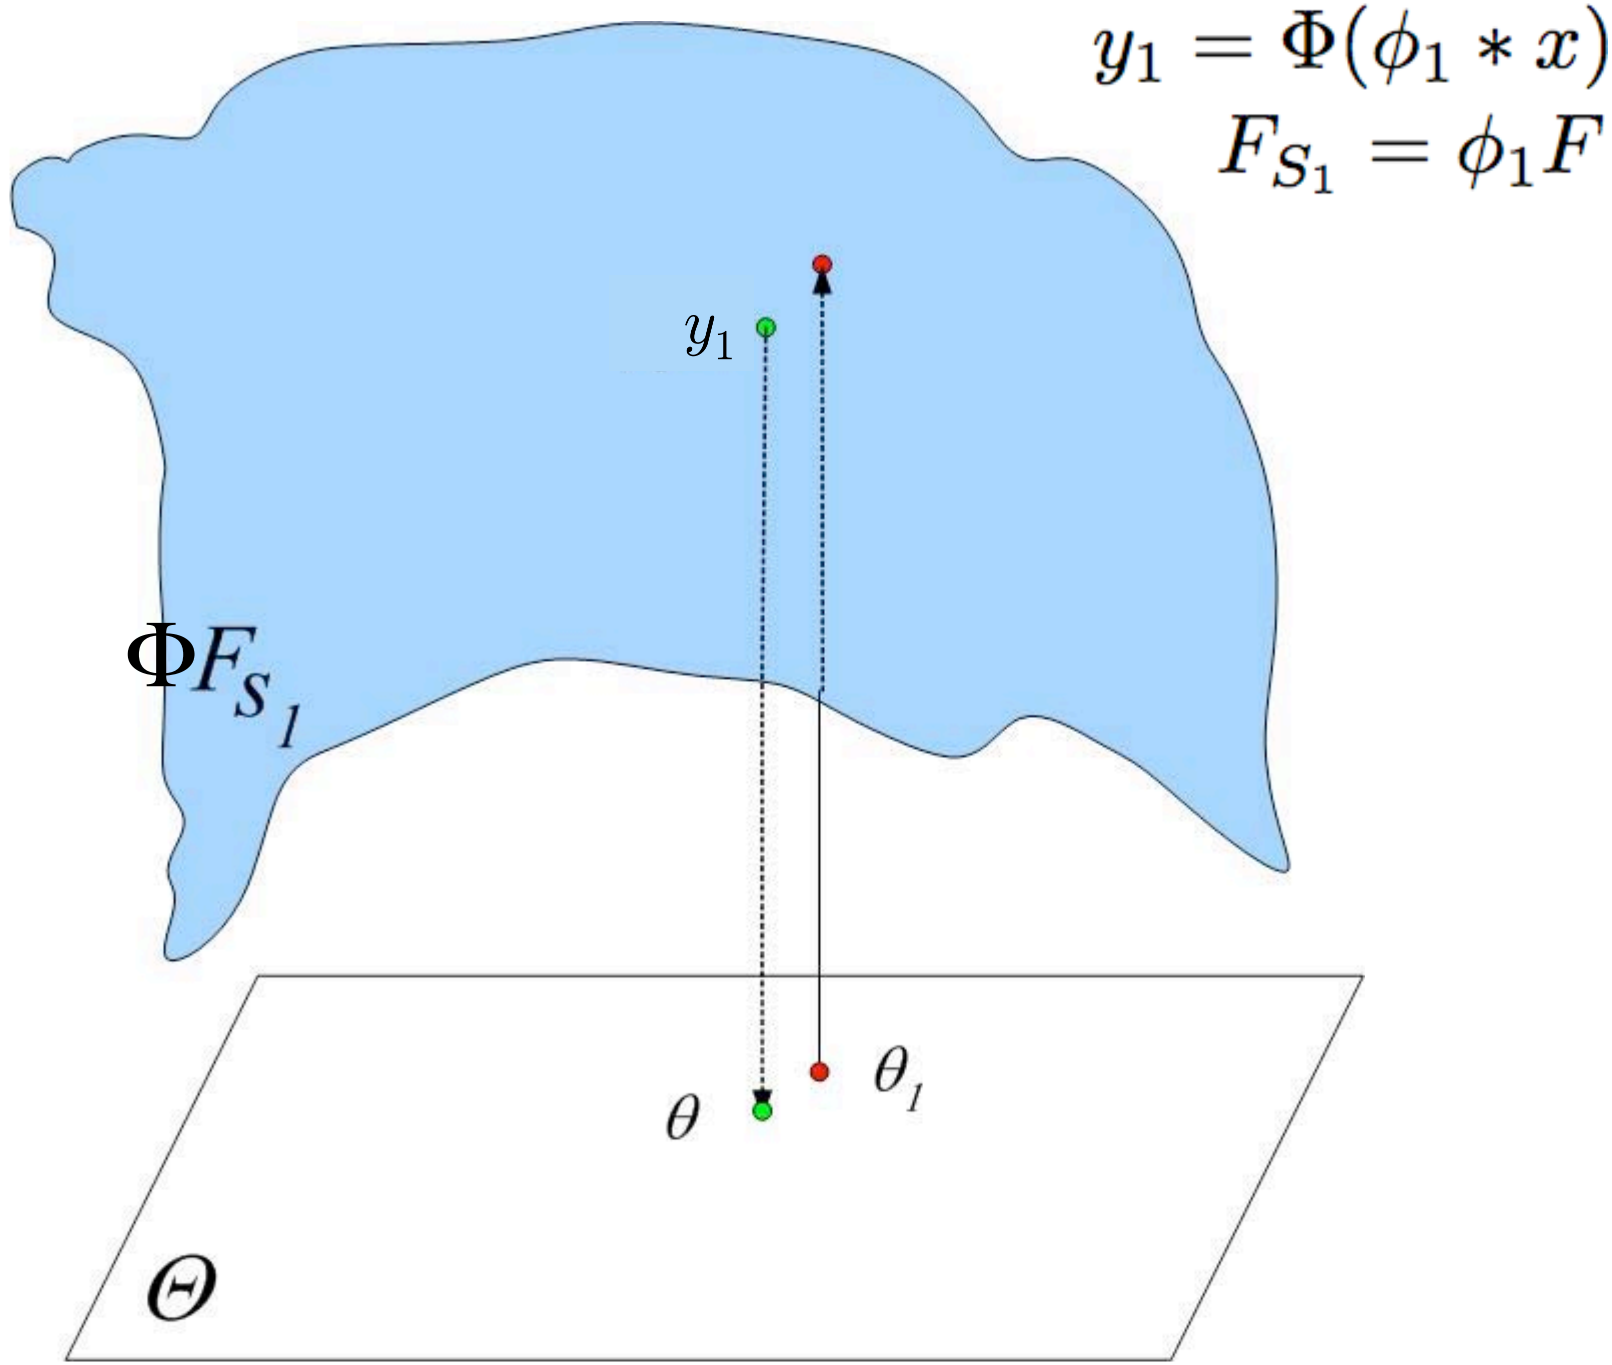
$y_1 = \Phi(\phi_1 * x)$
$F_{S_1} = \phi_1 F$
y_1
ΦF_{S_1}
θ_1
θ
Θ

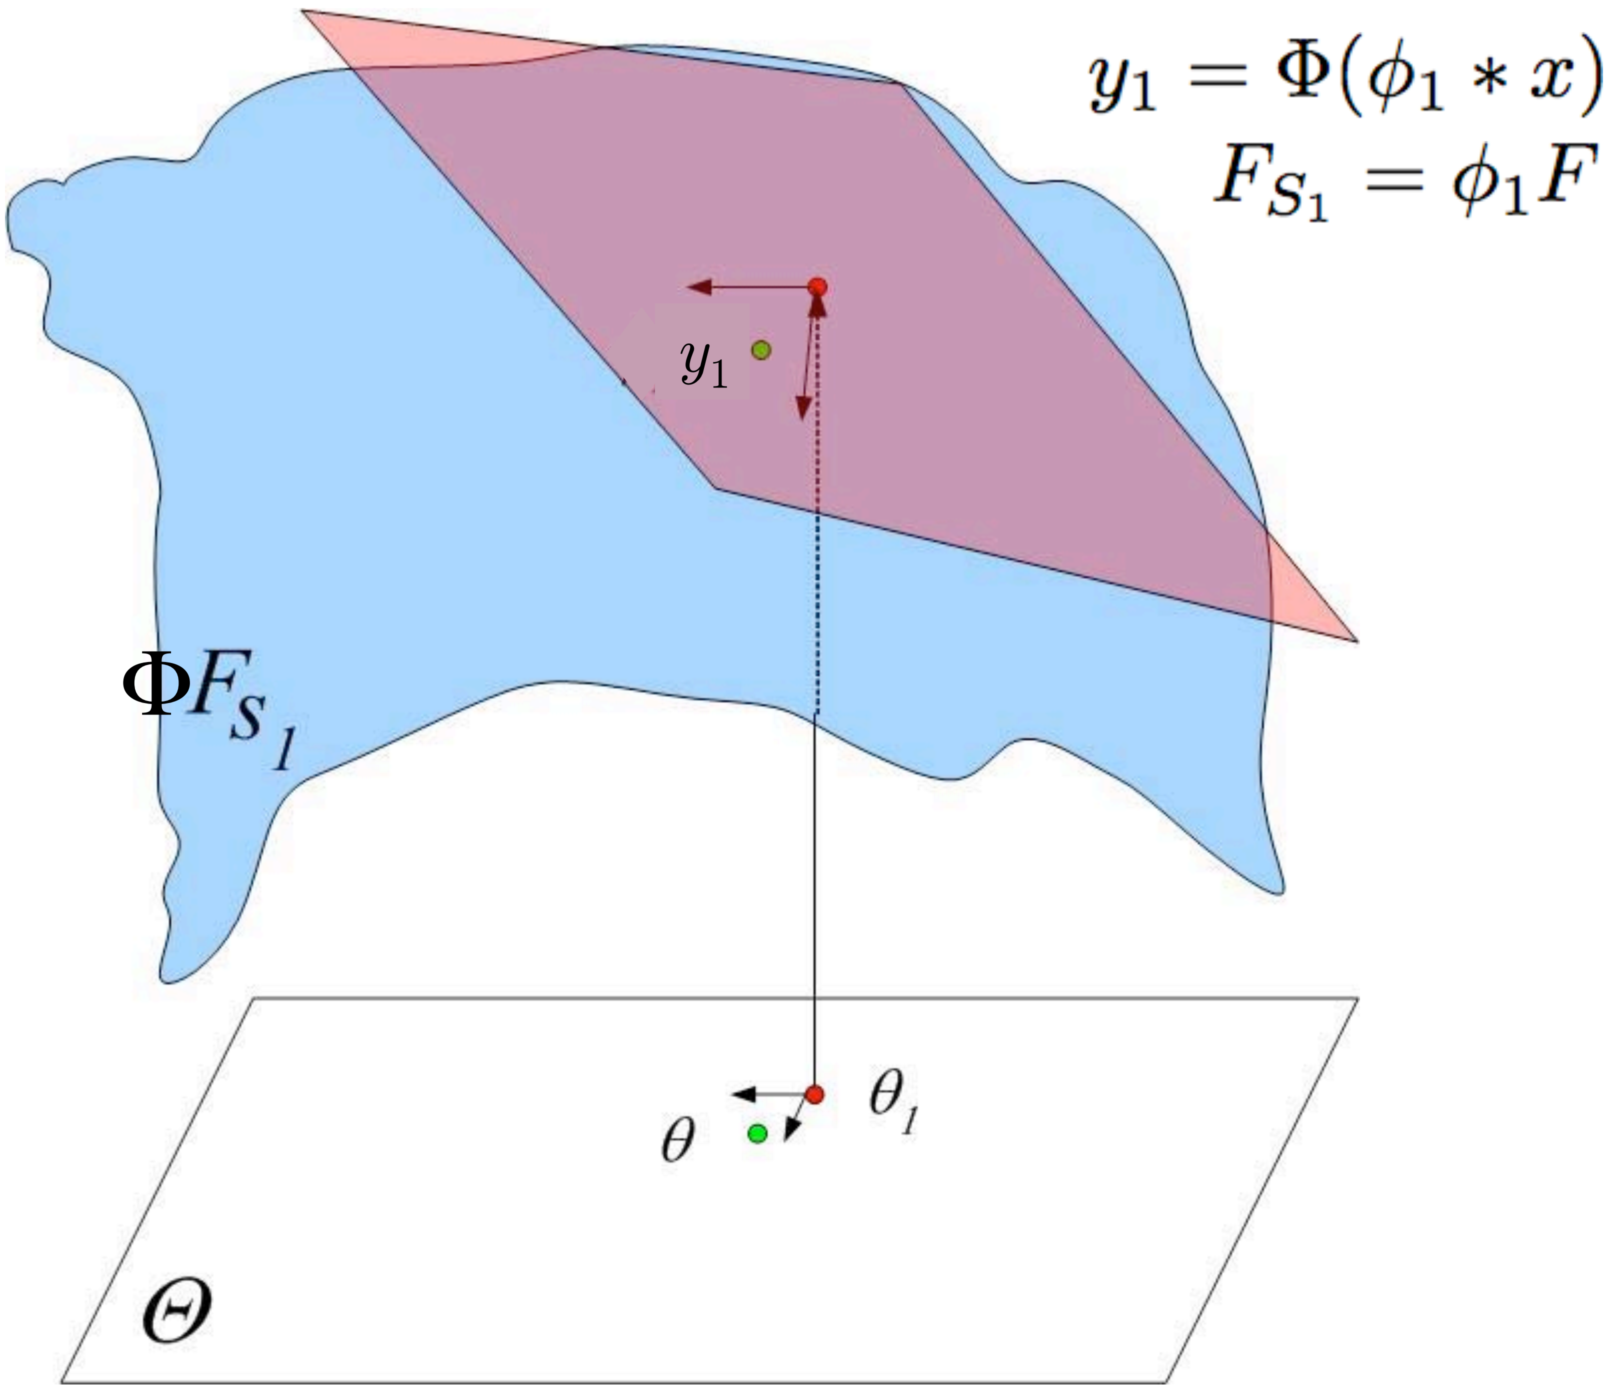
$y_1 = \Phi(\phi_1 * x)$
$F_{S_1} = \phi_1 F$
y_1
ΦF_{S_1}
θ_1
θ
Θ

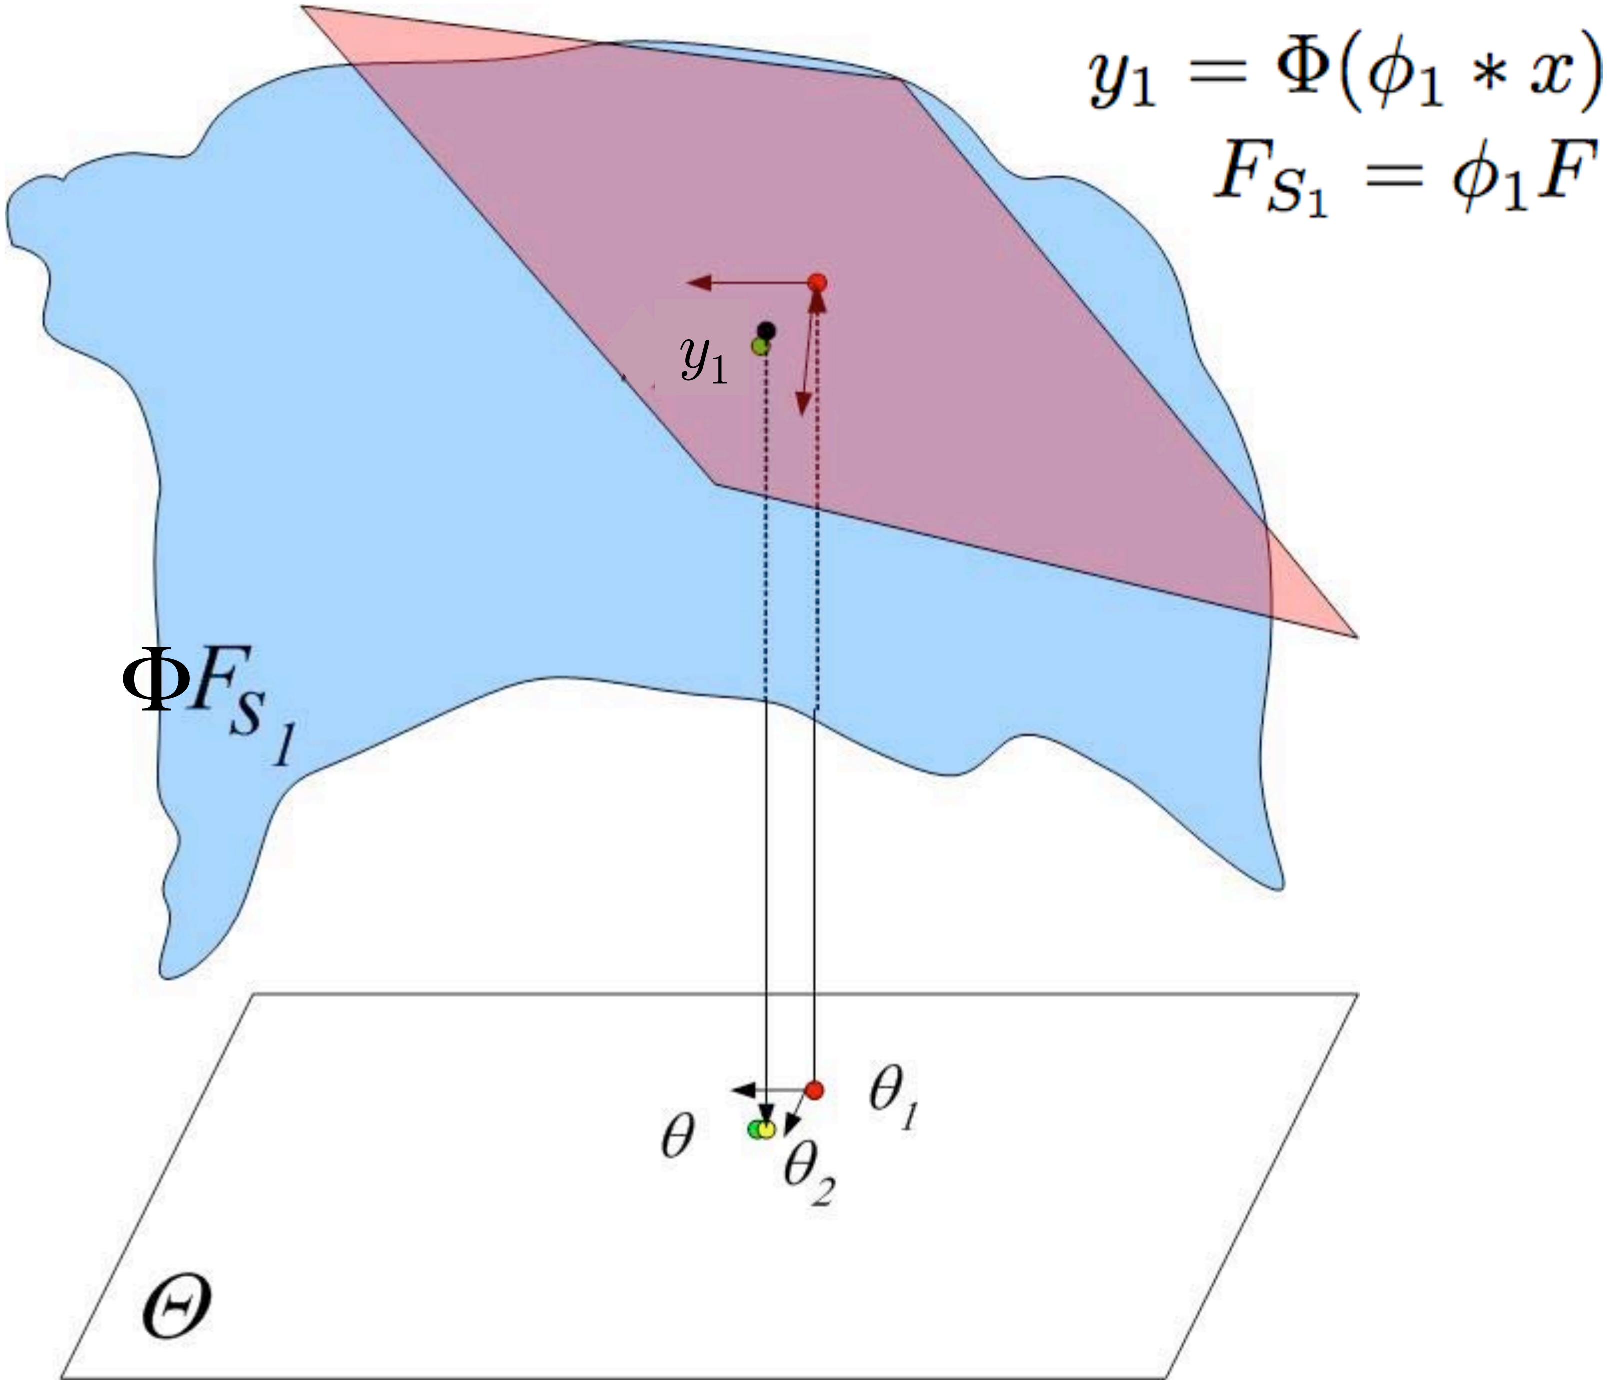

$y_1 = \Phi(\phi_1 * x)$
$F_{S_1} = \phi_1 F$
y_1
ΦF_{S_1}
θ_1
θ
θ_2
Θ

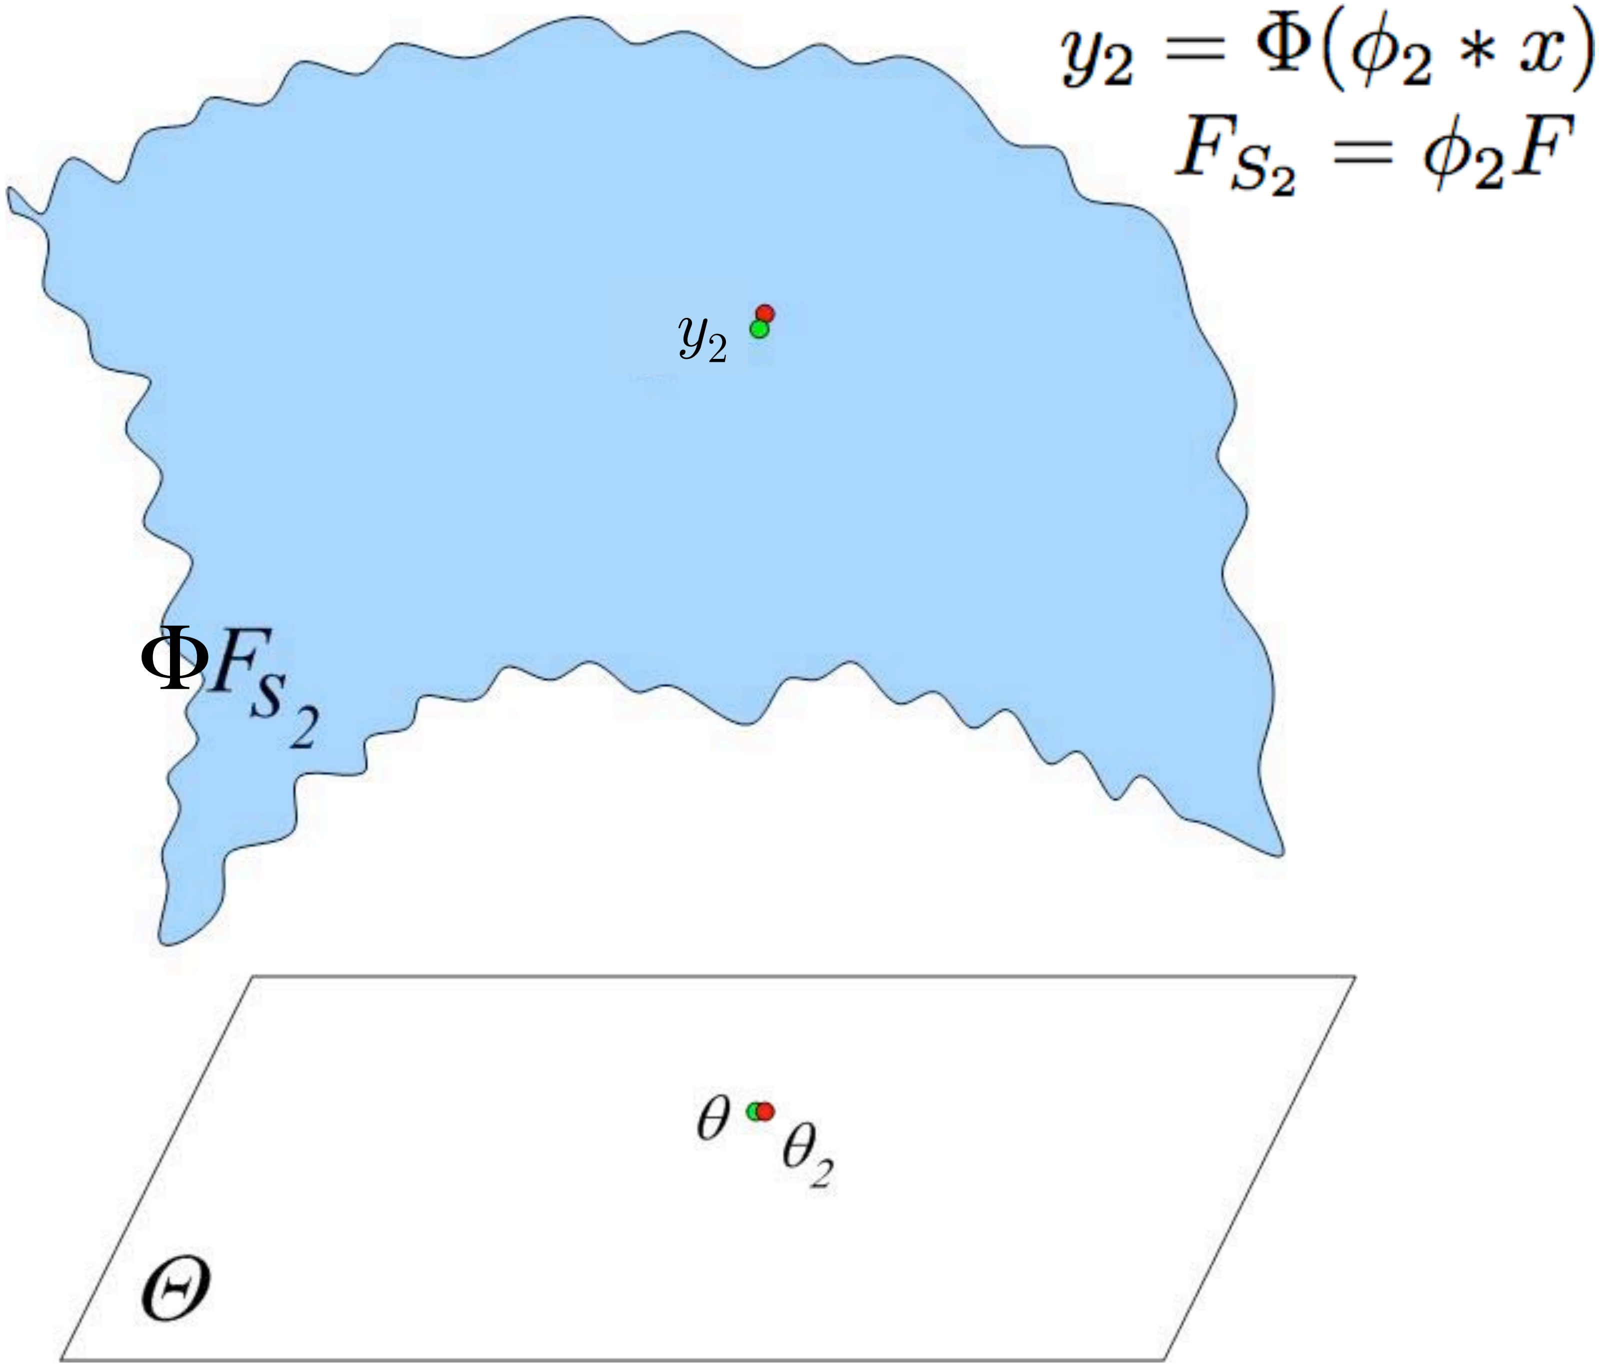

$y_2 = \Phi(\phi_2 * x)$
$F_{S_2} = \phi_2 F$
y_2
ΦF_{S_2}
θ
θ_2
Θ

Multiscale Smashed Filter

$$y_1 = \Phi \phi_1 x$$

$$y_2 = \Phi \phi_2 x$$

$$\vdots$$

$$y_P = \Phi \phi_P x$$

Multiscale Smashed Filter

$$y_1 = \Phi_1 x$$

$$y_2 = \Phi_2 x$$

$$\vdots$$

$$y_P = \Phi_P x$$

$$\Phi_i = \Phi \phi_i$$

Multiscale Smashed Filter

$$y = \begin{bmatrix} \Phi_1 \\ \Phi_2 \\ \vdots \\ \Phi_P \end{bmatrix} x$$

$$\Phi_i = \Phi \phi_i$$

Multiscale Smashed Filter

$$y = \begin{bmatrix} \Phi_1 \\ \Phi_2 \\ \vdots \\ \Phi_P \end{bmatrix} x$$

If $\mathbf{\Phi}$ has binary random entries,
x can be regularized using pixelation

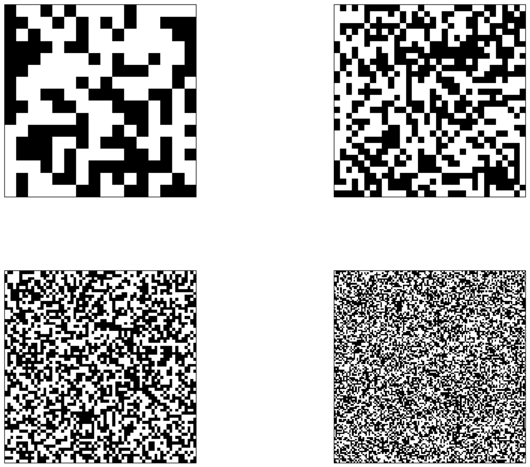

Rice Single-Pixel Camera

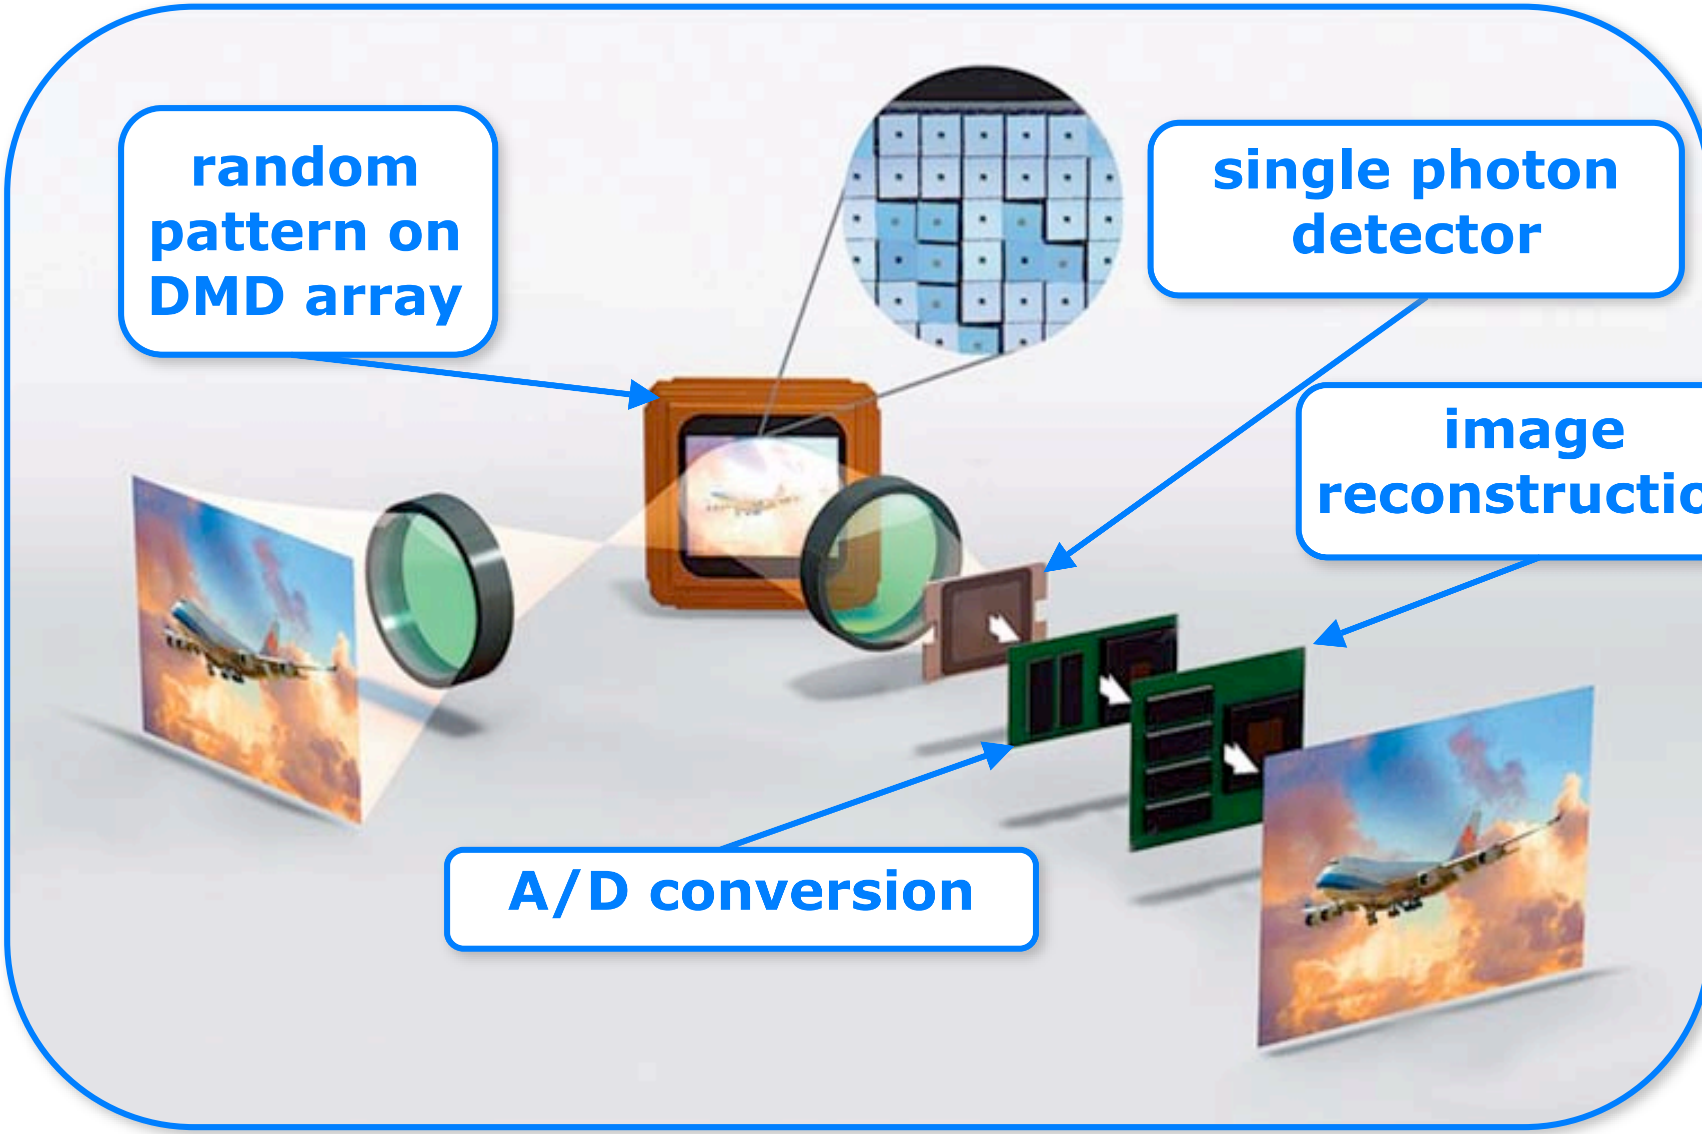

Experiments

- 3 image classes imaged using single-pixel camera
  - rotations 2^o, 4^o, …, 360^o
  - binary random measurements
  - 5 regularization kernels through pixelation(16, 8, 4, …)
- Training set for each class: CS measurements
  - estimate rotation using multiscale projections
  - identify most likely class using nearest-neighbor test

Tank

Bus

SUV

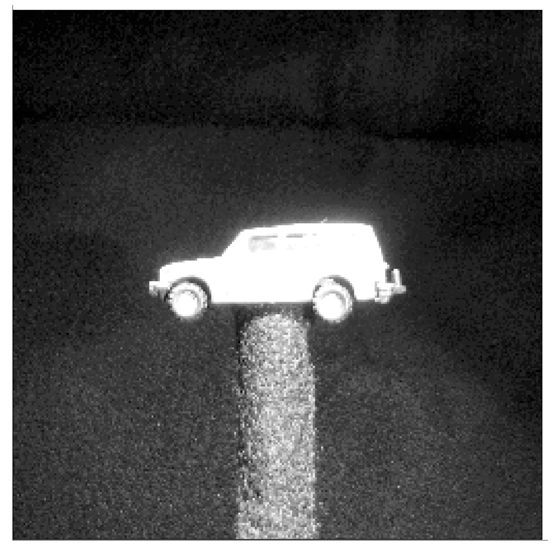

Classification results

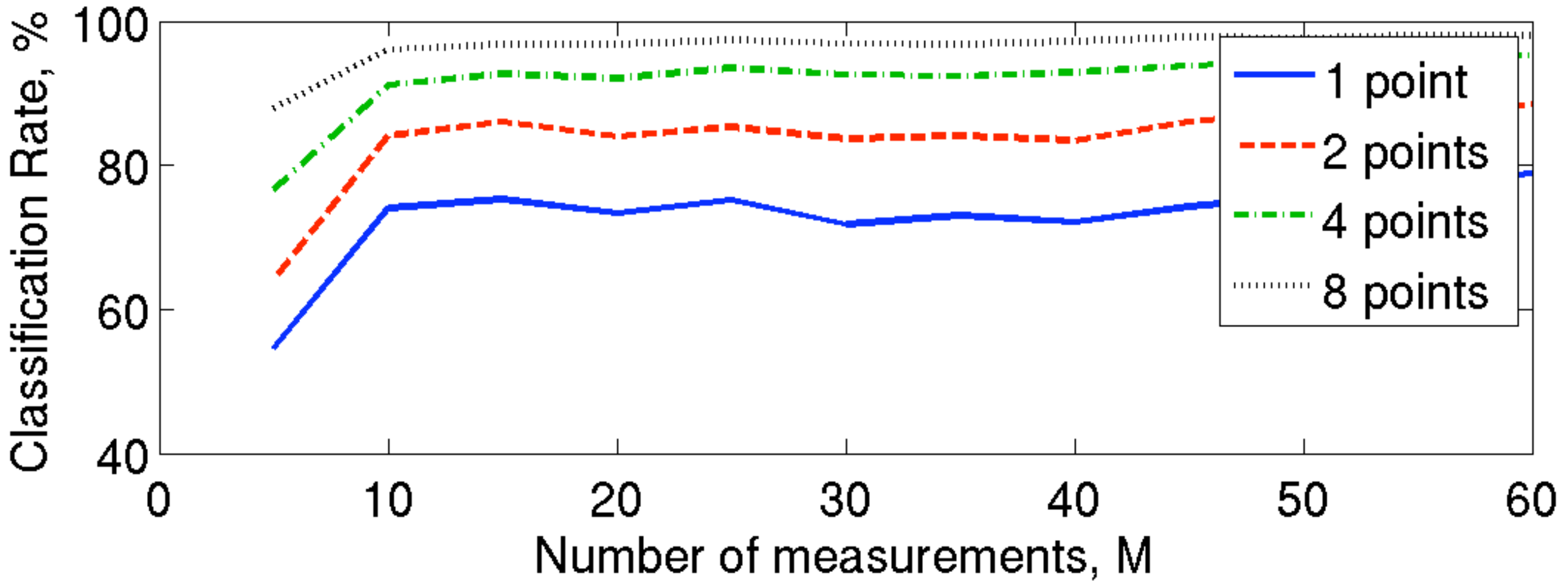

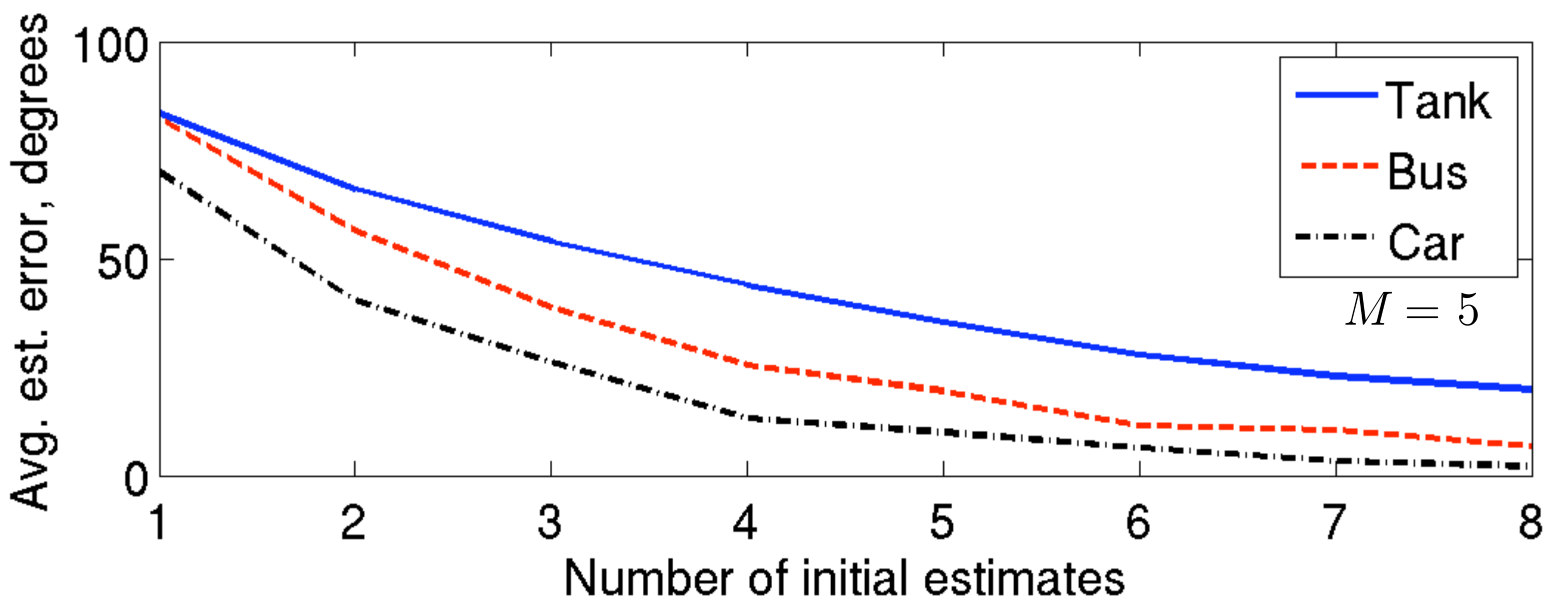

Tolerance to added noise

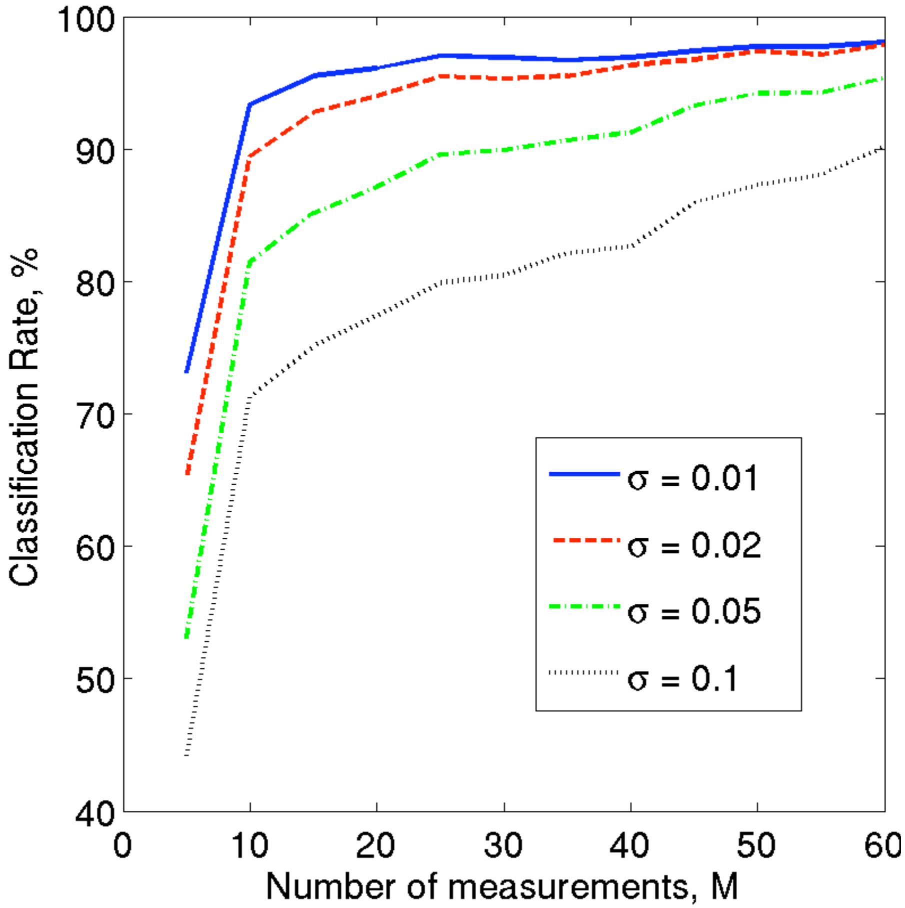

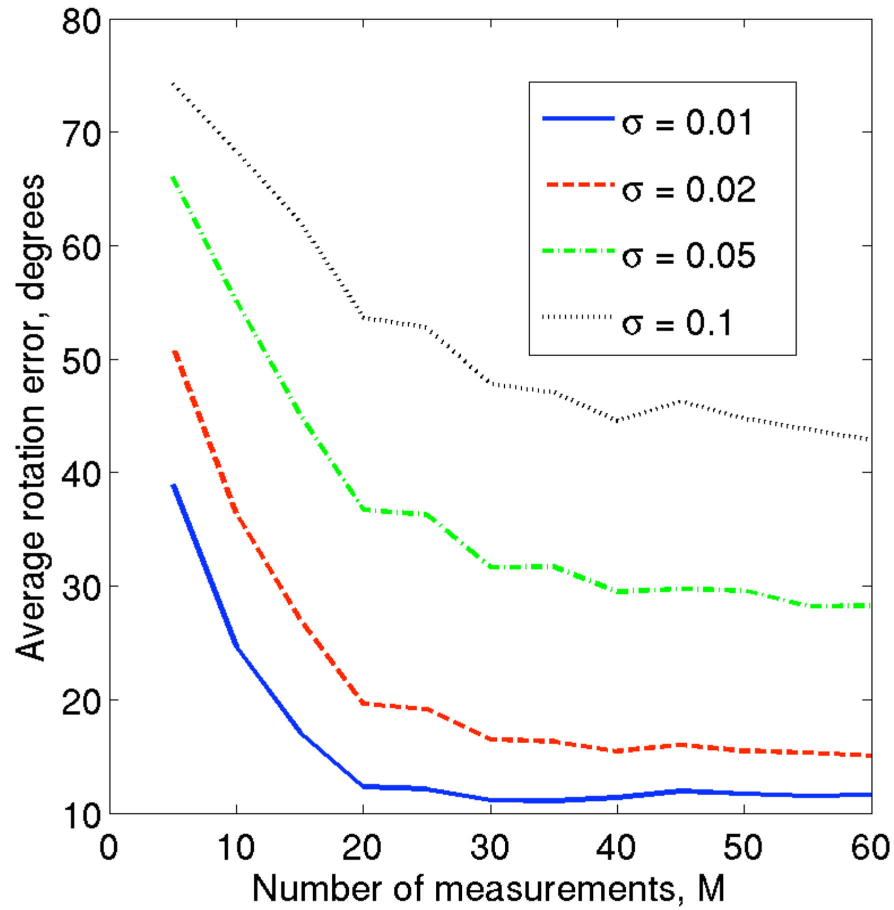

- 8 initial estimates
- Higher noise requires more measurements for accurate parameter estimation
- Accurate classification requires reliable parameter estimation

Conclusions

- *Multiscale Smashed Filter*
  - efficiently exploits compressive measurements
  - Reduced computational burden
  - broadly applicable
  - effective for image classification when combined with single-pixel camera

- Current work:
  - extension to support vector machines, other algorithms
  - noise analysis (signal dependent noise)
  - collaborative compressive classification
  - compressive signal processing

Rice University DSP

dsp.rice.edu/cs