A Globally Optimal Data-driven Approach for Image Distortion Estimation

Yuandong Tian Srinivasa Narasimhan

THE ROBOTICS INSTITUTE

Carnegie Mellon University

Sponsors: NSF, ONR

Distortions in the real world

Water fluctuation

Turbulence

Cloth deformation

Optical scanning of text

Problem statement

Template $T = I_0$

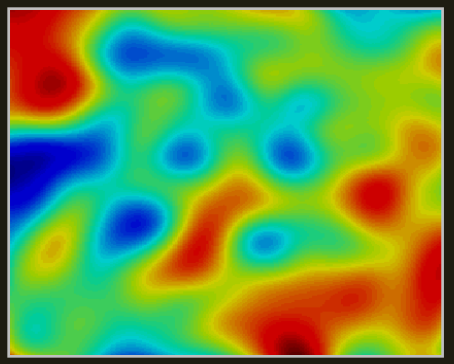

Dense Warping Field
$W(x; p)$

Distorted image I_p

Problem statement

Template $T = I_0$

Distorted image I_p

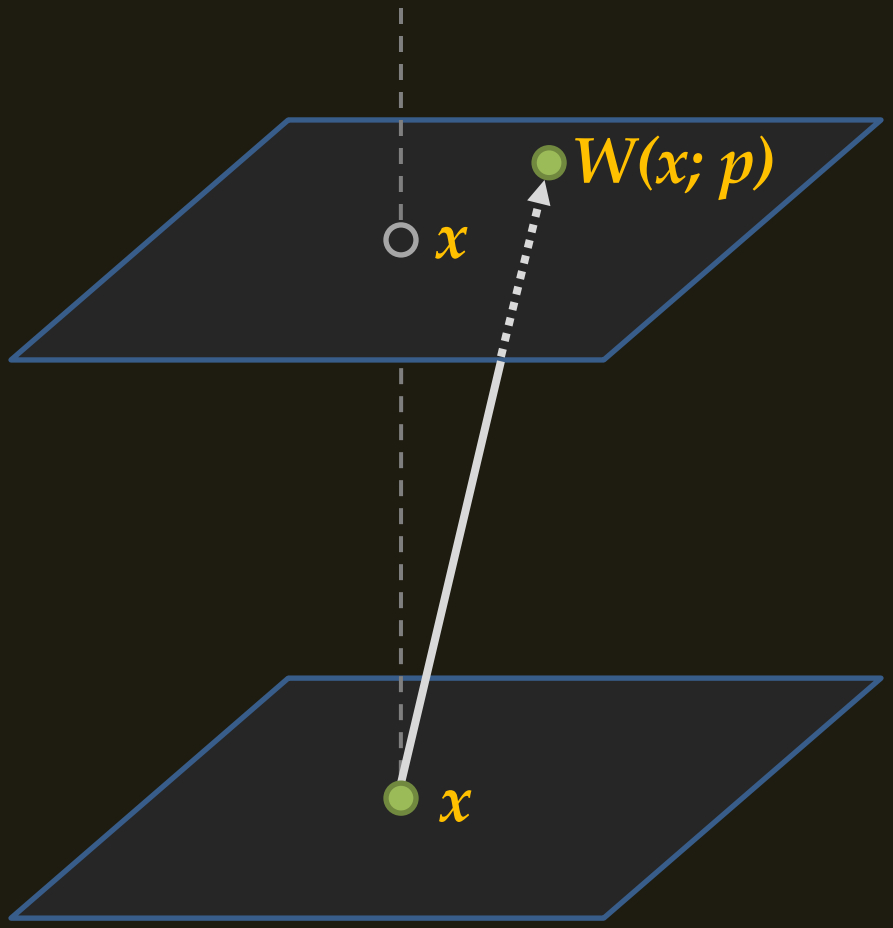

Distortion model

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

Bases *Parameters*

Choice of Bases:

Closed-form (e.g. Affine)

From Physical Simulation

$$\frac{\partial^2 h(x,t)}{\partial t^2} = c^2 \nabla^2 h(x,t)$$

From Measurement

Related work

- Generative Approach

Local Minima

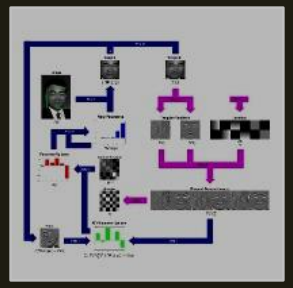

Lucas-Kanade [1981]

AAM [Cootes et al, 2001]

- Discriminative Approach

Exponential #samples

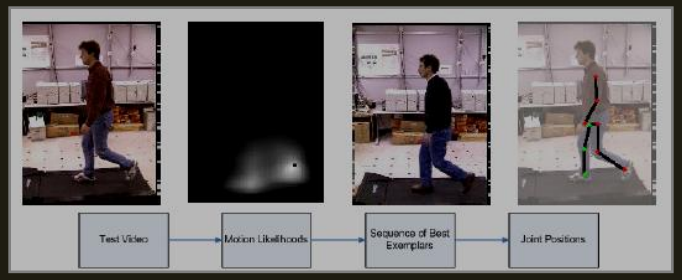

Exemplar-based [Fathi et al, 2007]

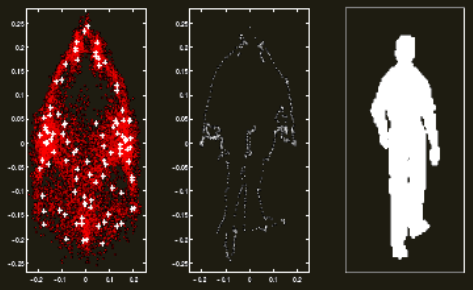

RVM [Agarwal et al, 2004]

Our algorithm overcomes both problems

Spaceship returning to the Earth

Spaceship

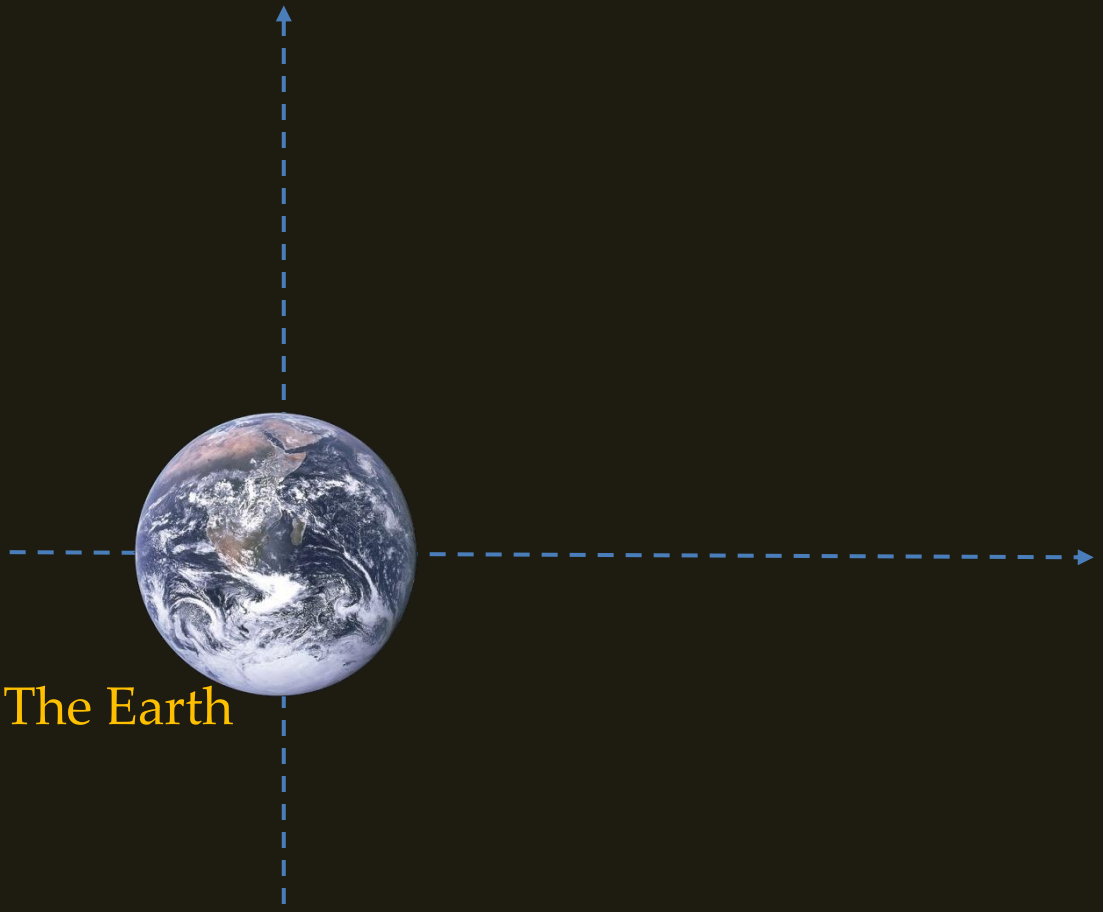

The Earth

Spaceship returning to the Earth

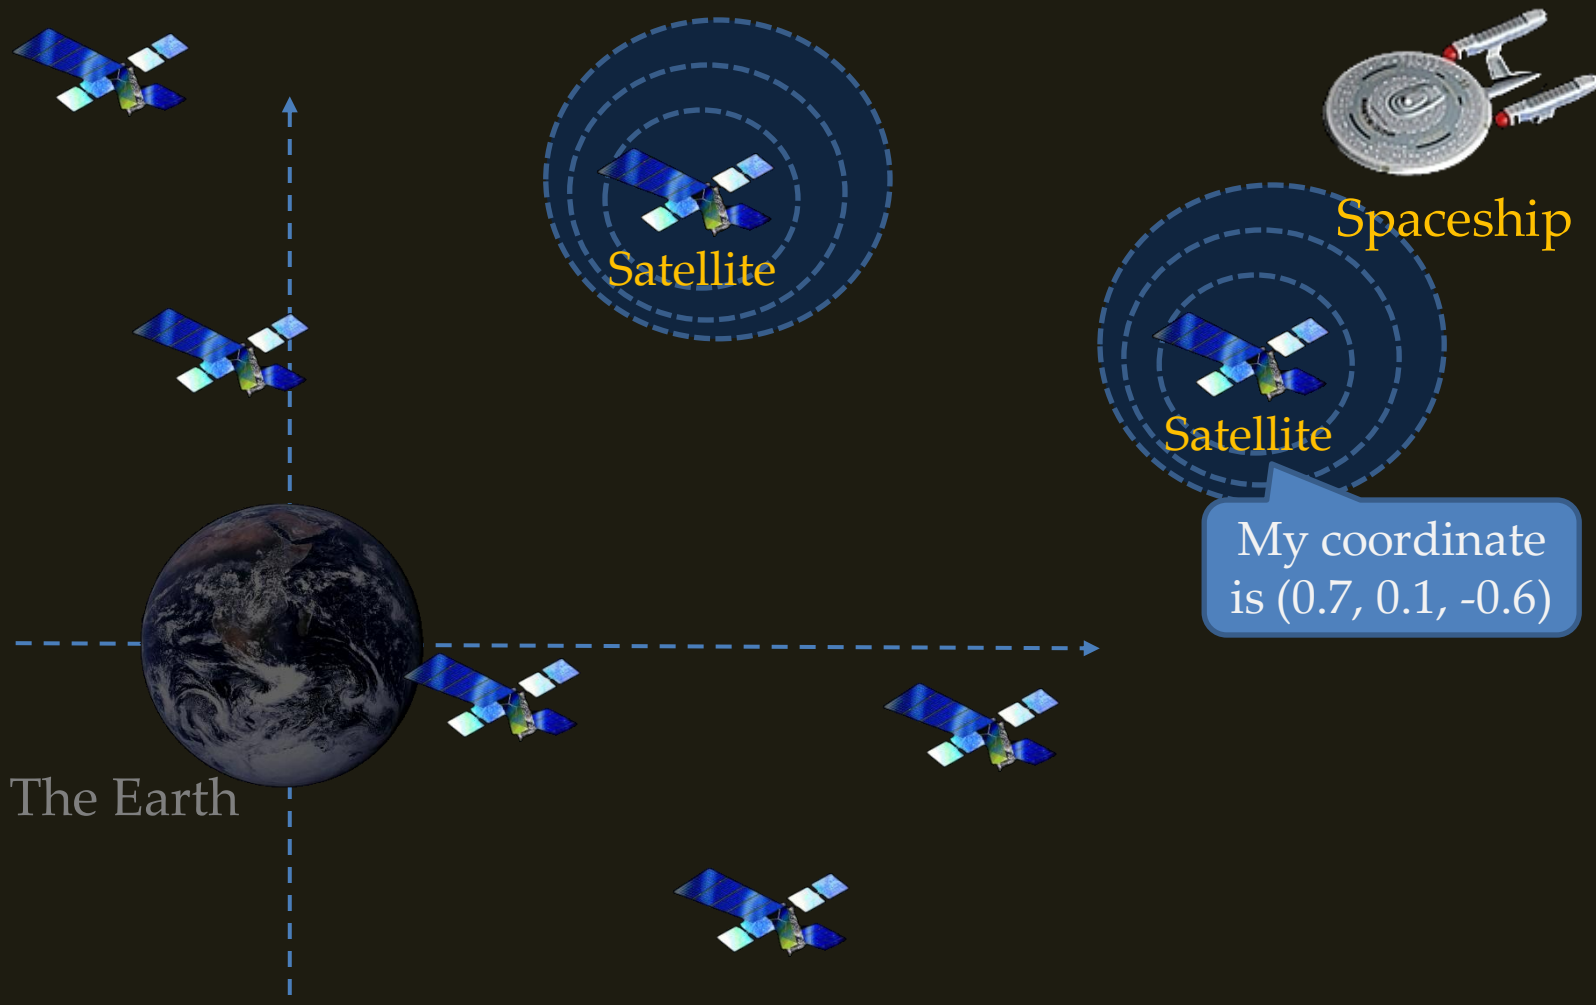

Spaceship returning to the Earth

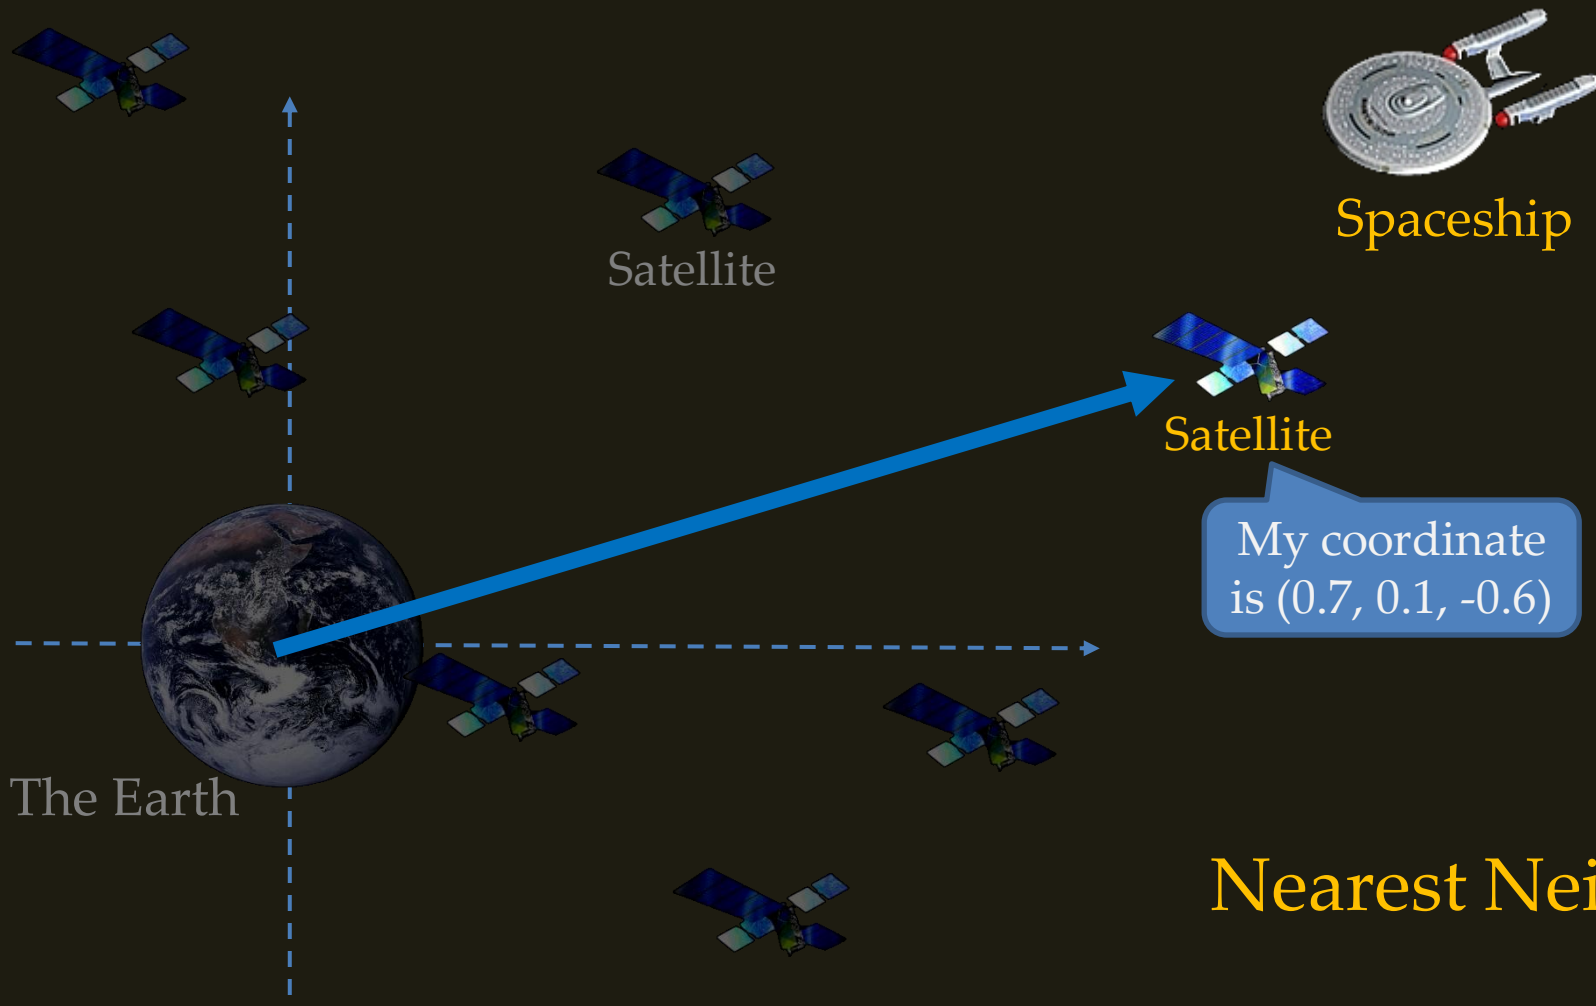

Nearest Neighbor

Spaceship returning to the Earth

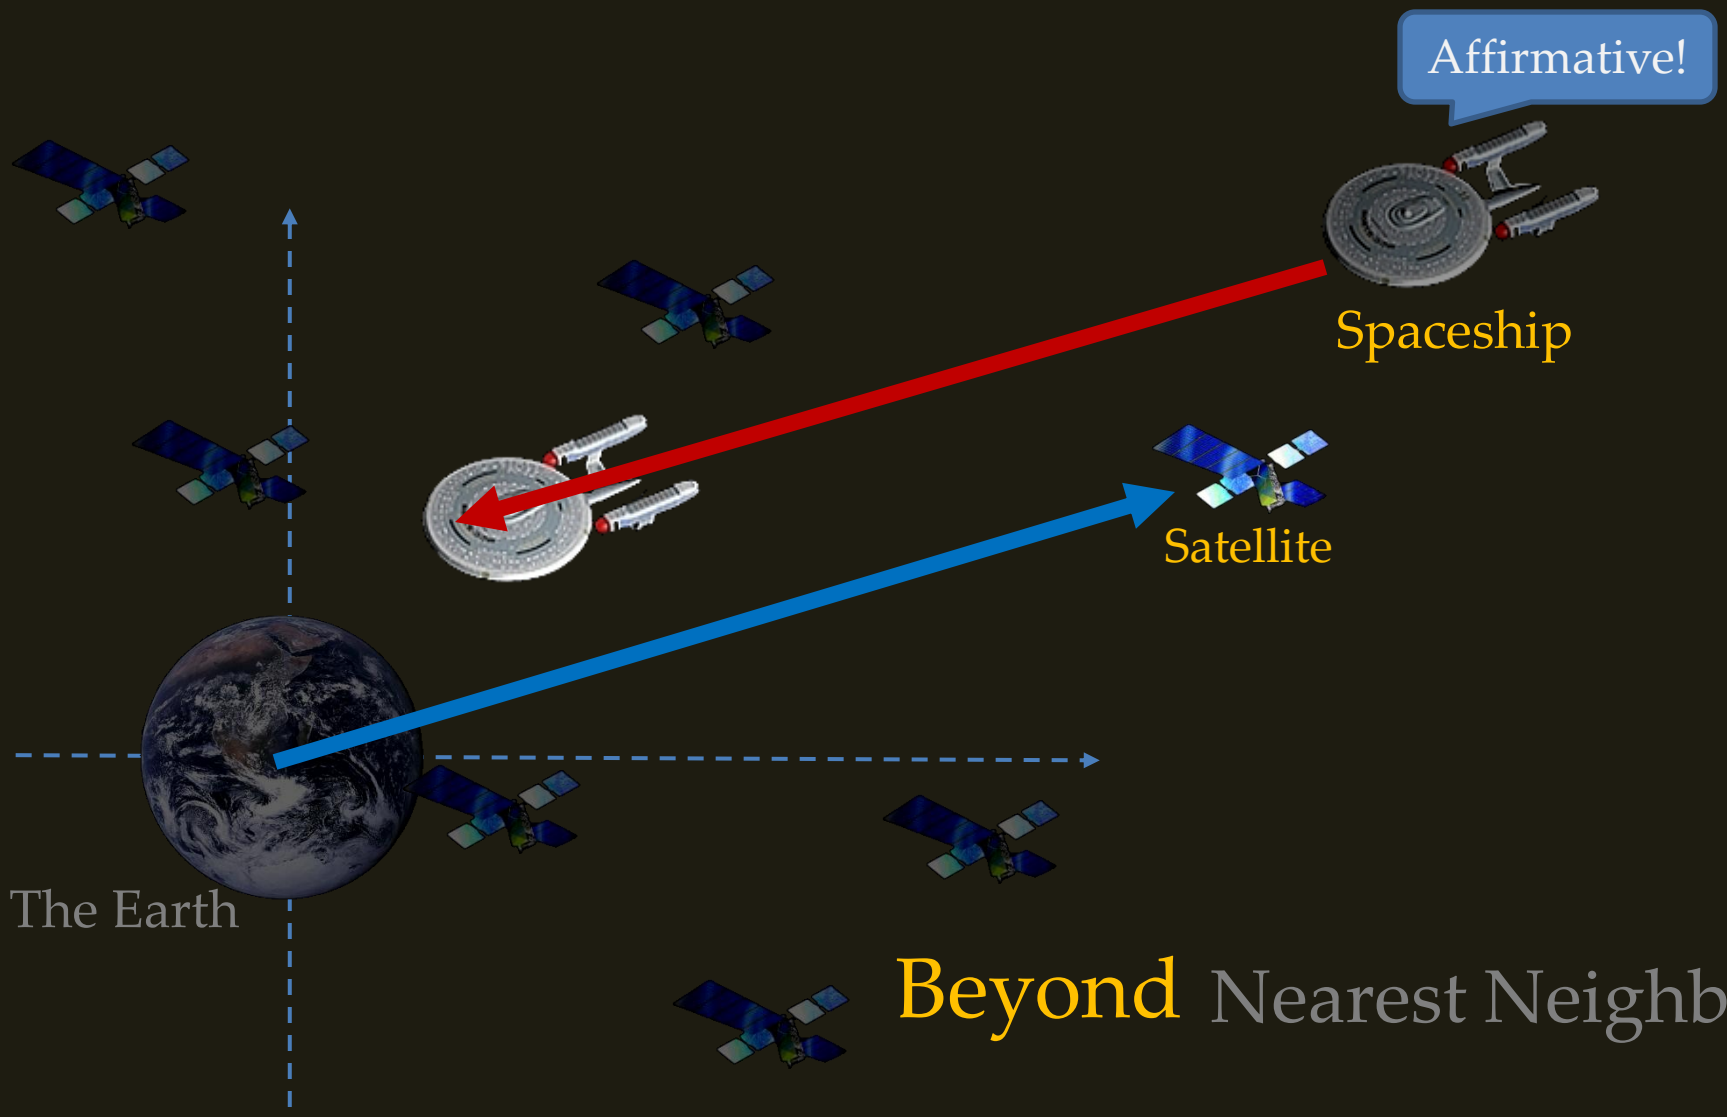

Beyond Nearest Neighbor

Spaceship returning to the Earth

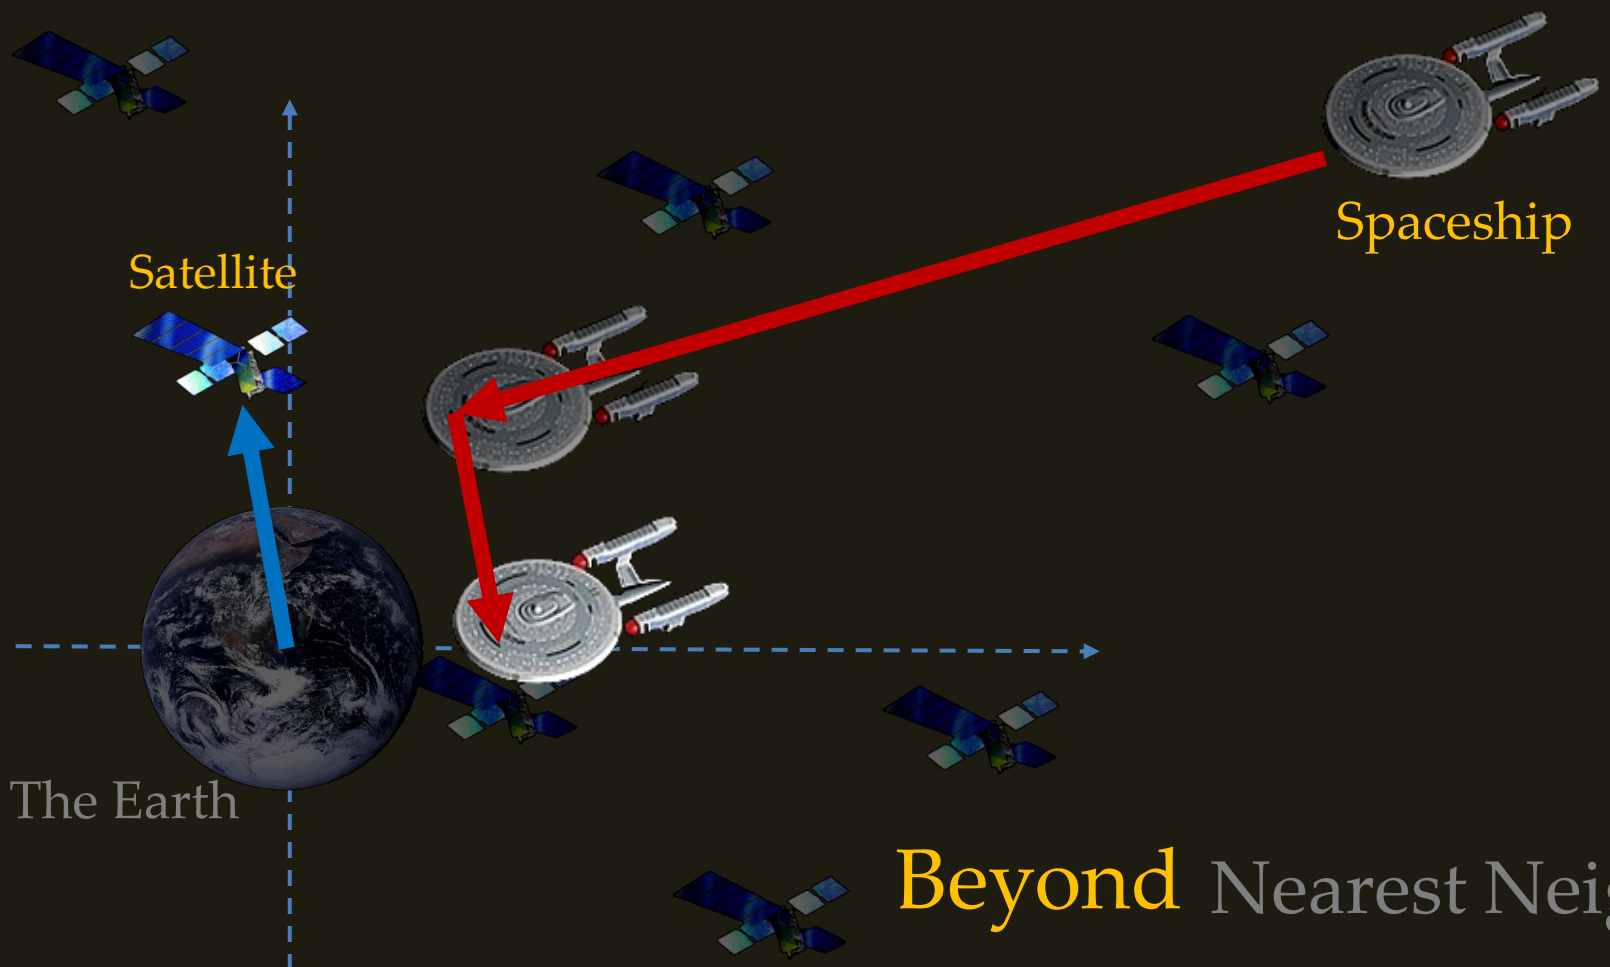

Spaceship returning to the Earth

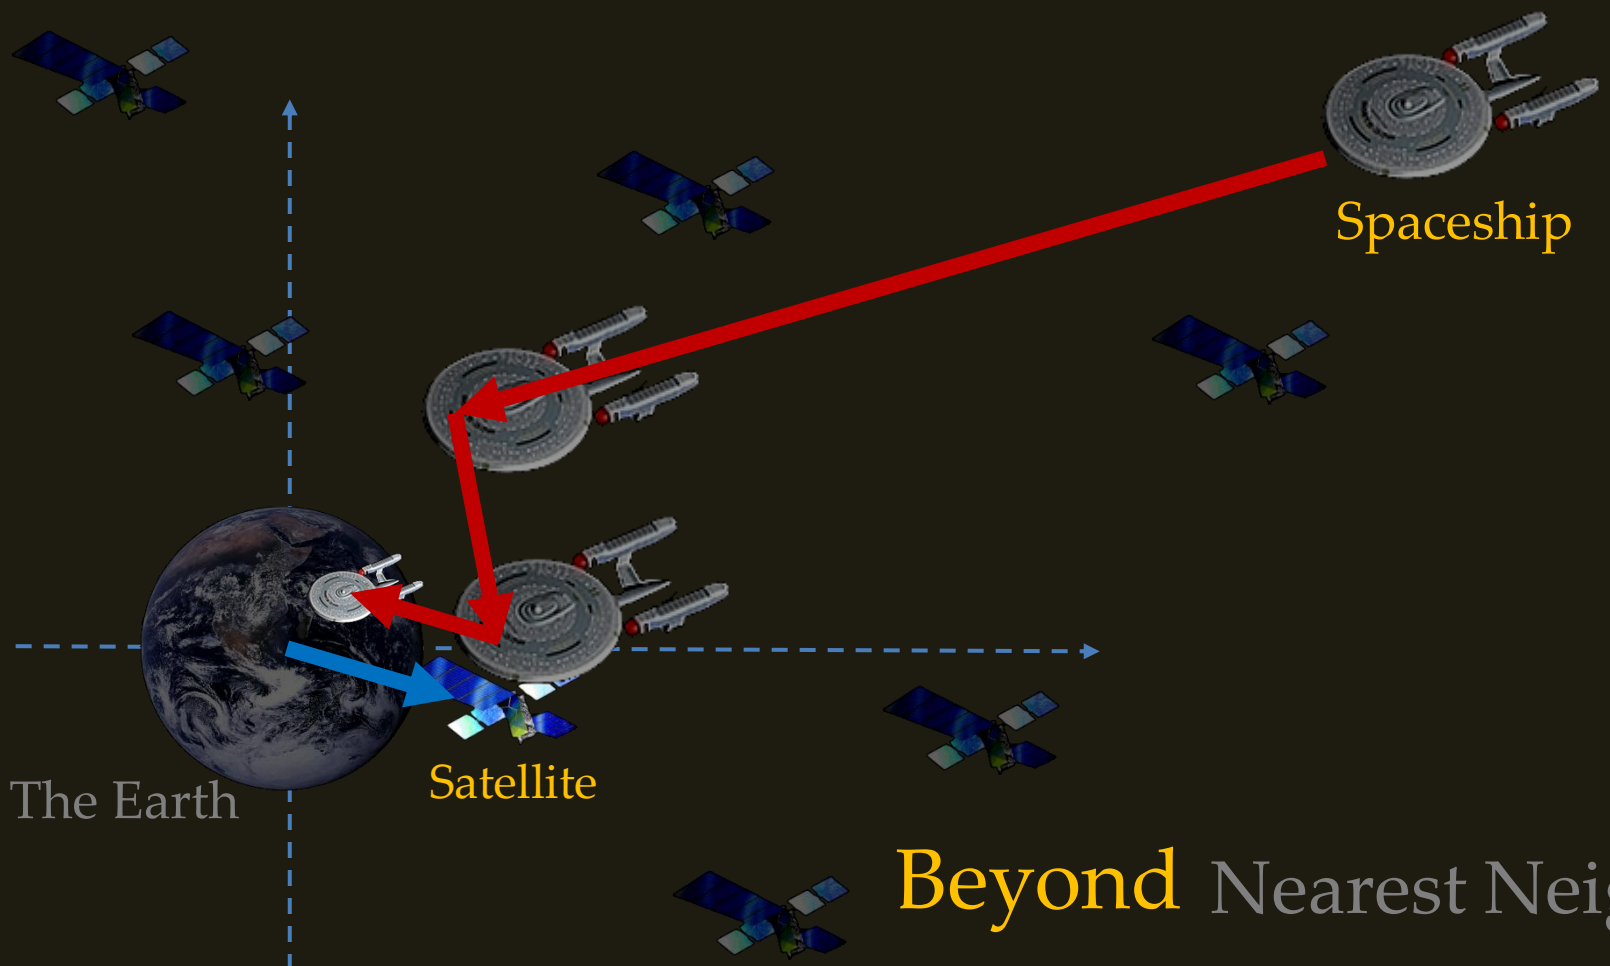

Beyond Nearest Neighbor

Similar operations for images

Parameter space

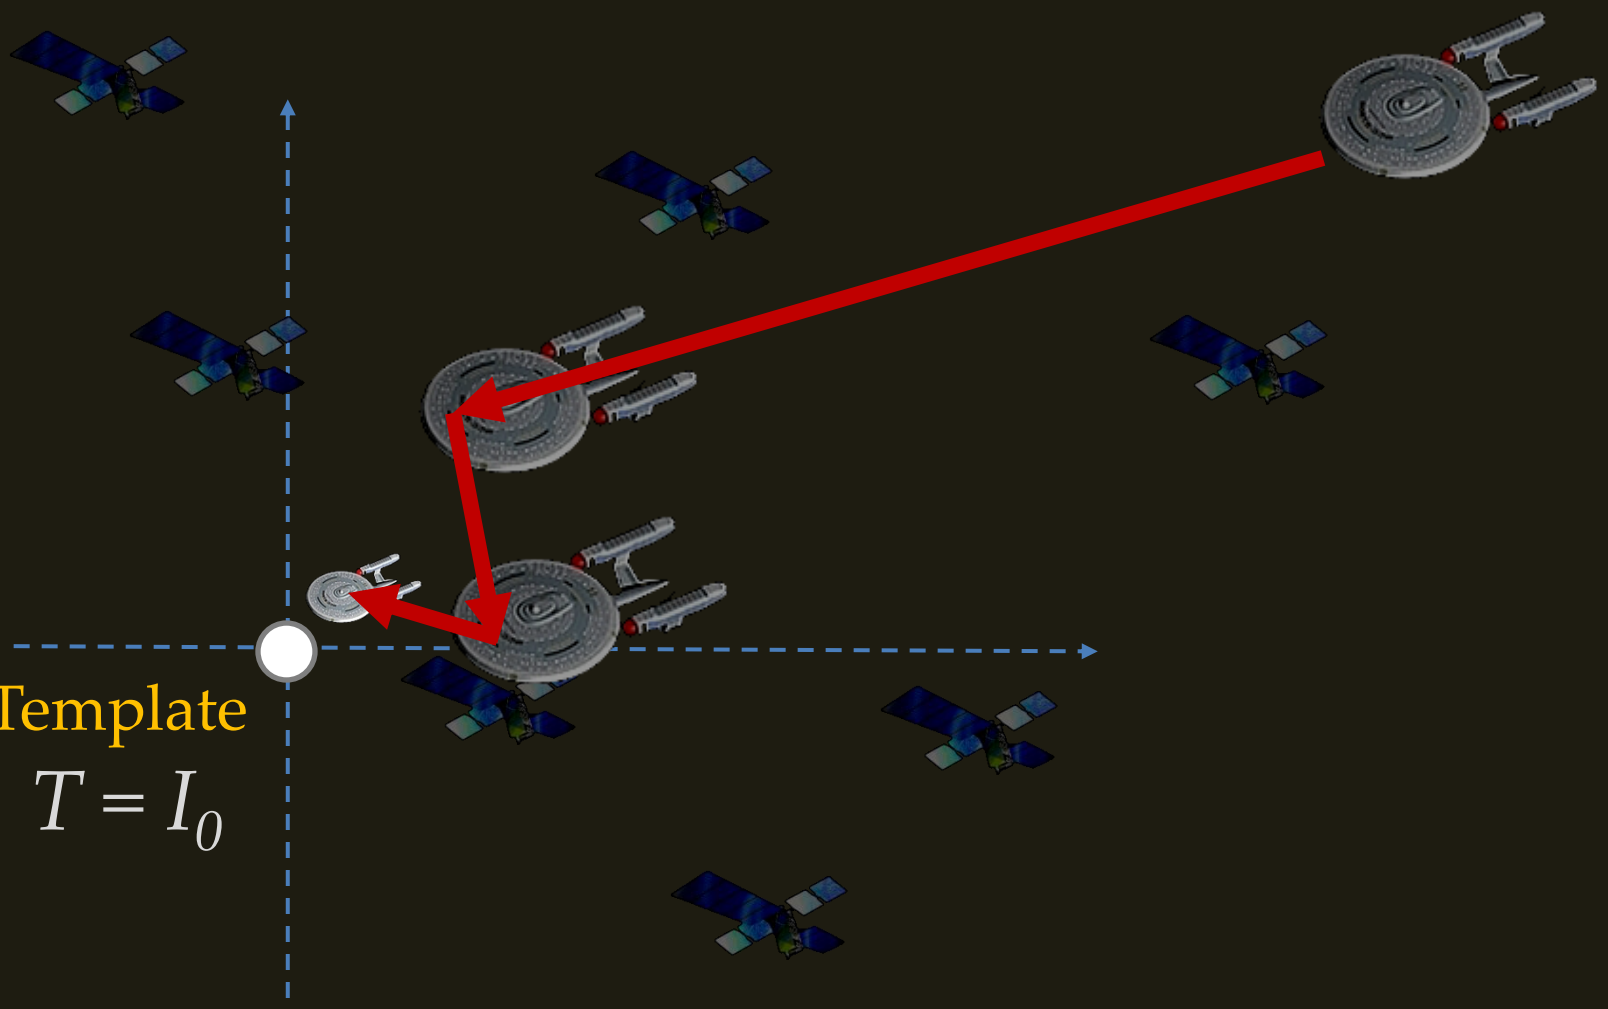

Similar operations for images

Parameter space

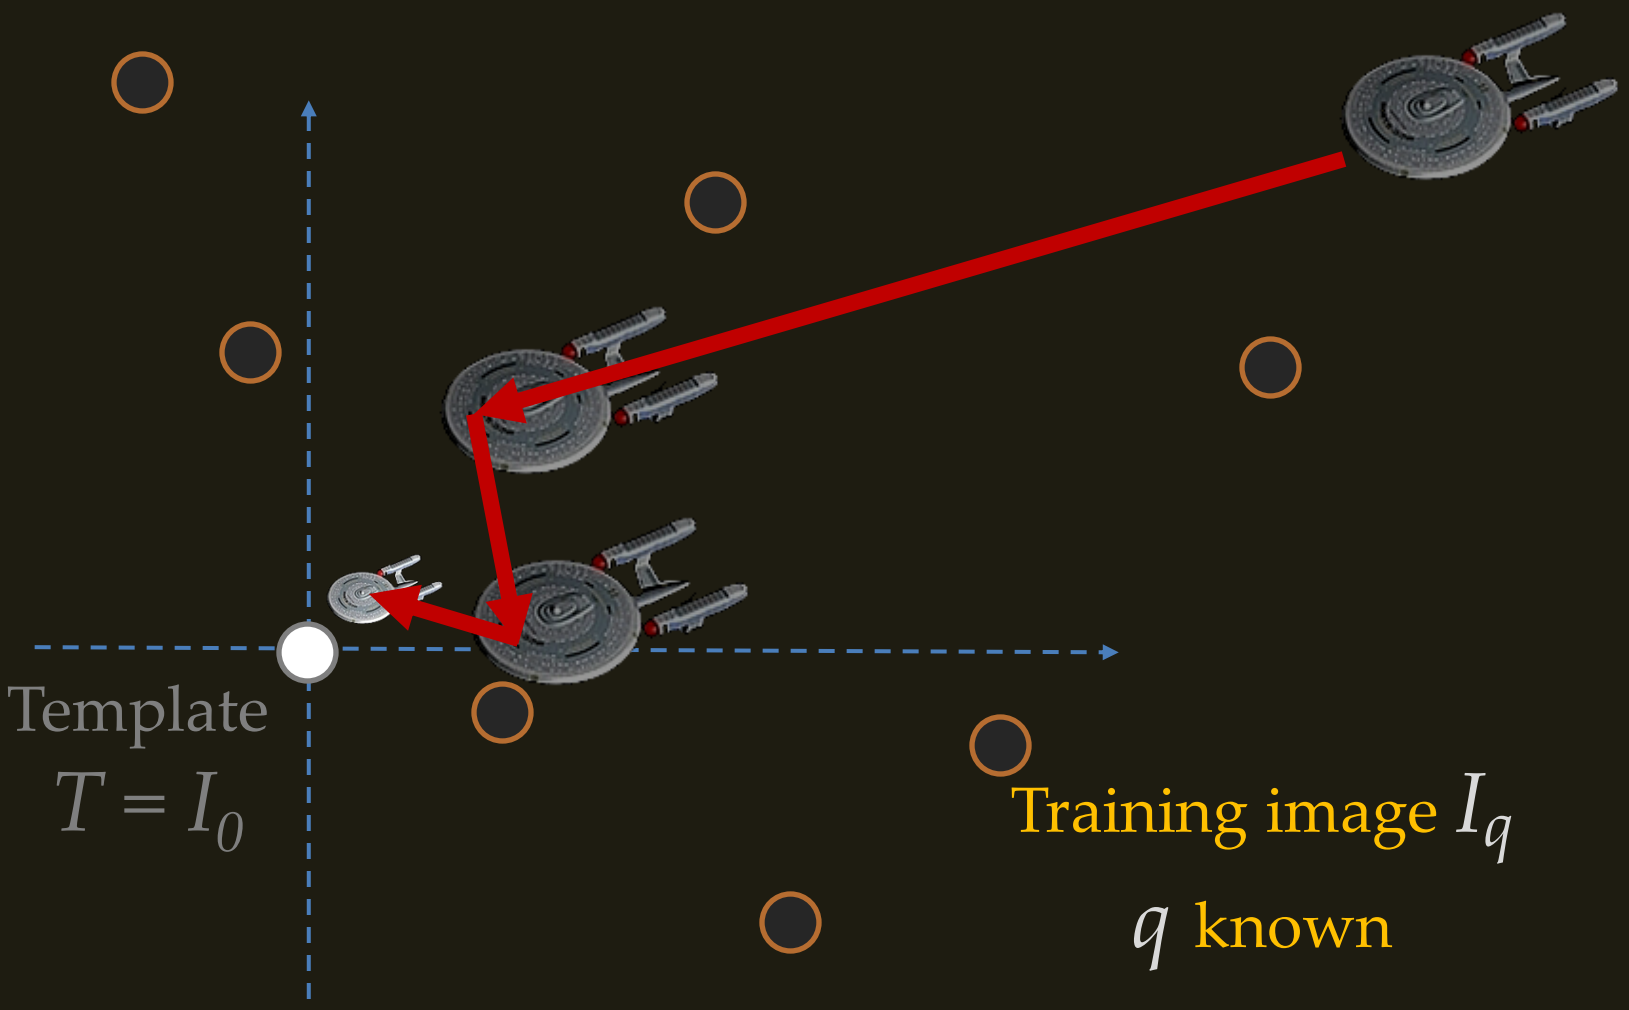

Similar operations for images

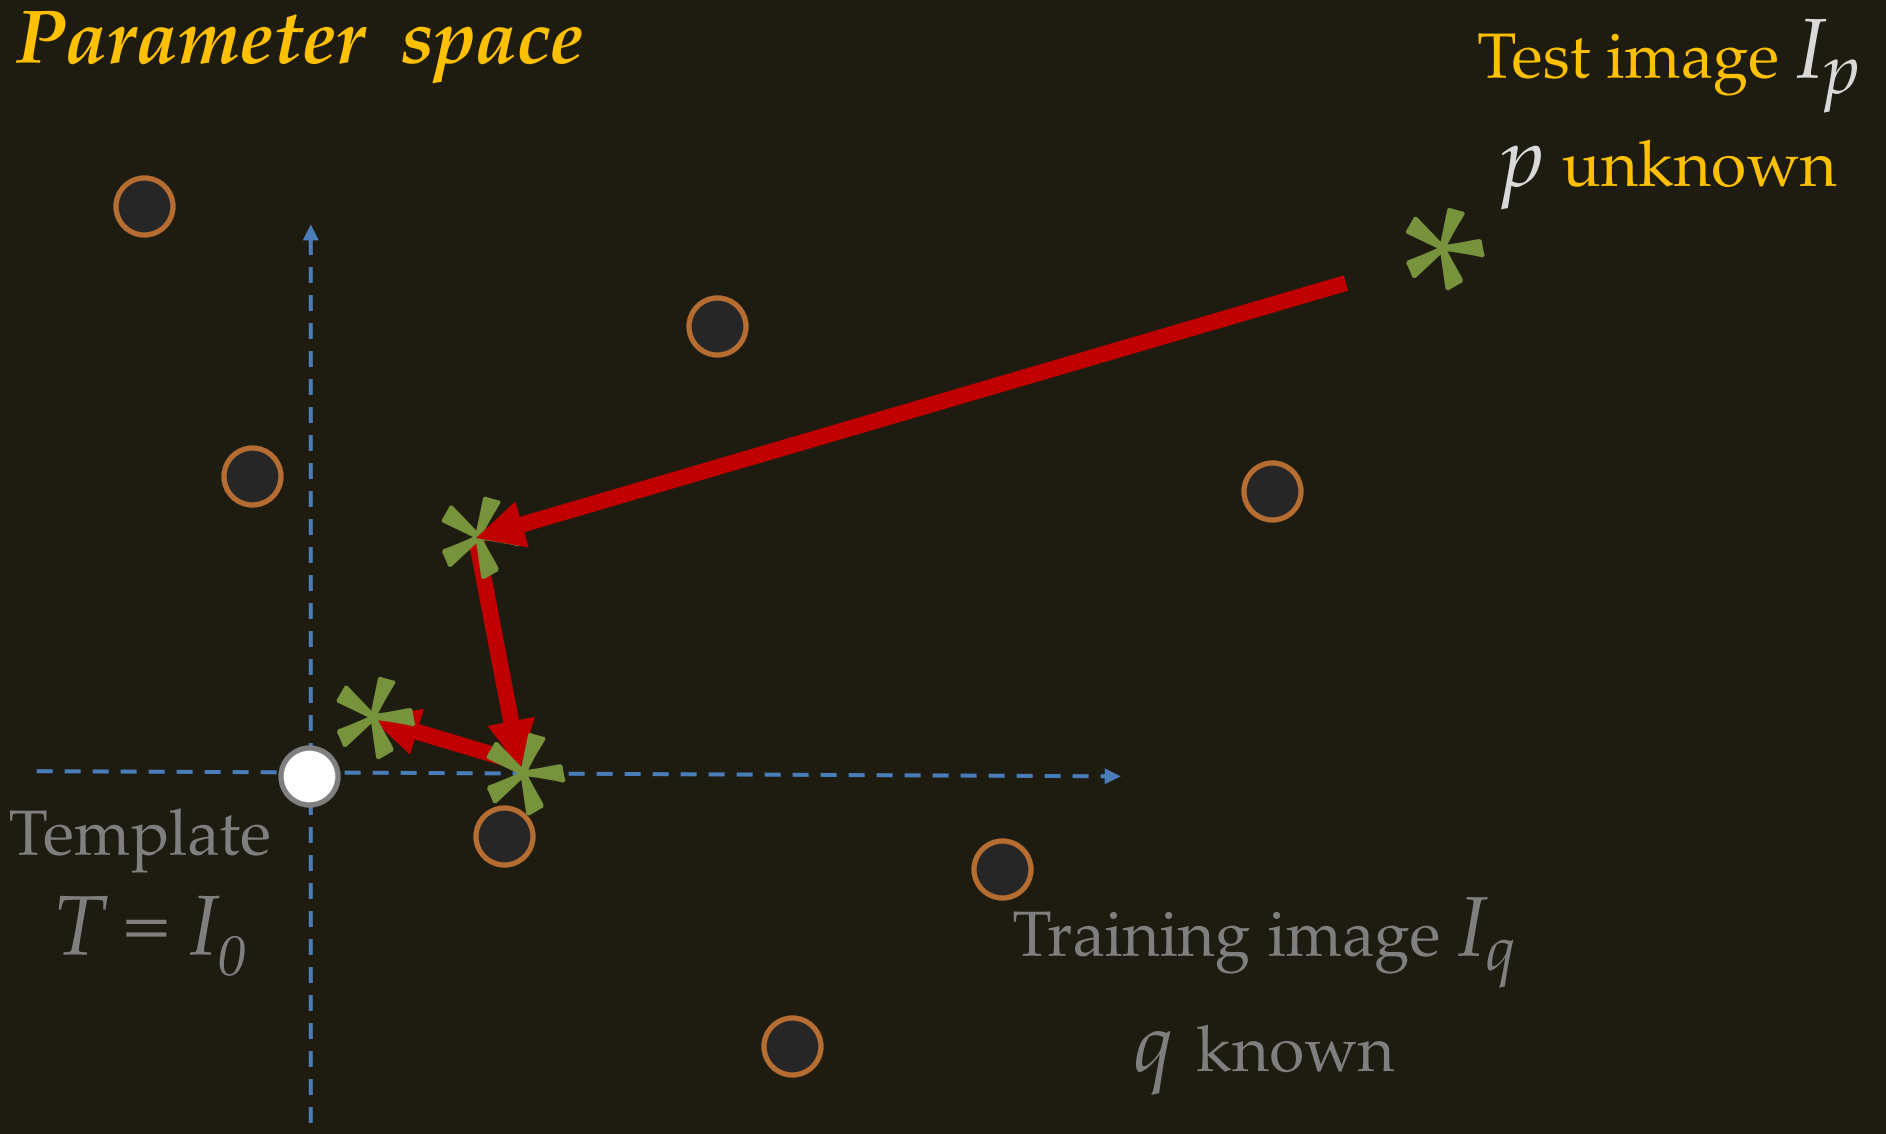

Similar operations for images

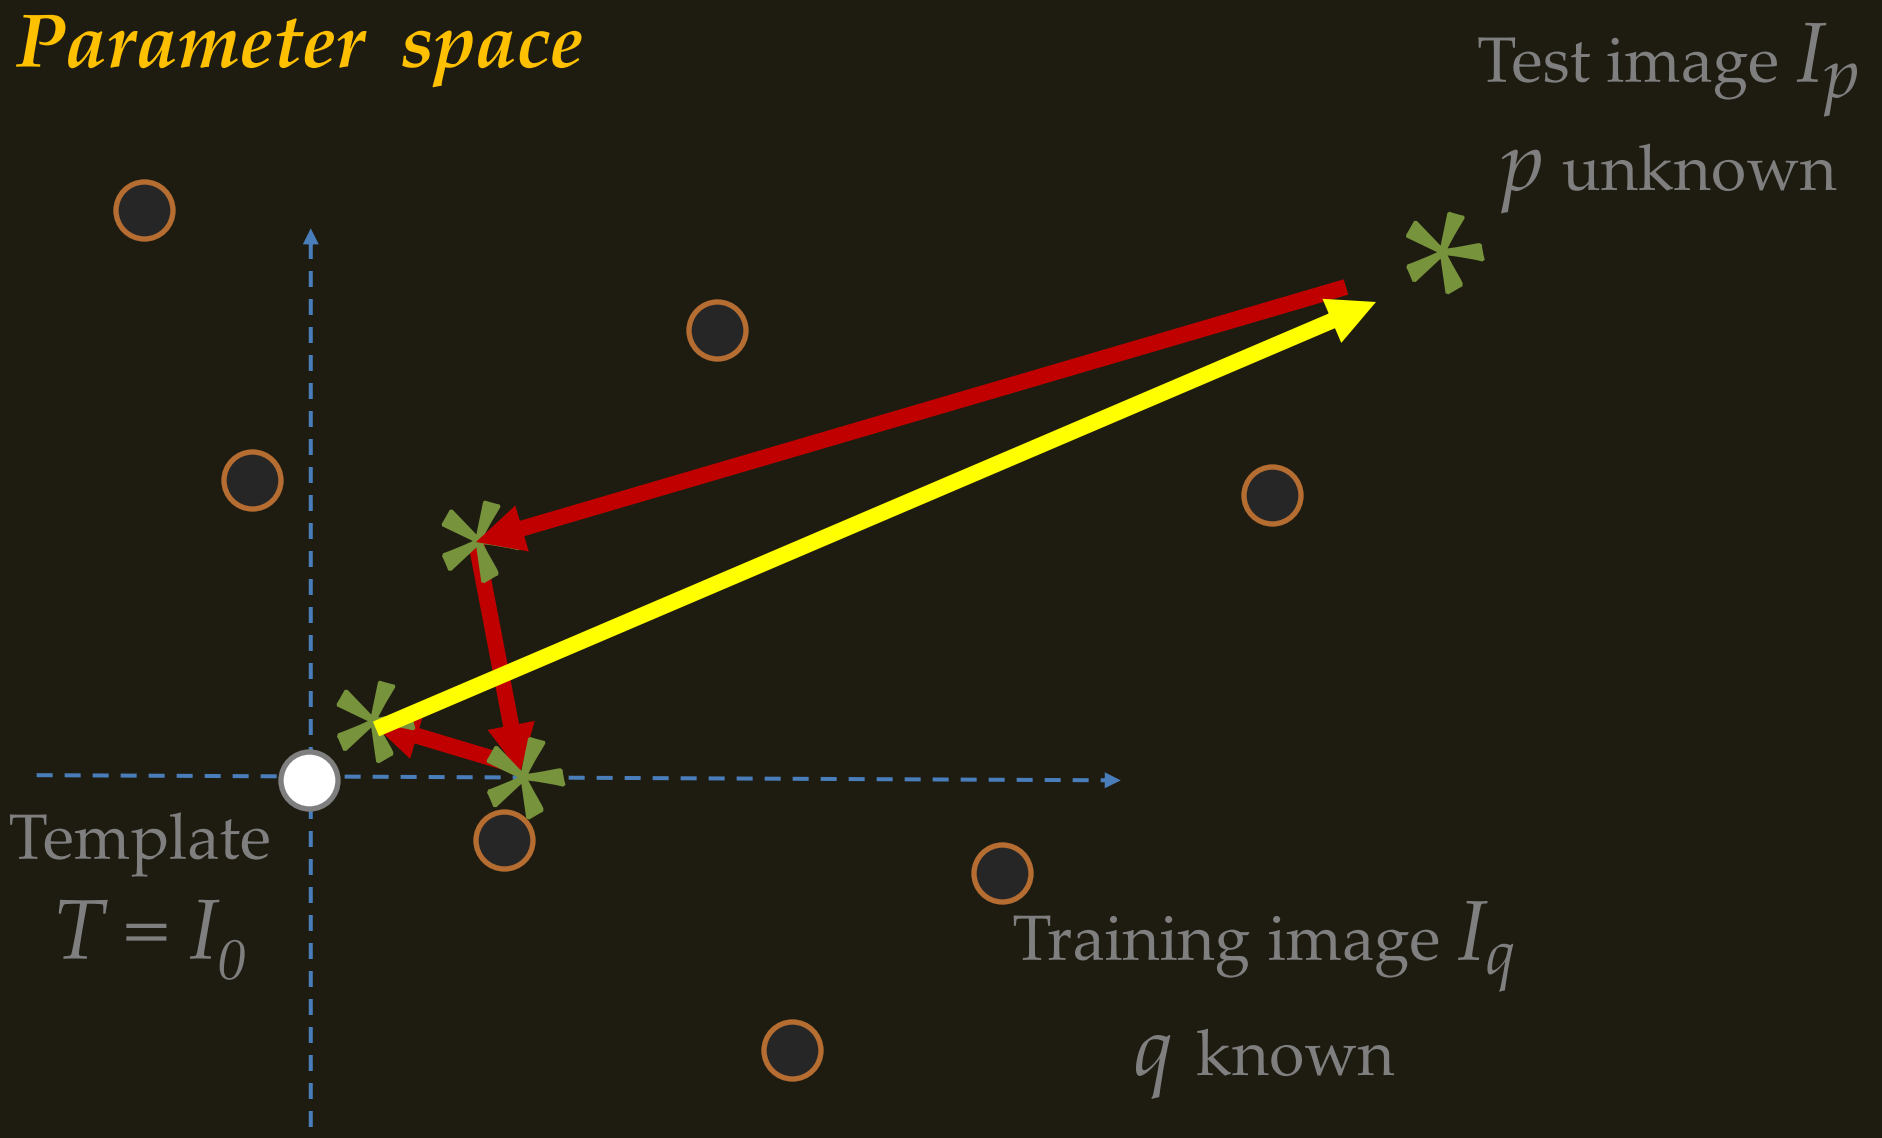

The three components of our algorithm

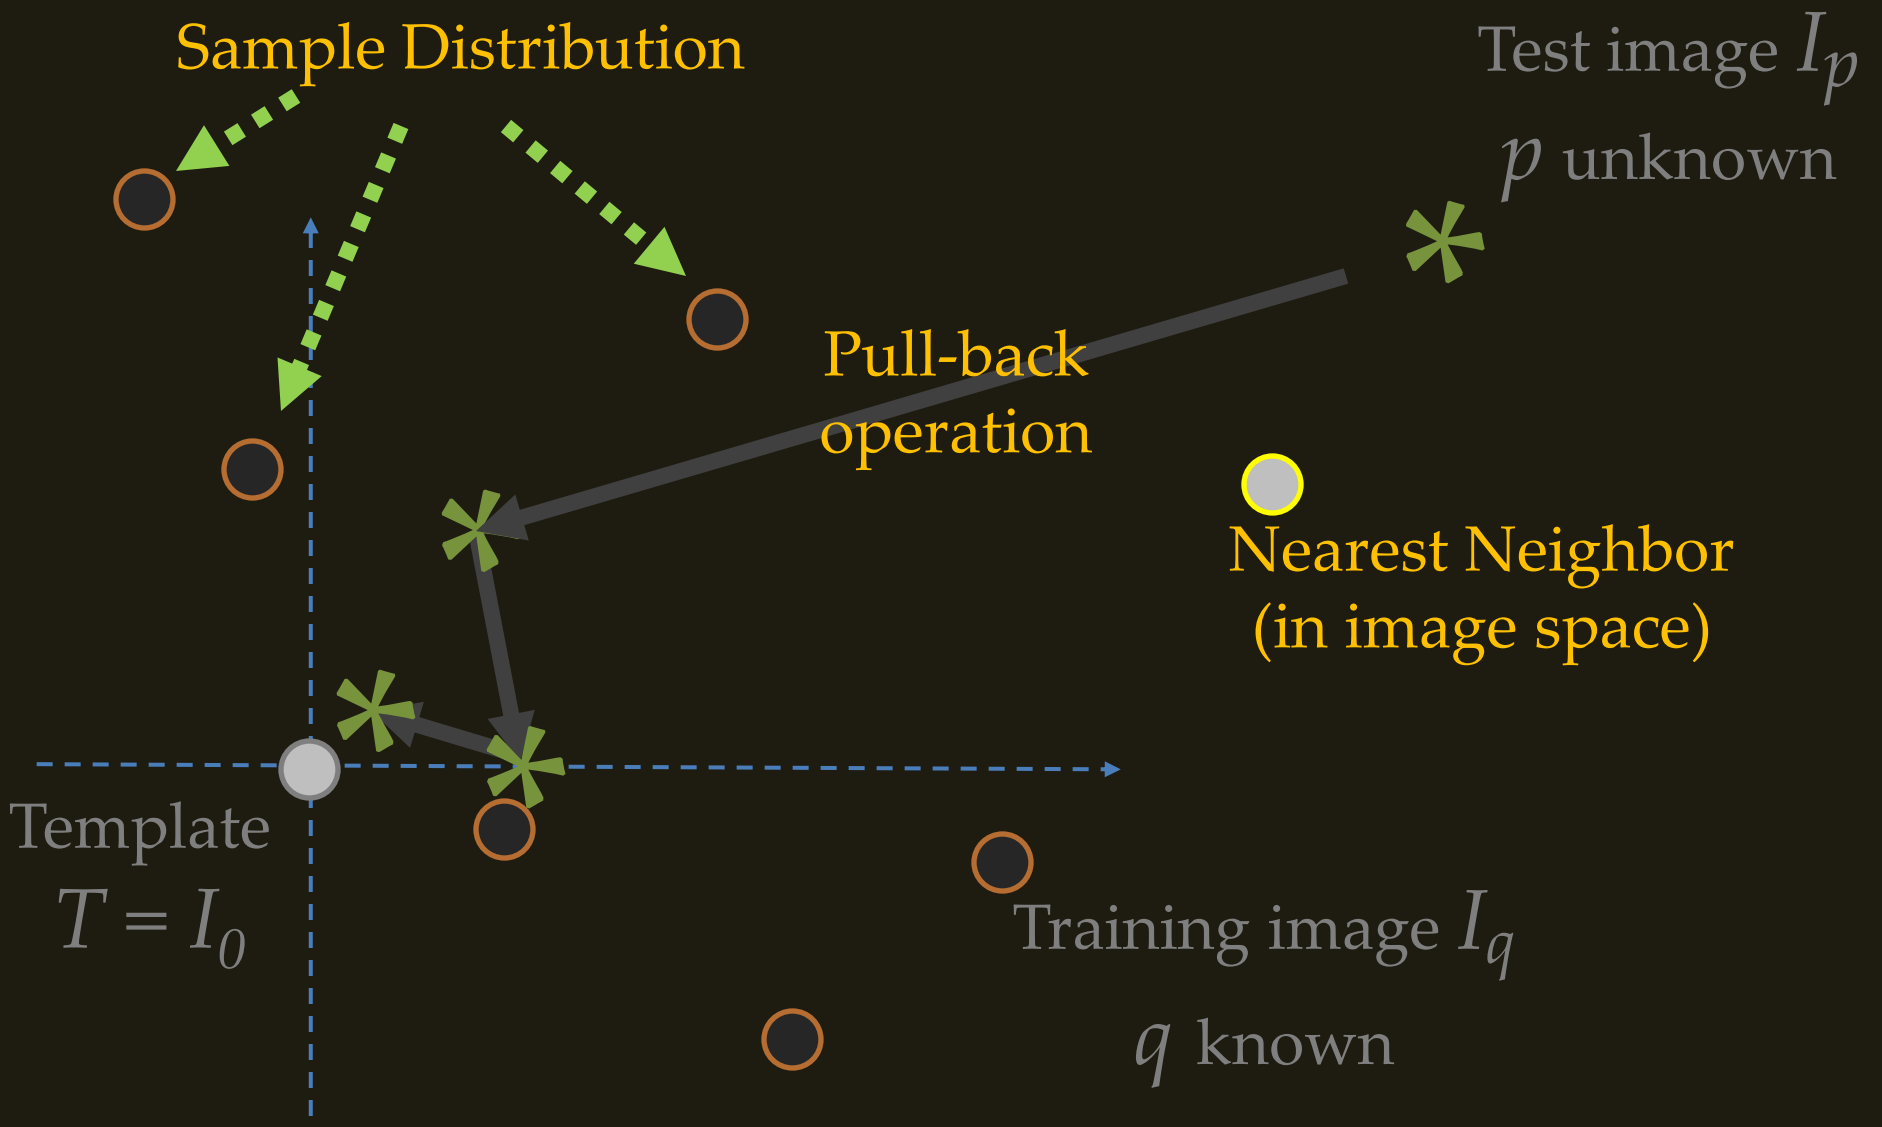

NN in image vs. parameter space

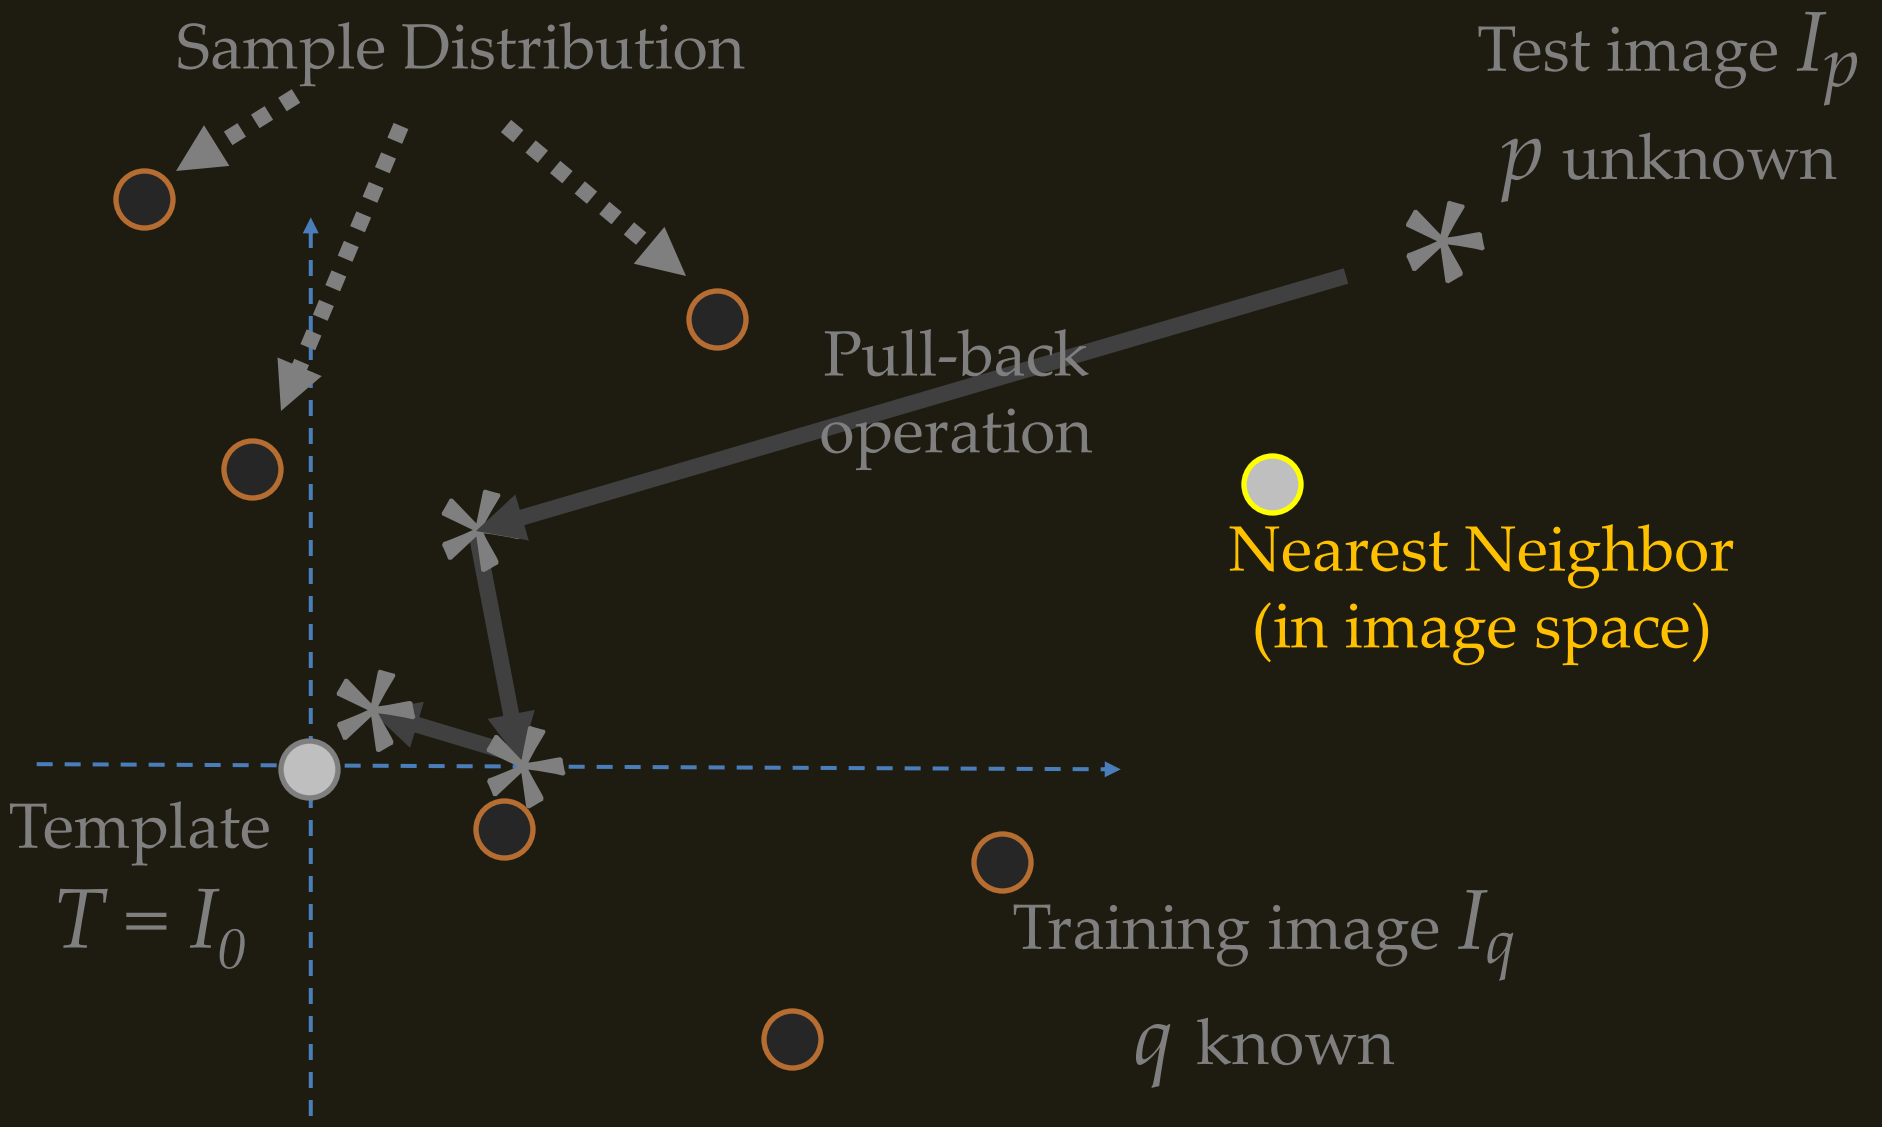

NN in image vs. parameter space

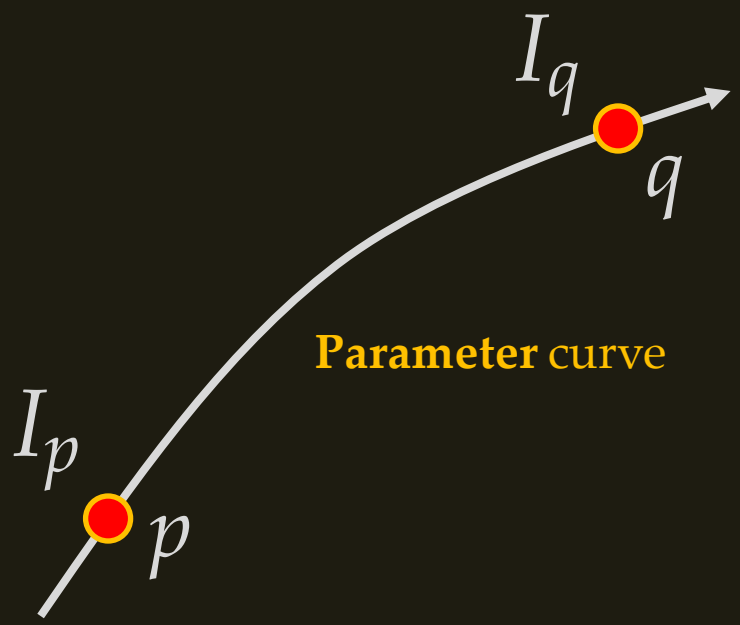

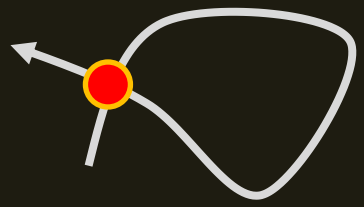

Failure case
(one-to-many mapping)

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

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

The three components of our algorithm

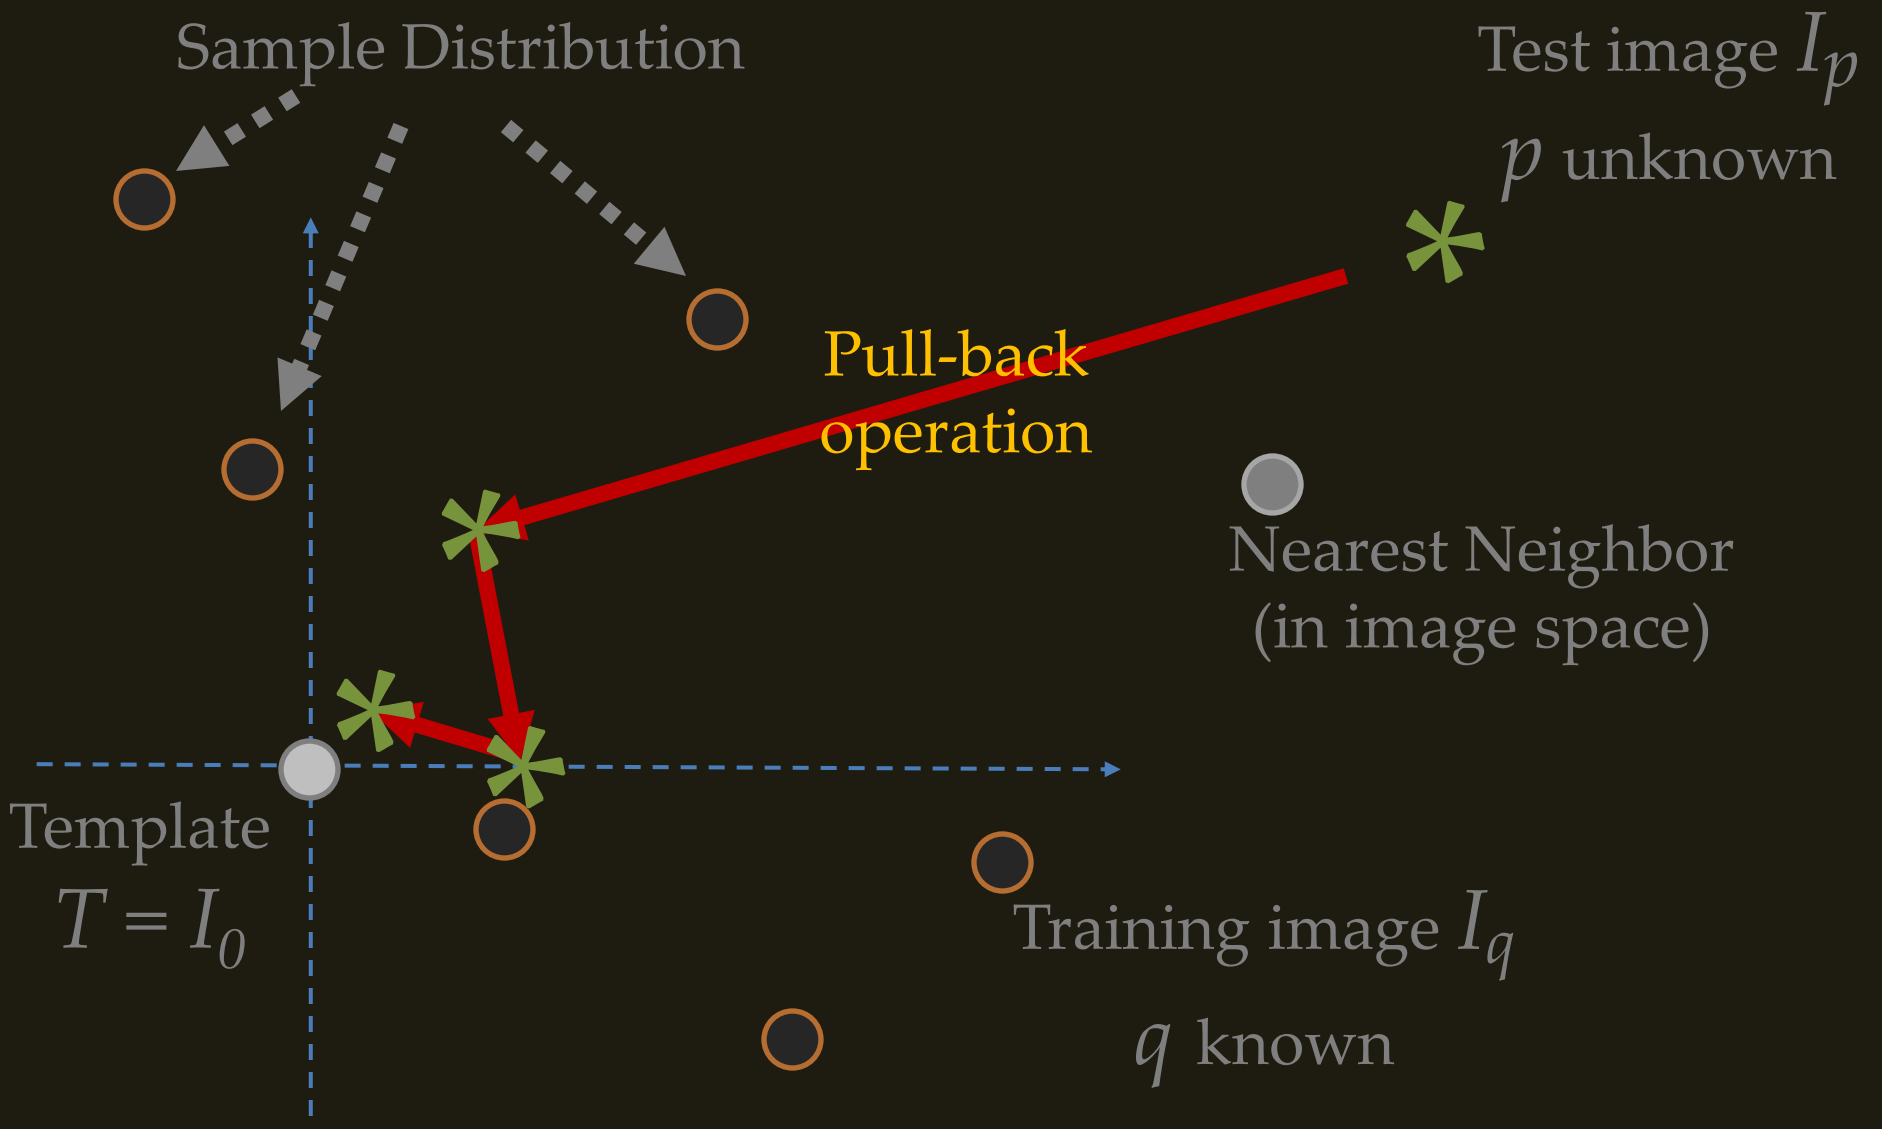

The pull-back operation H

If distortion is invertible, e.g. affine

$$H(I_p, q) = \text{Inverse}\,(I_p, q) = I_{p\text{-}q}$$

[S. Baker and I. Matthews, CVPR 2001]

Non-invertible distortions

For the case of

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

We prove the following upper bound:

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

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

Failure case → Large resampling artifacts:

$$\| I_{p\text{-}q} - H(I_p\,,\, q) \| \leq R \, \| p\text{-}q \| + \ldots\ldots$$

The distribution of training samples

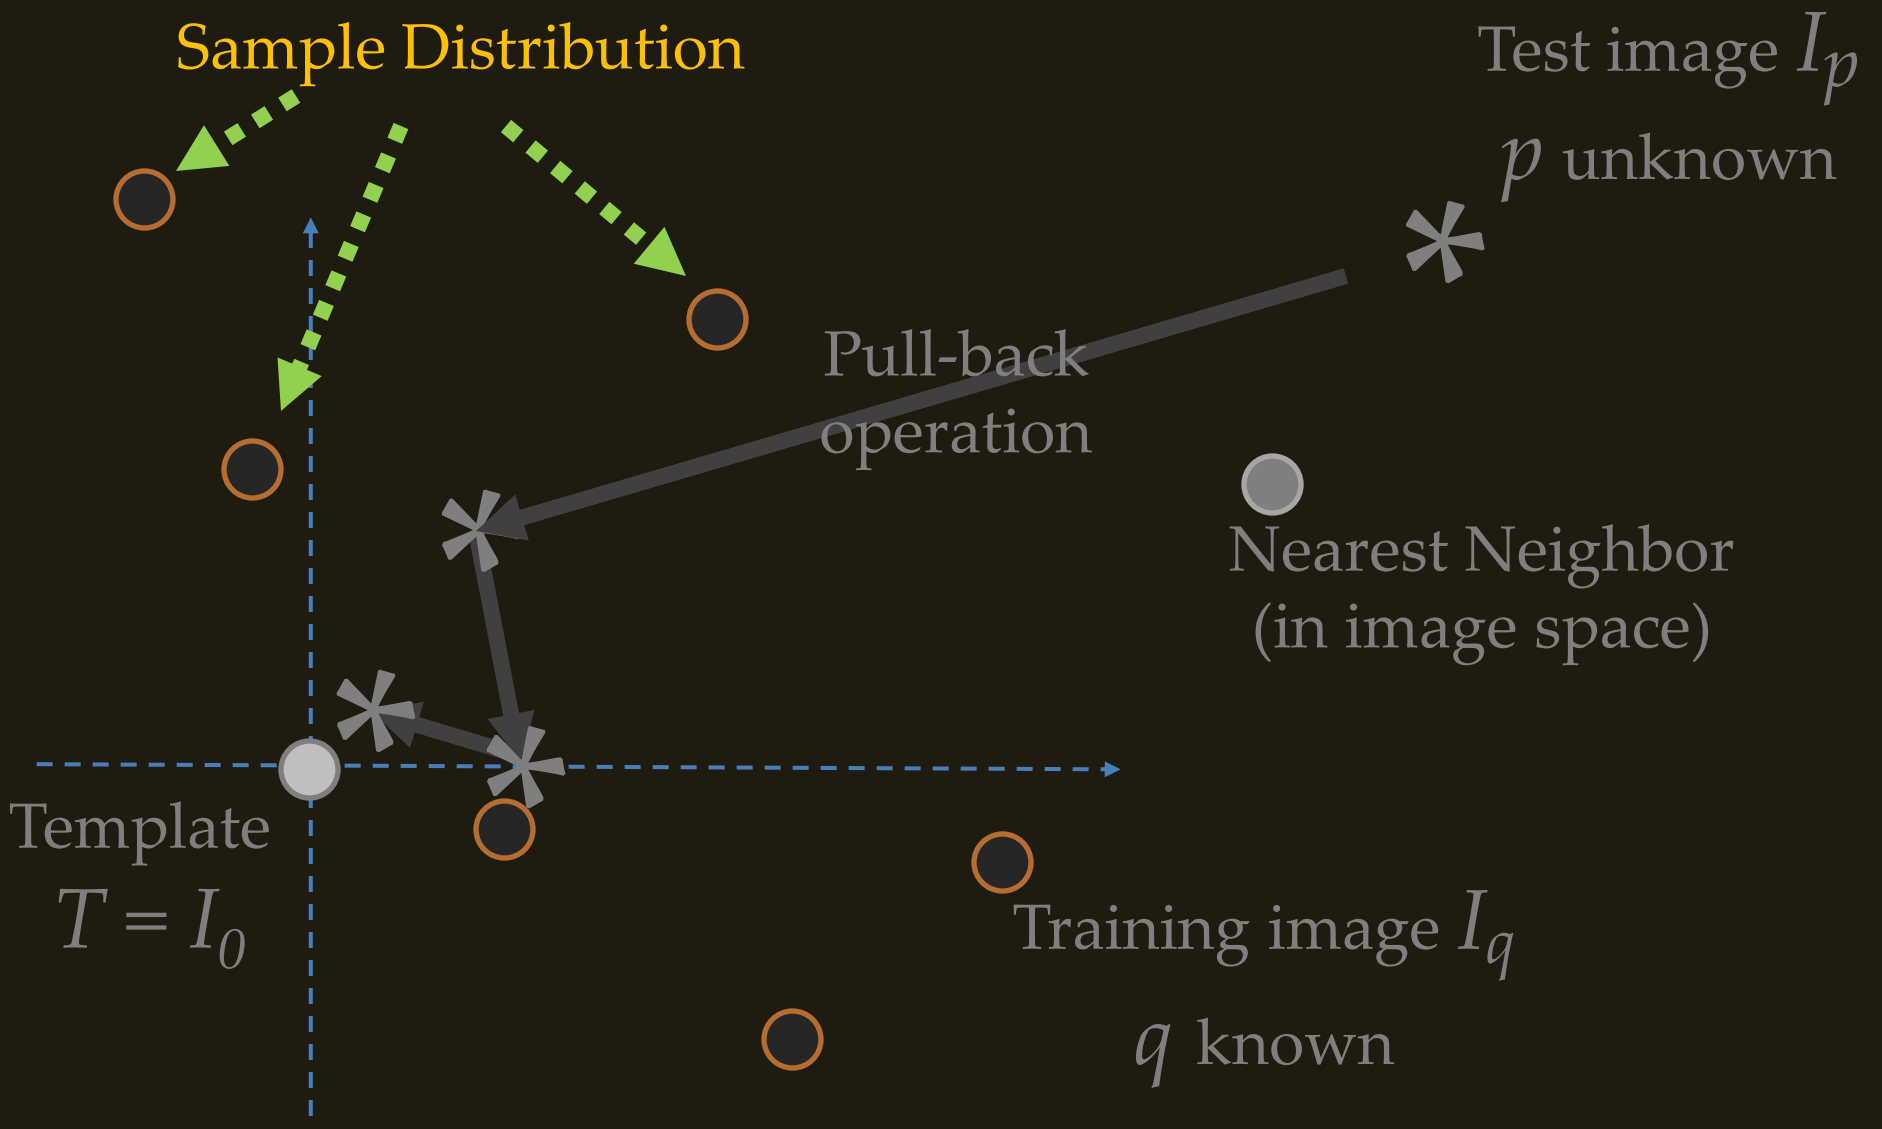

Training sample distribution

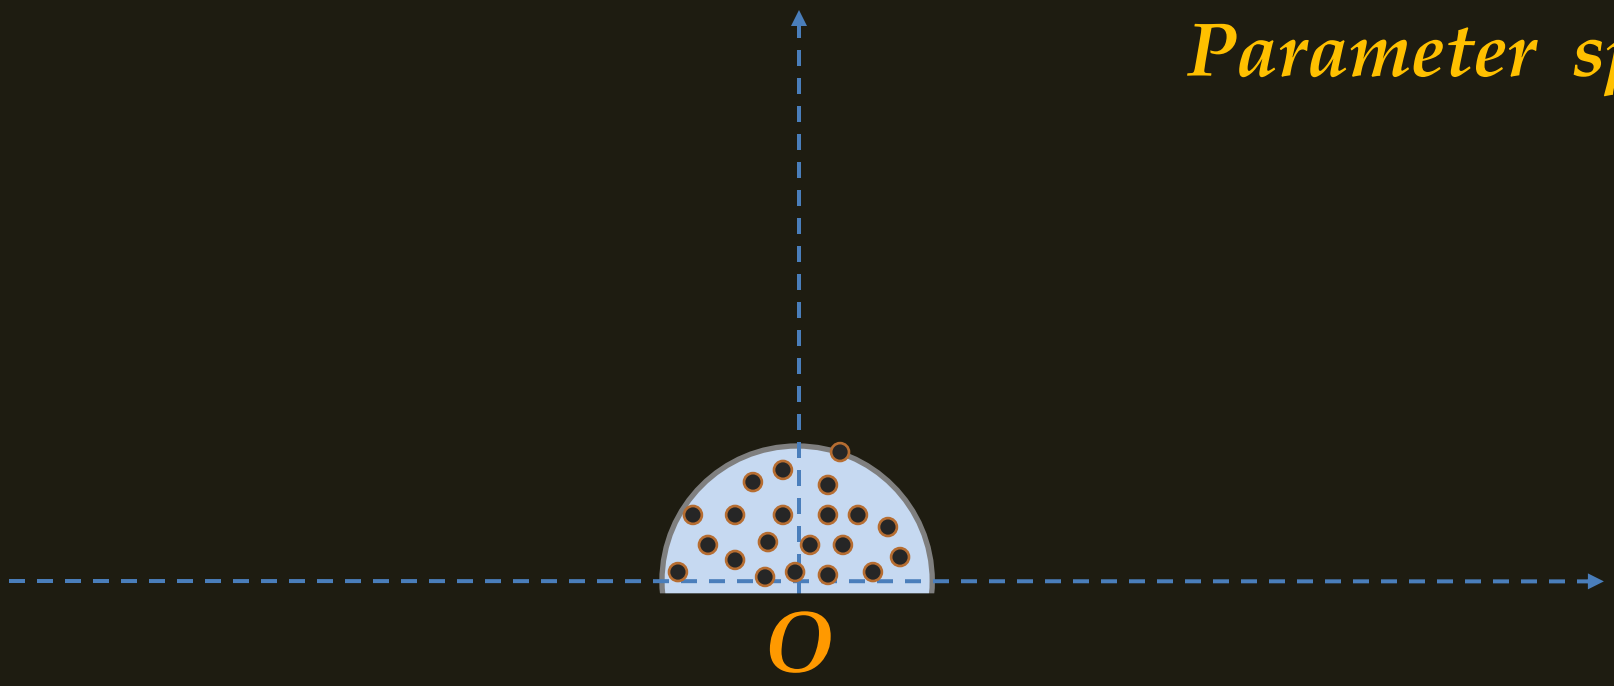

Training sample distribution

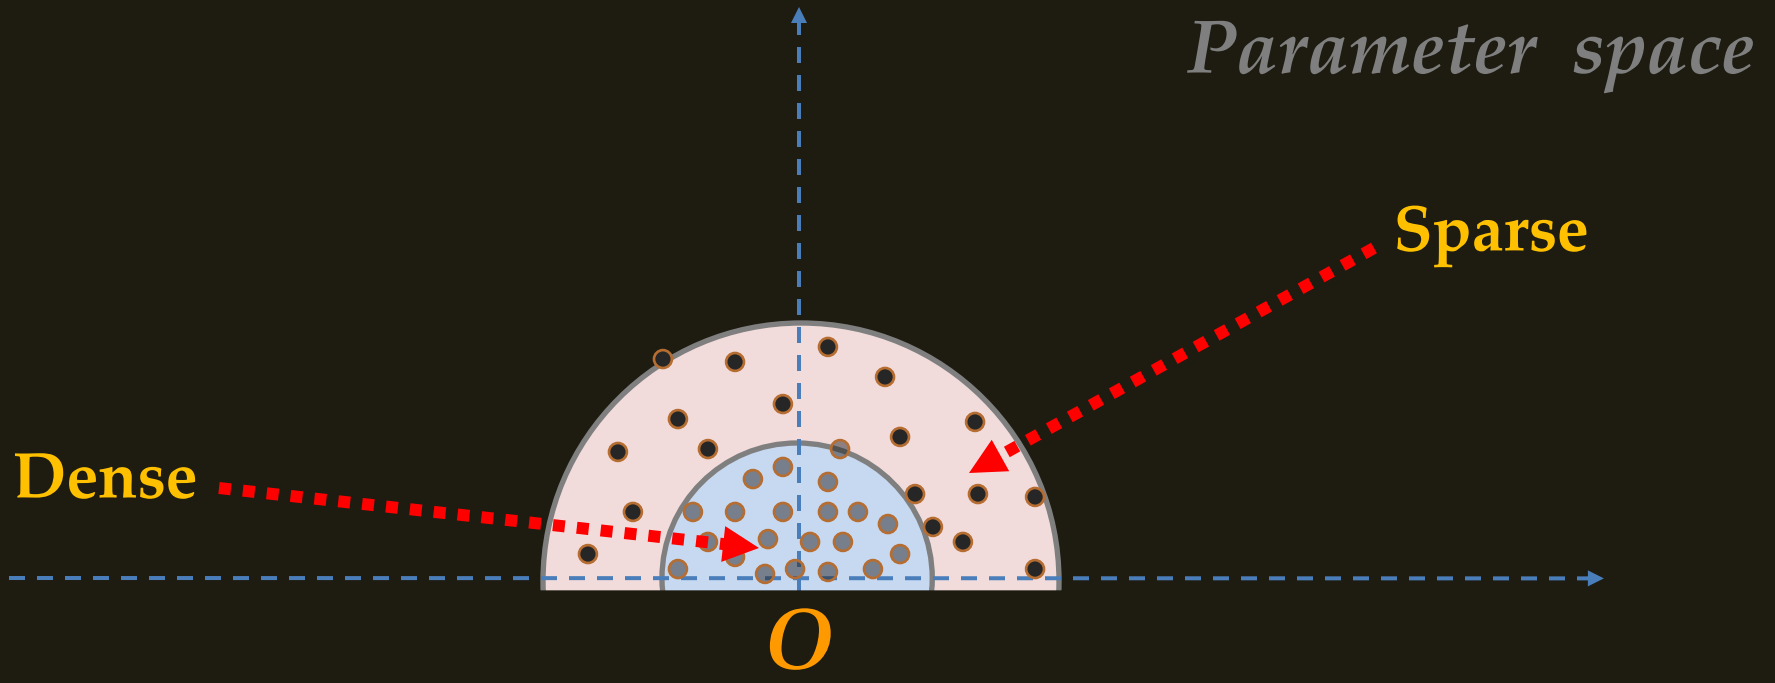

Training sample distribution

$$Density \propto \frac{1}{\|\mathbf{p}\|^d}$$

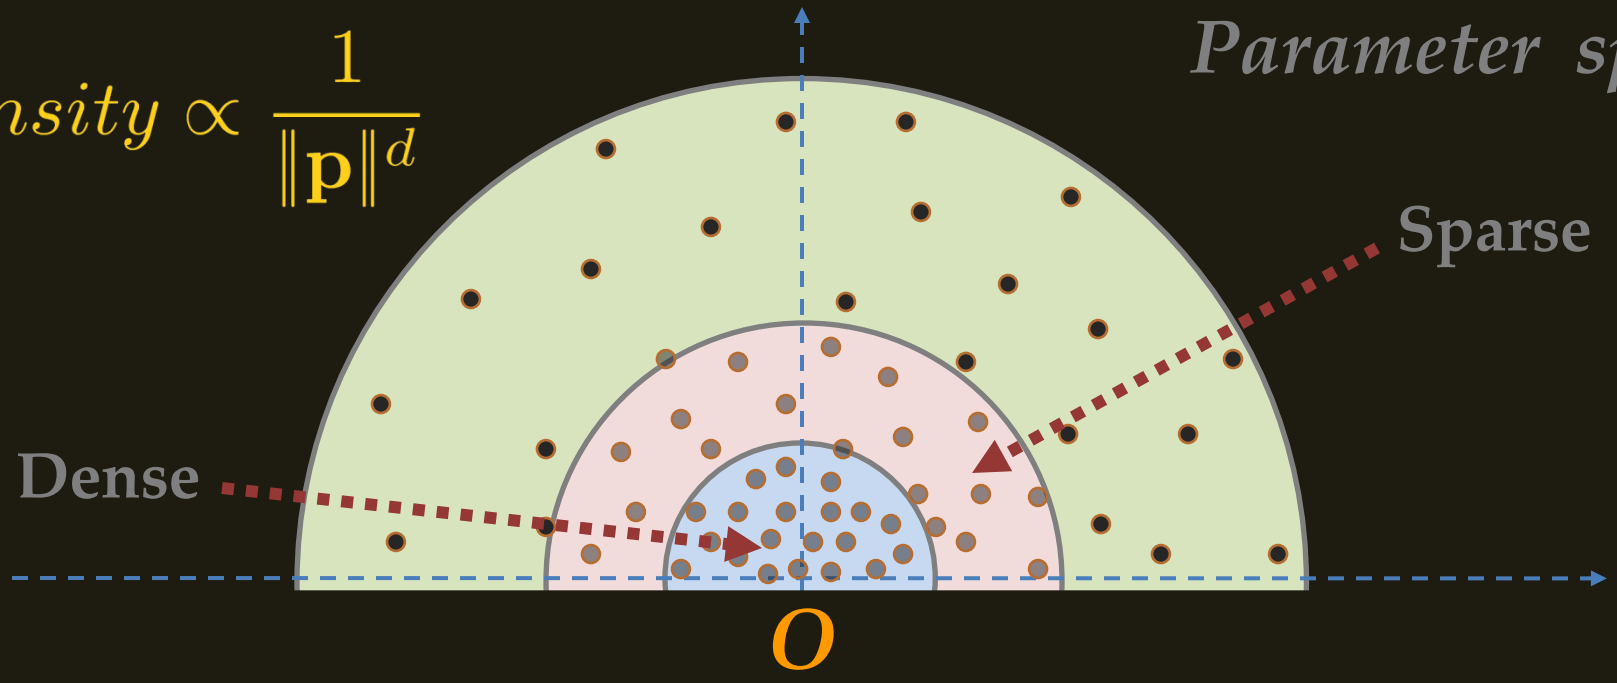

Training sample distribution

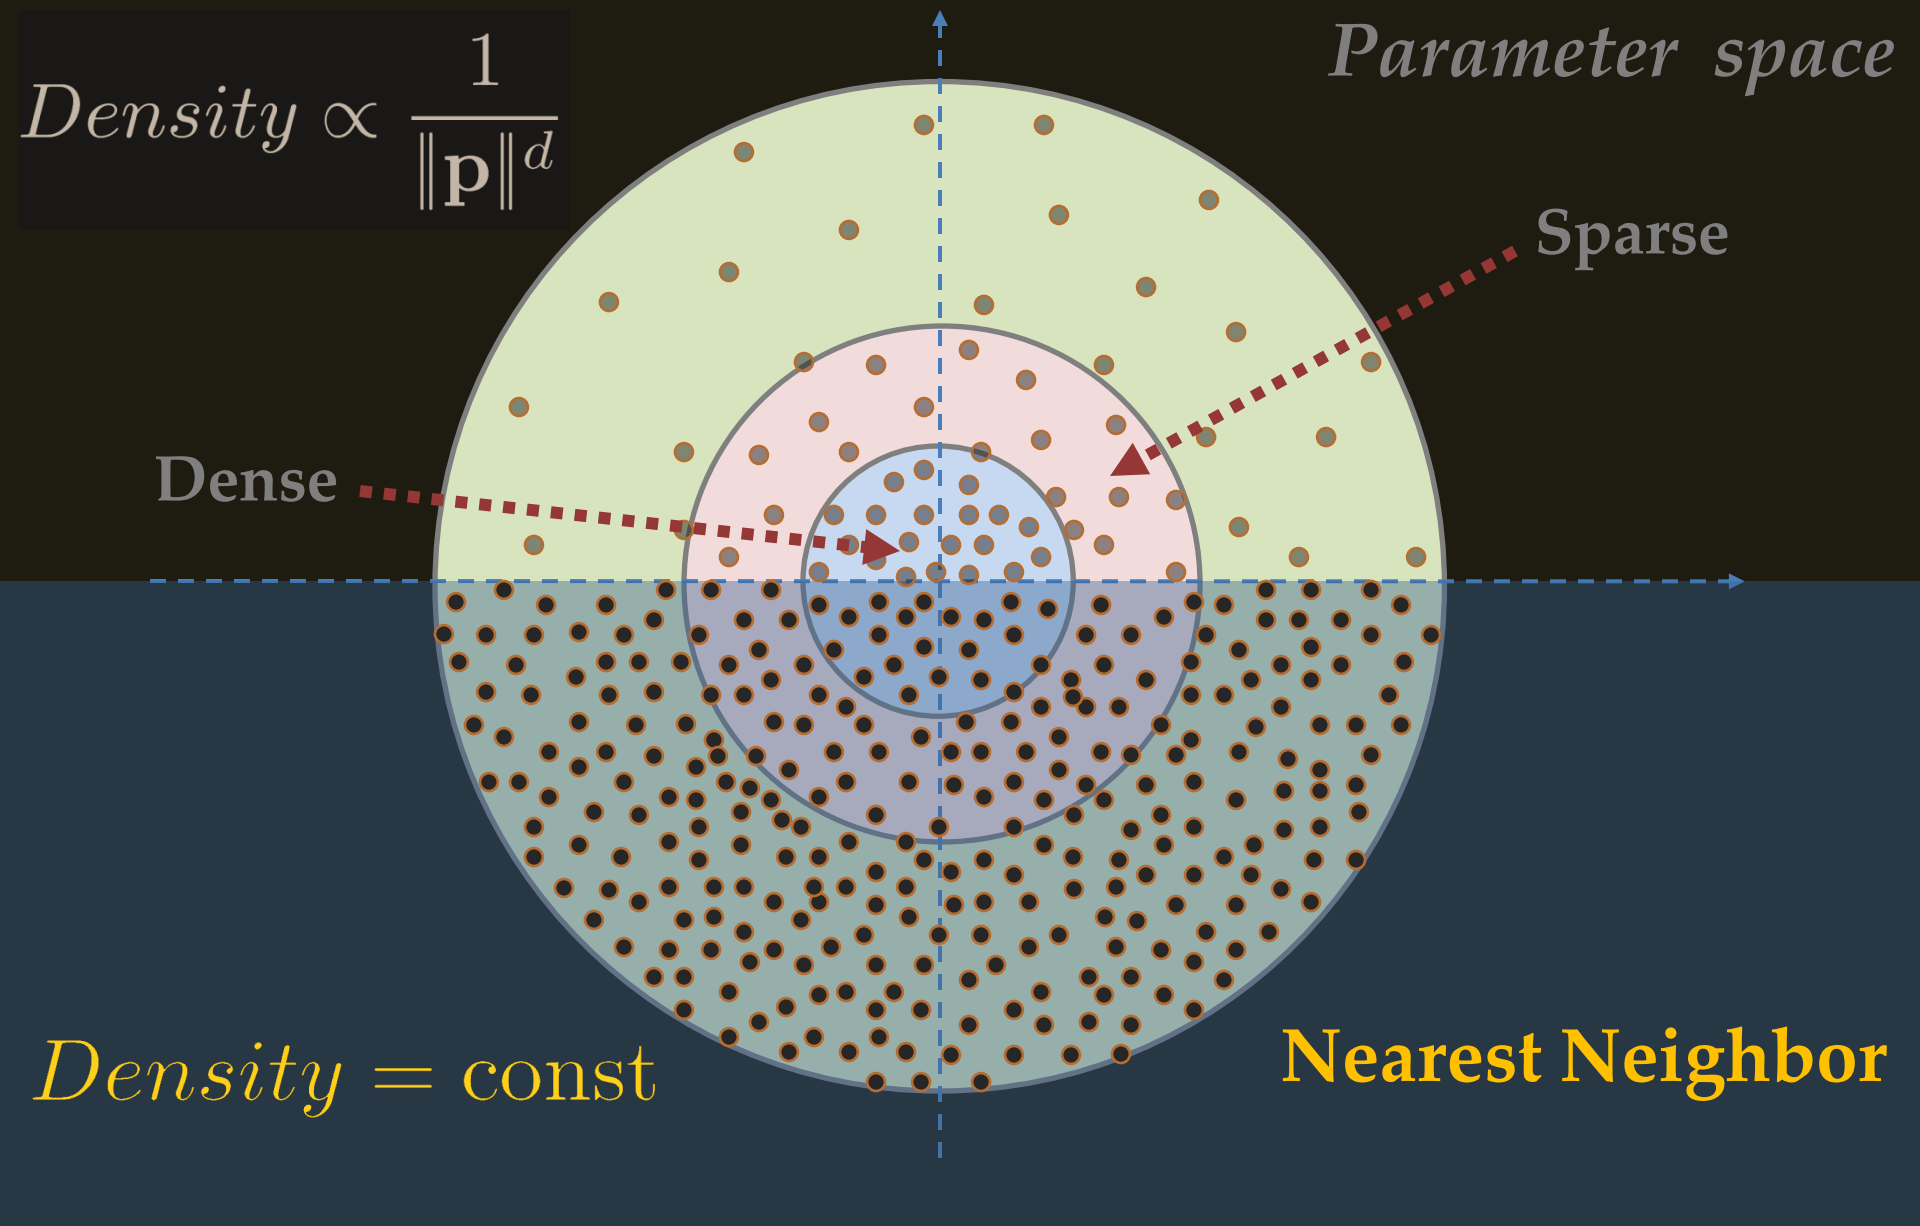

Number of training samples

$O(C^d \log(1/\varepsilon))$

Parameter space

Accuracy = $1/\varepsilon$

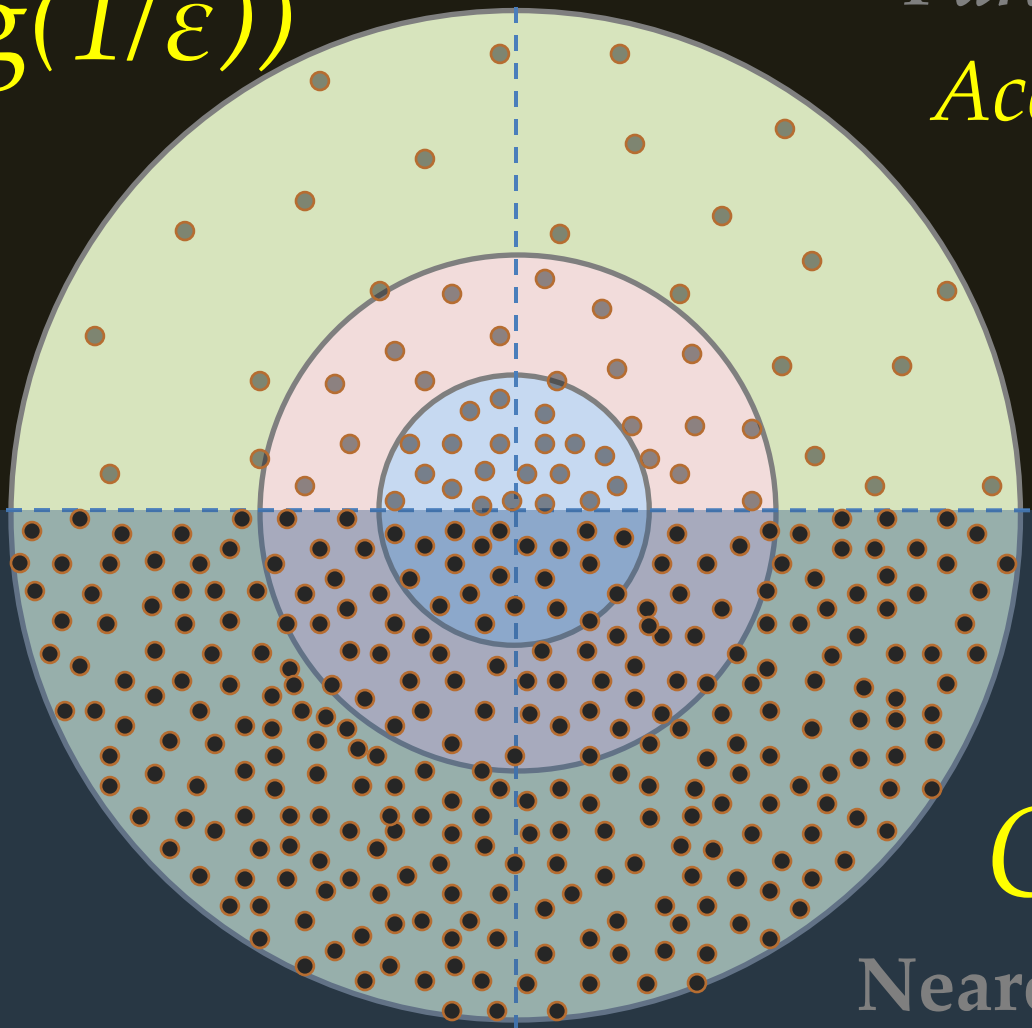

$O(1/\varepsilon)^d$

Nearest Neighbor

Simulations

Distorted Rectified

Distorted Rectified

Convergence Behavior

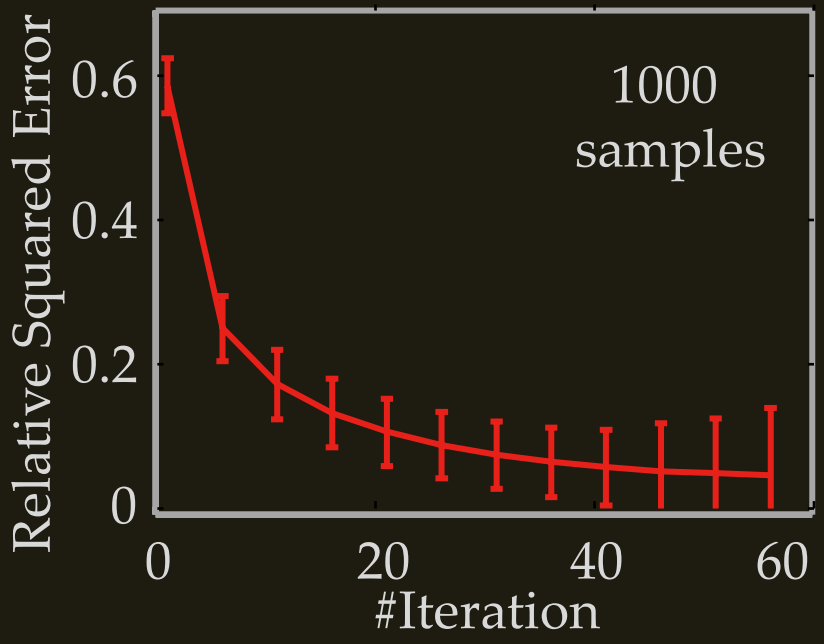

Comparison with NN

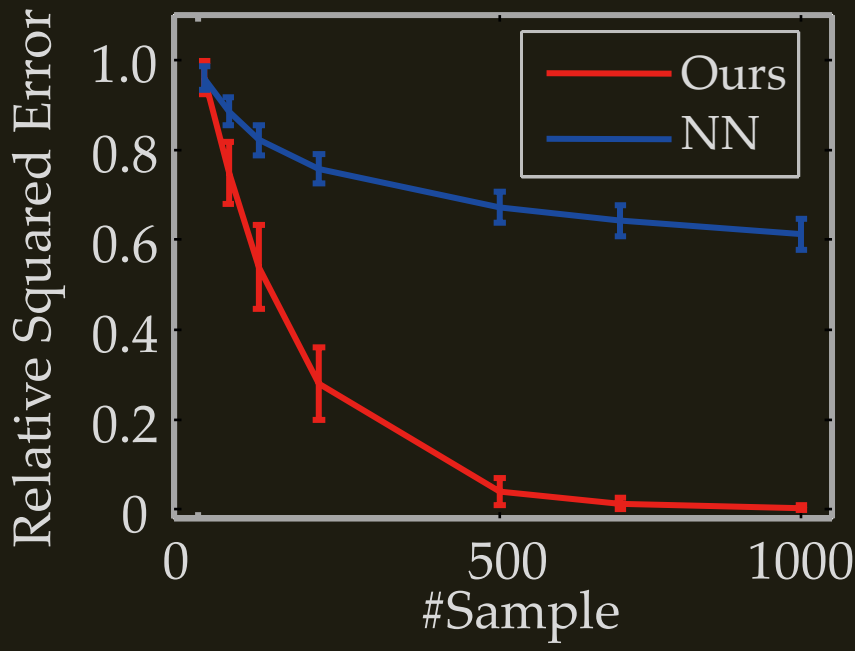

Drift-free video tracking

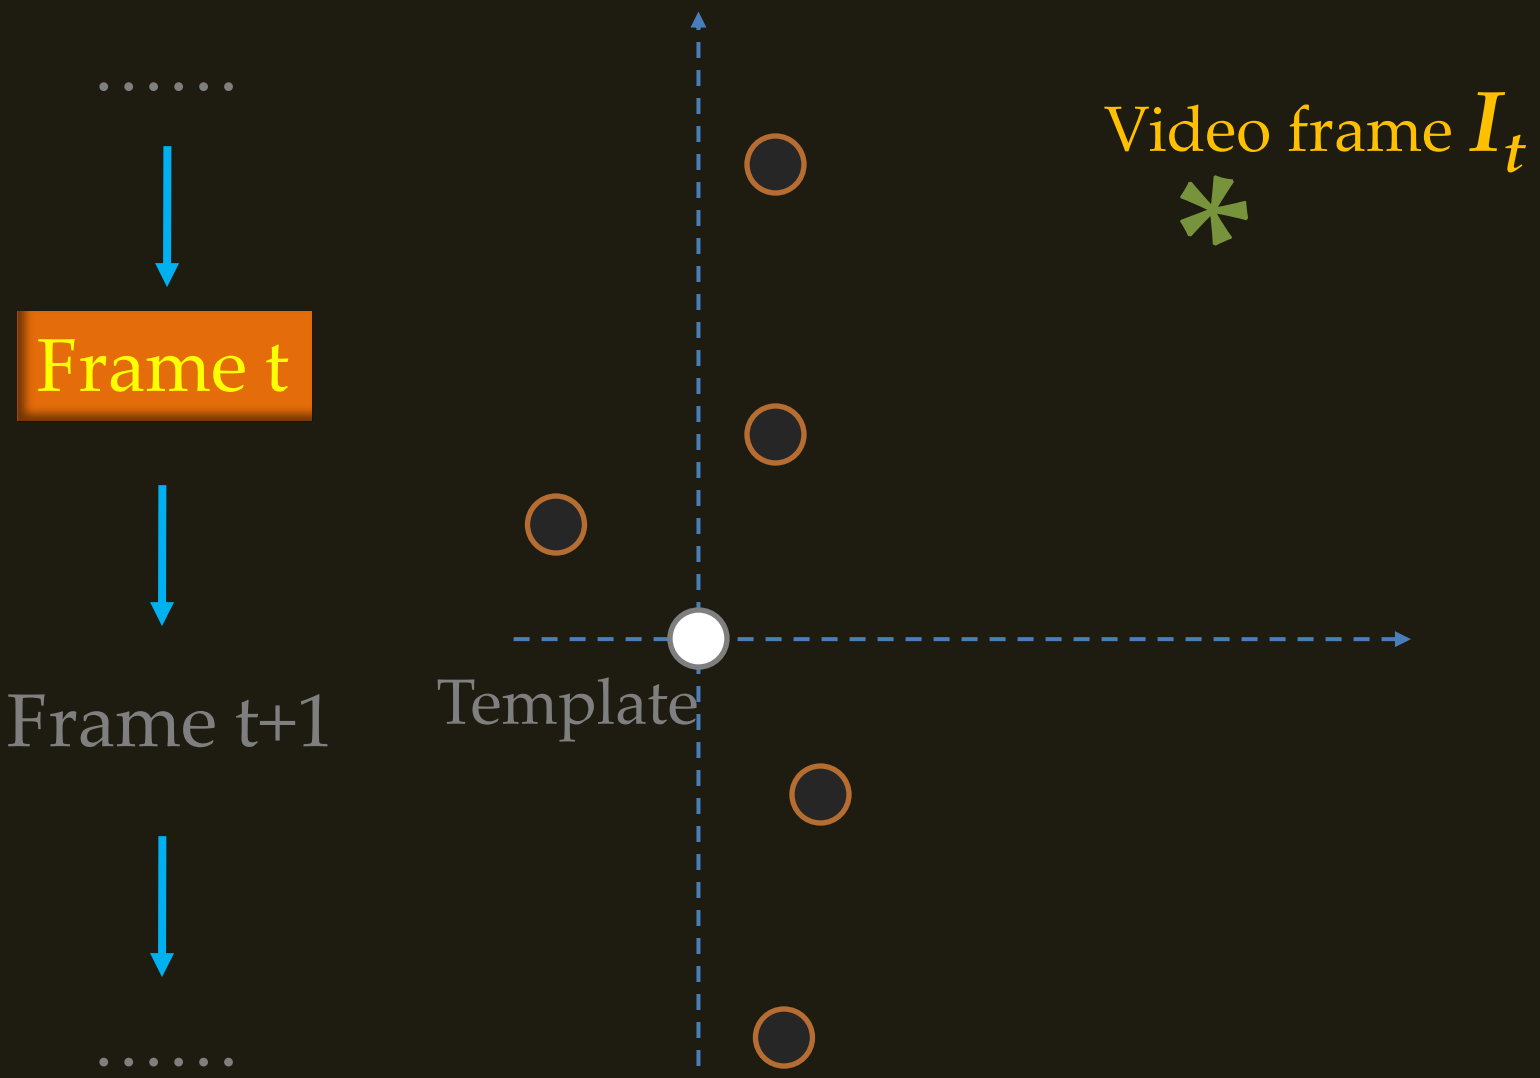

Drift-free video tracking

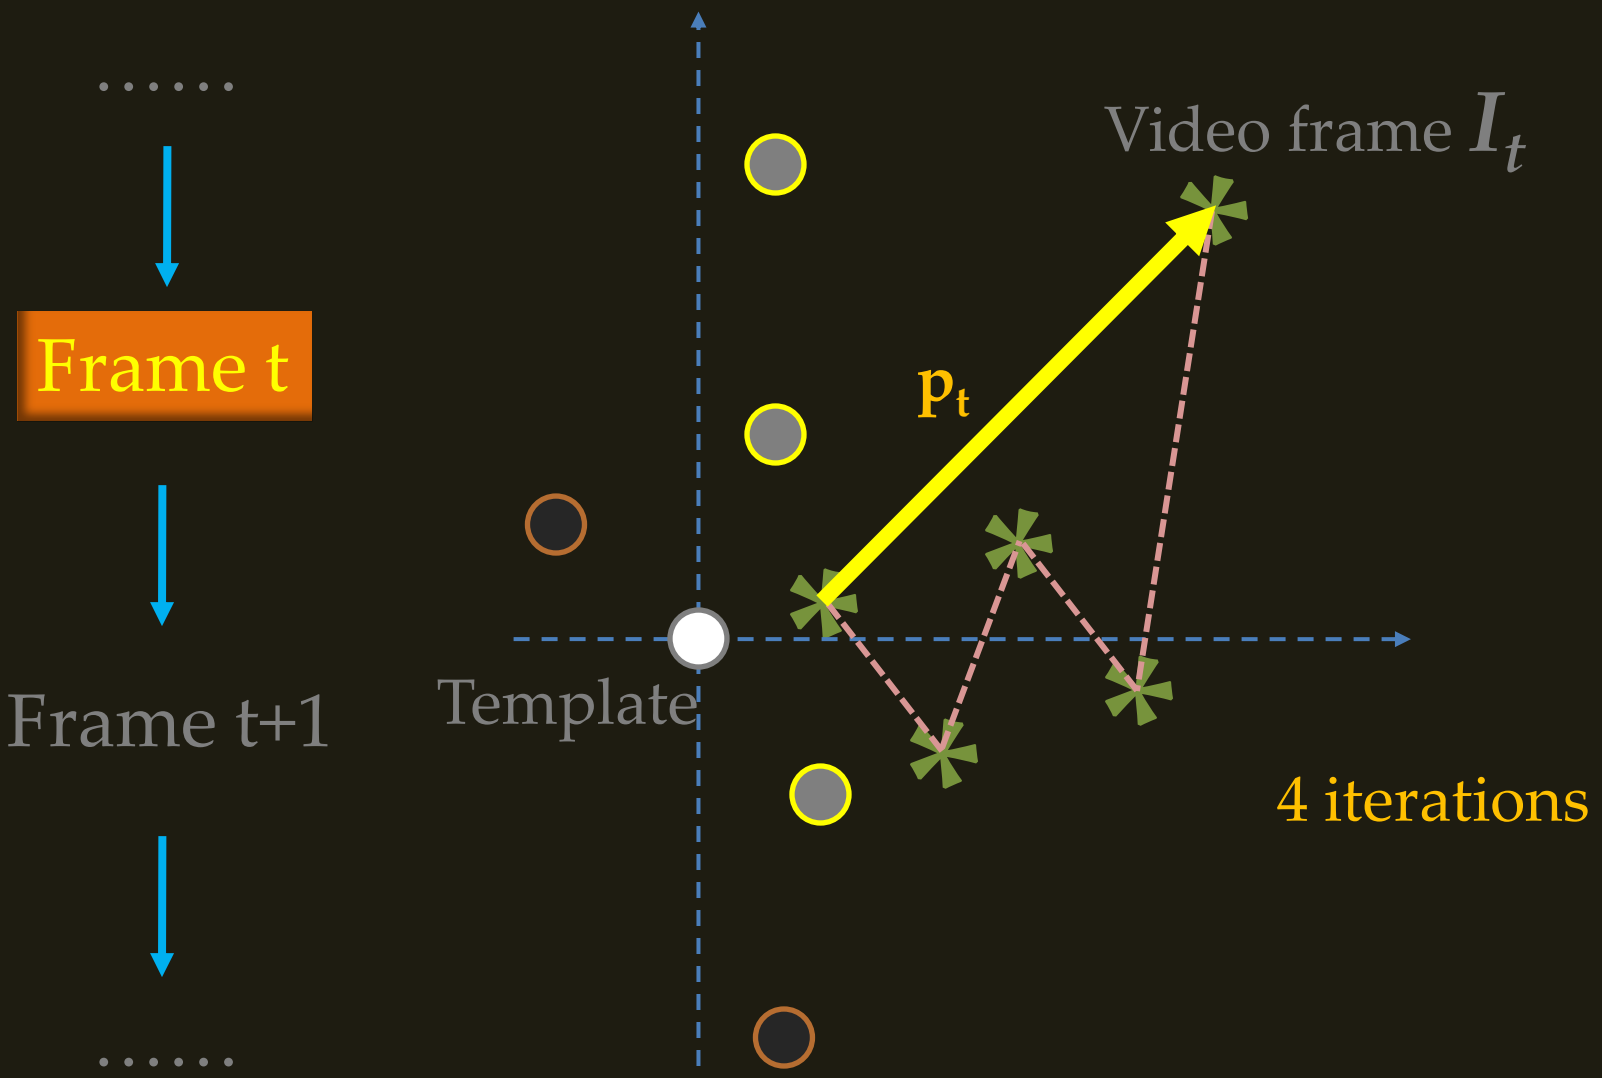

Drift-free video tracking

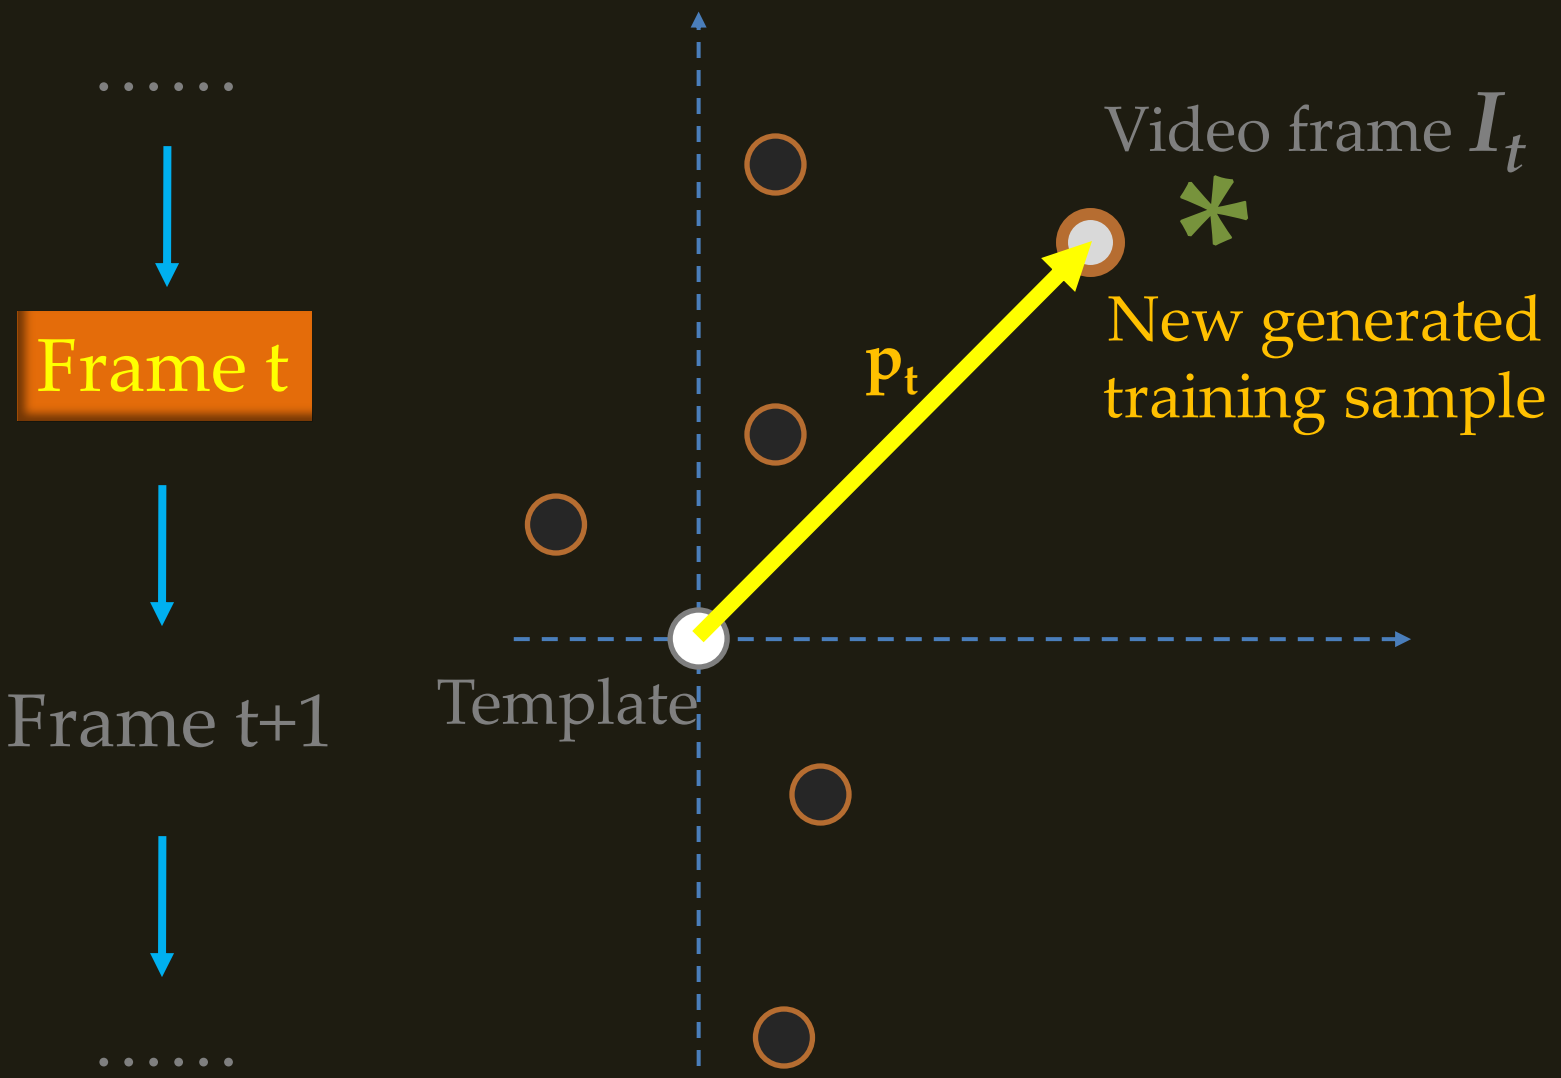

Drift-free video tracking

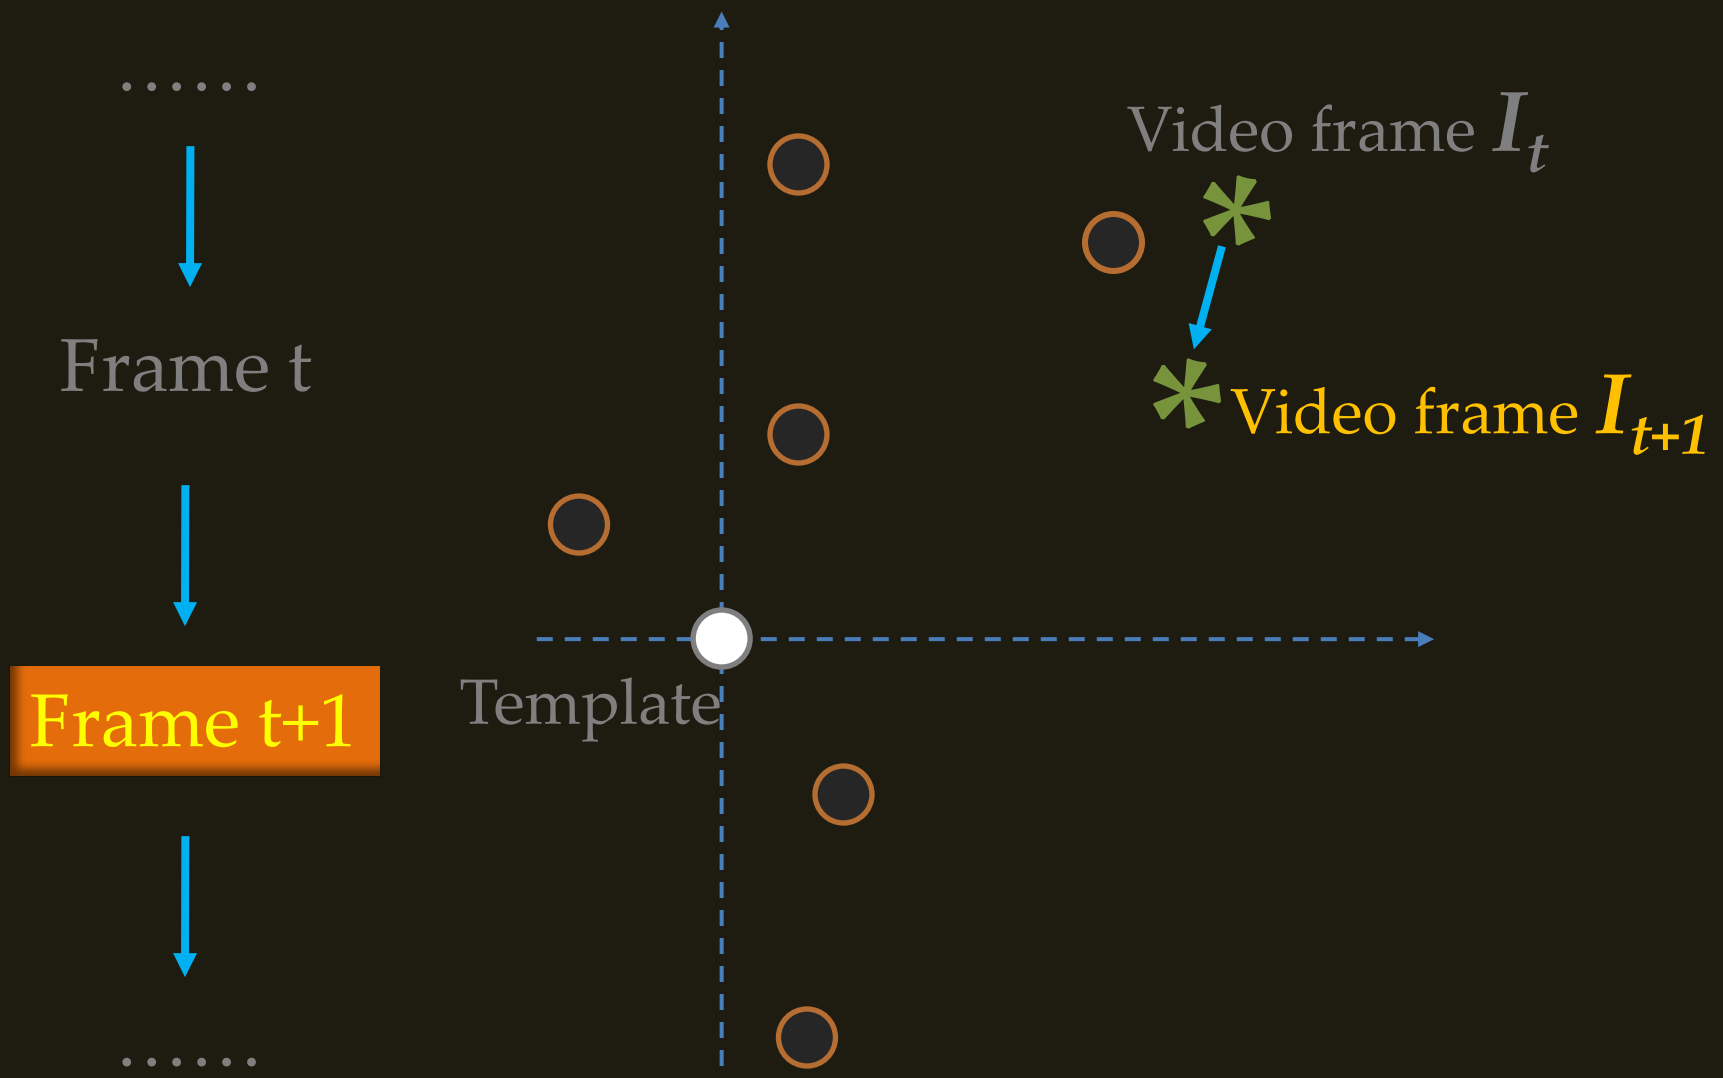

Drift-free video tracking

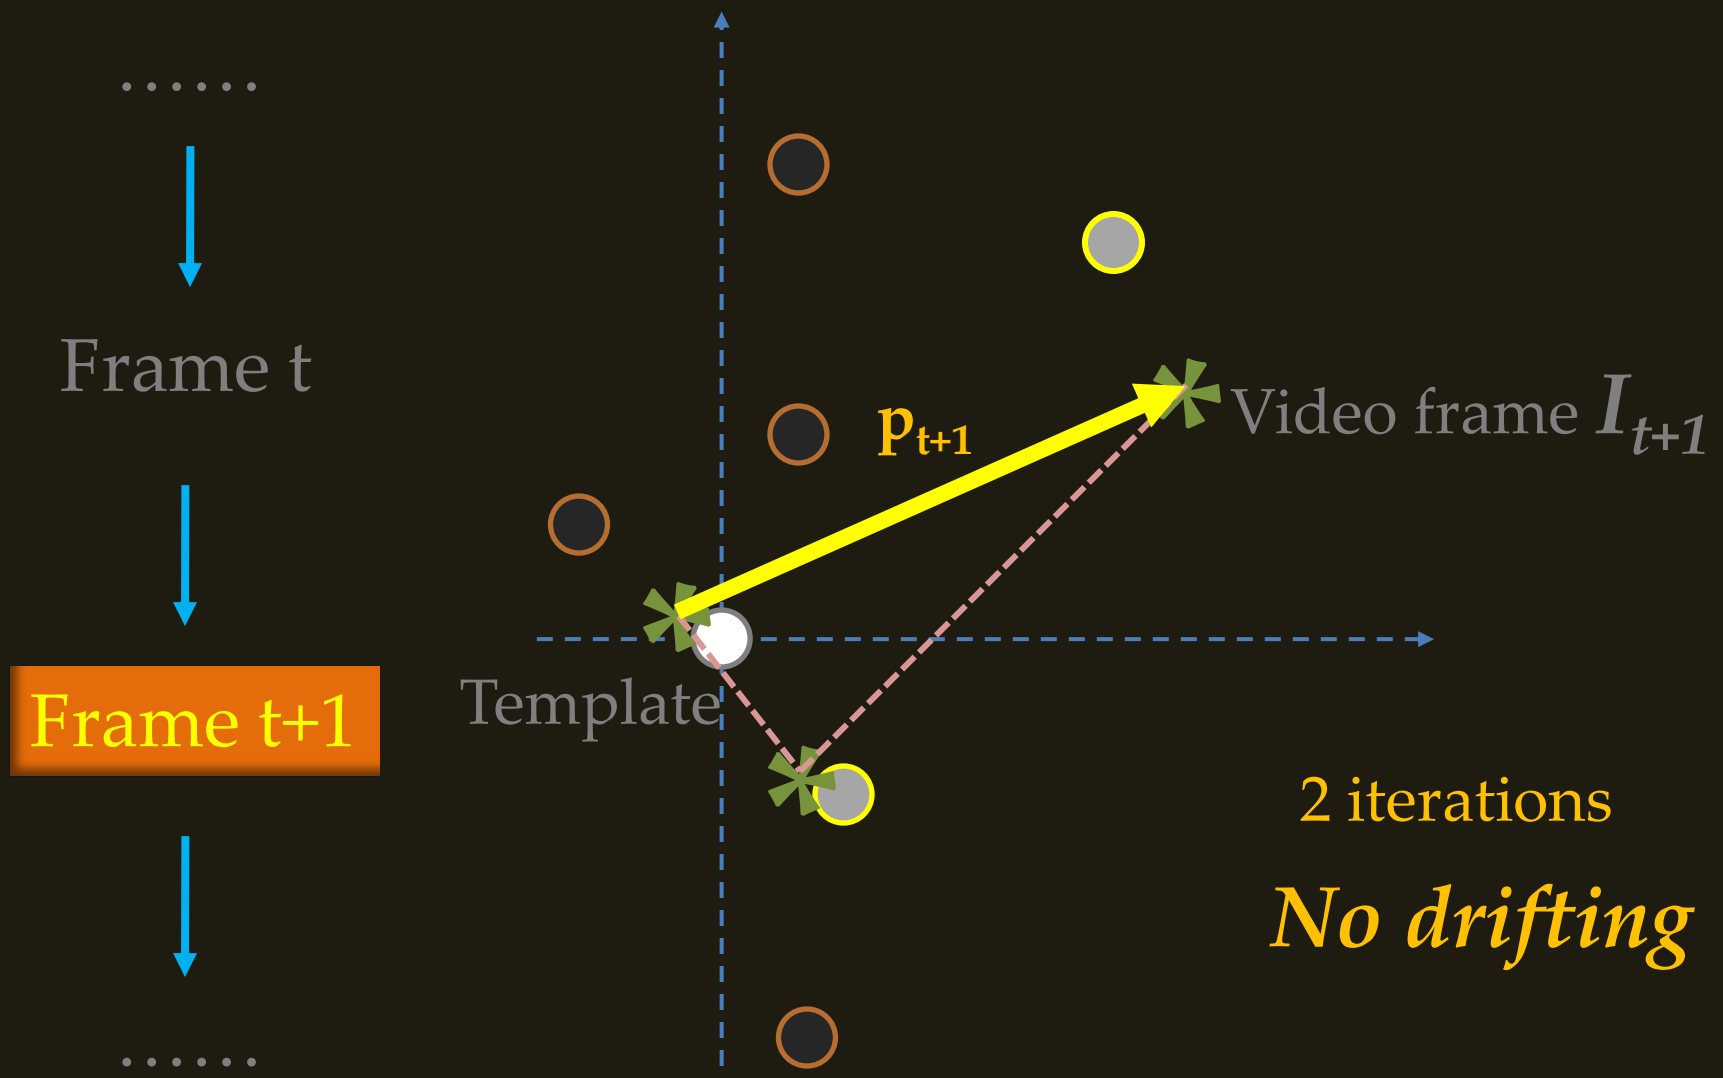

Water distortions

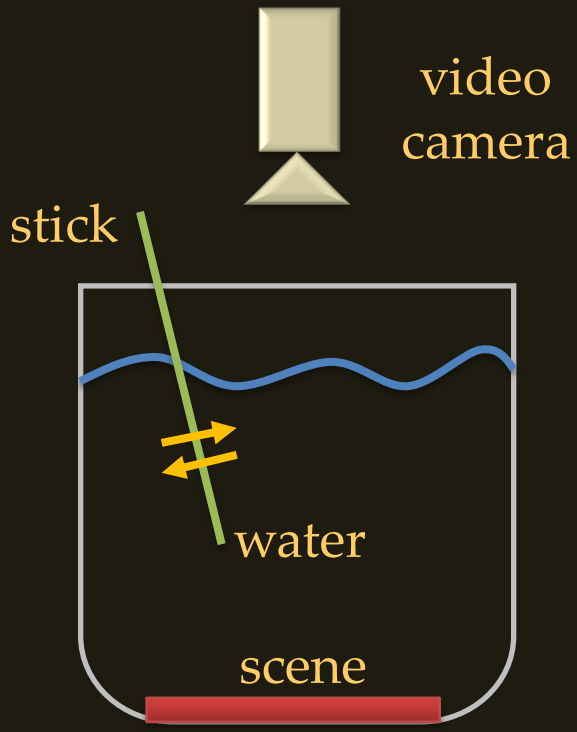

Experimental setup

Bases for water distortion

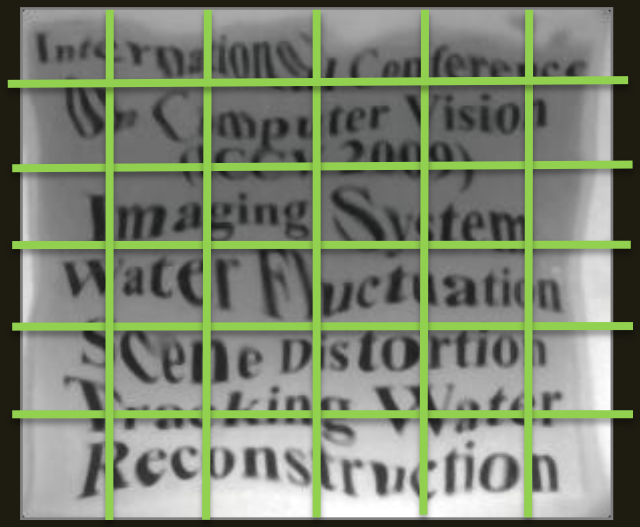

Overlapping partition

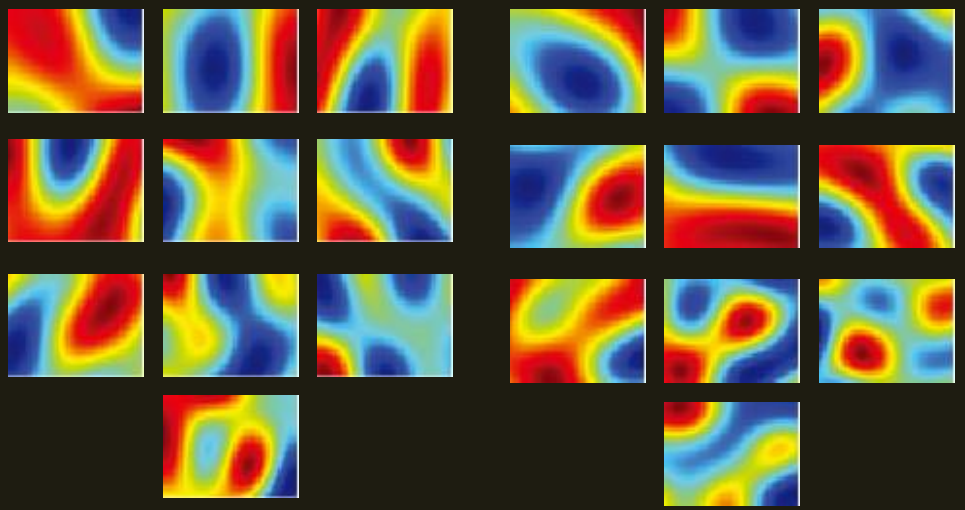

Water Bases

[Tian and Narasimhan, ICCV 09]

Correcting water distortions

Distorted image

Template

Our approach

Feature Matching

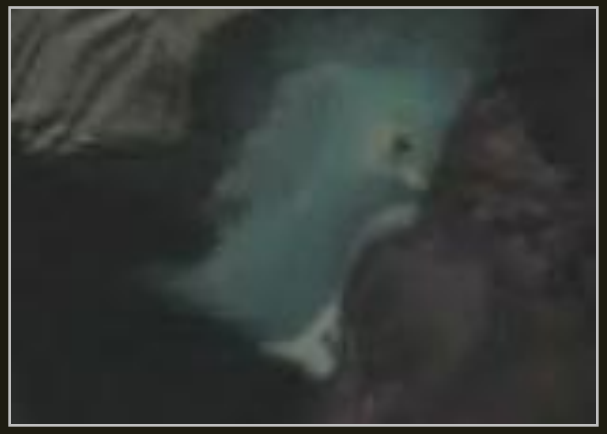

B-spline [Rueckert et al.]

Video rectification/Surface reconstruction

Original video

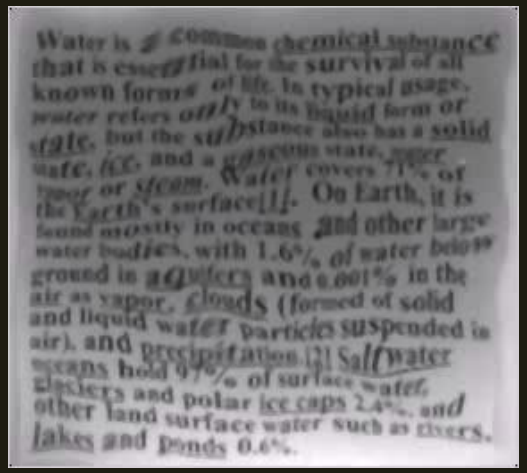

B-spline [Rueckert et al.]

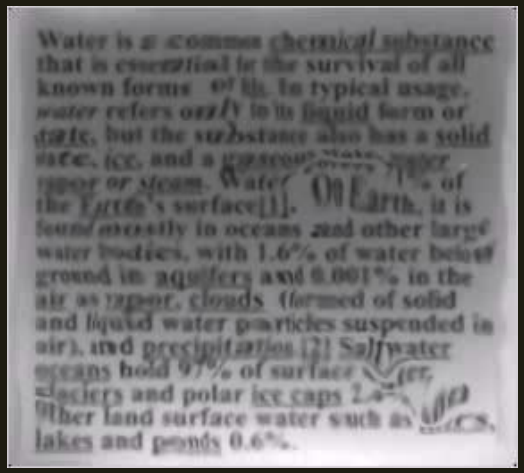

Our approach

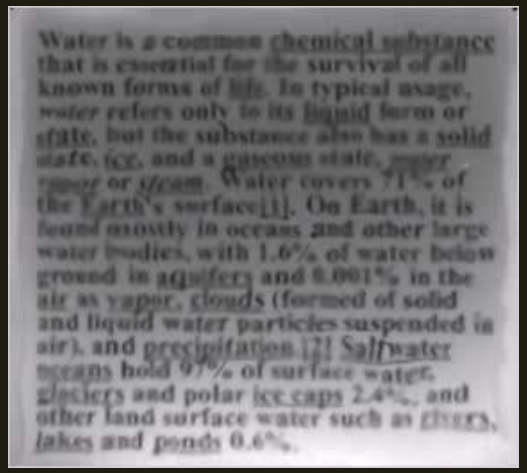

Water surface reconstruction

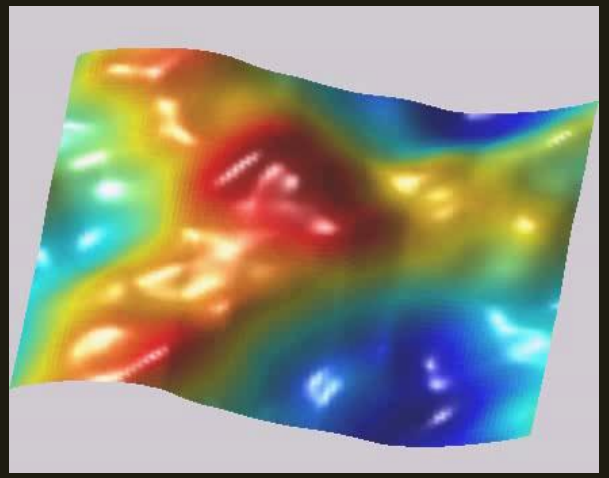

Video rectification

Original video

Rectified video

Video tracking

Video Tracking

Template tracking

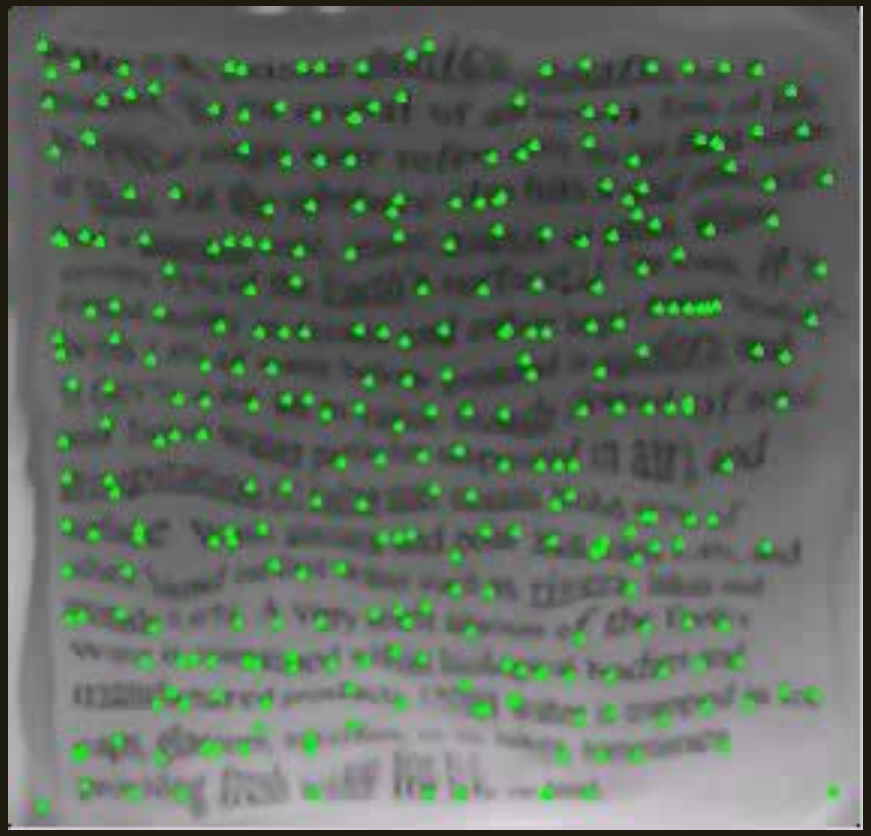

Our approach

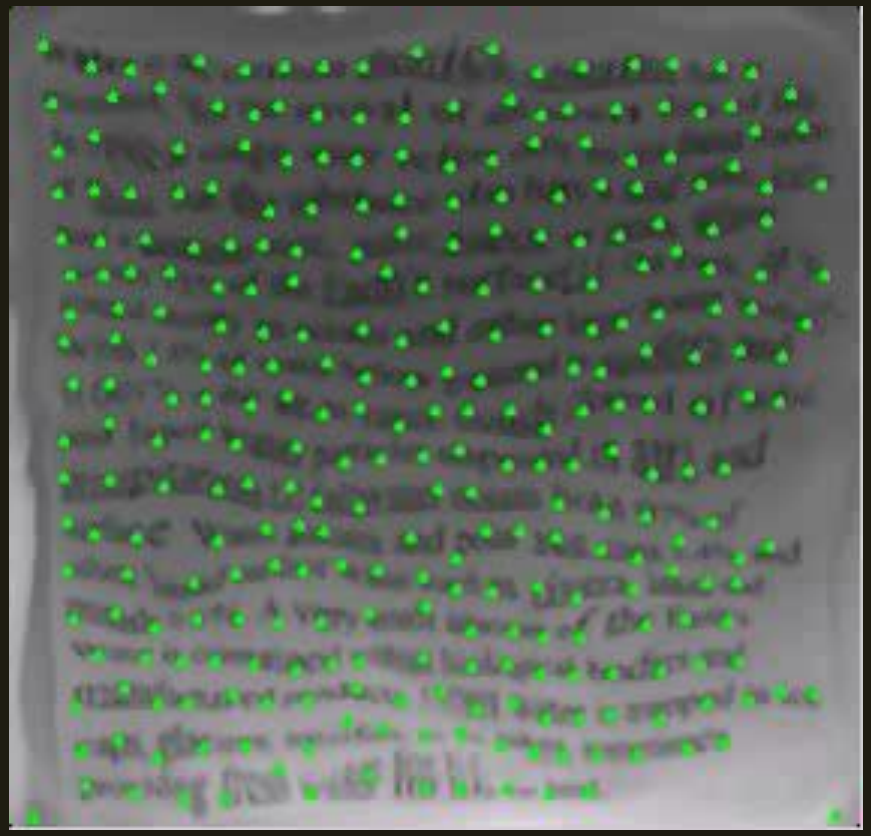

Cloth deformation

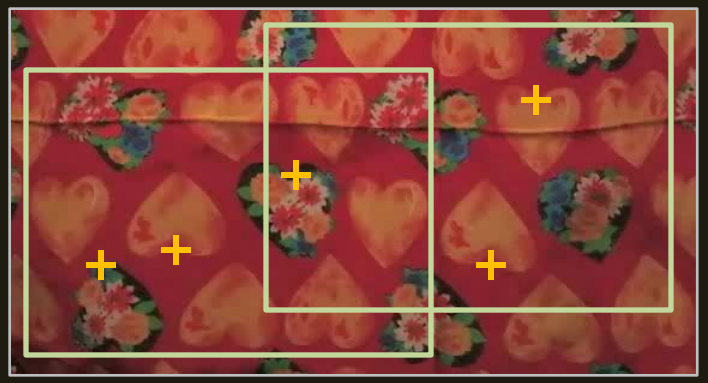

Affine Bases / Trackers

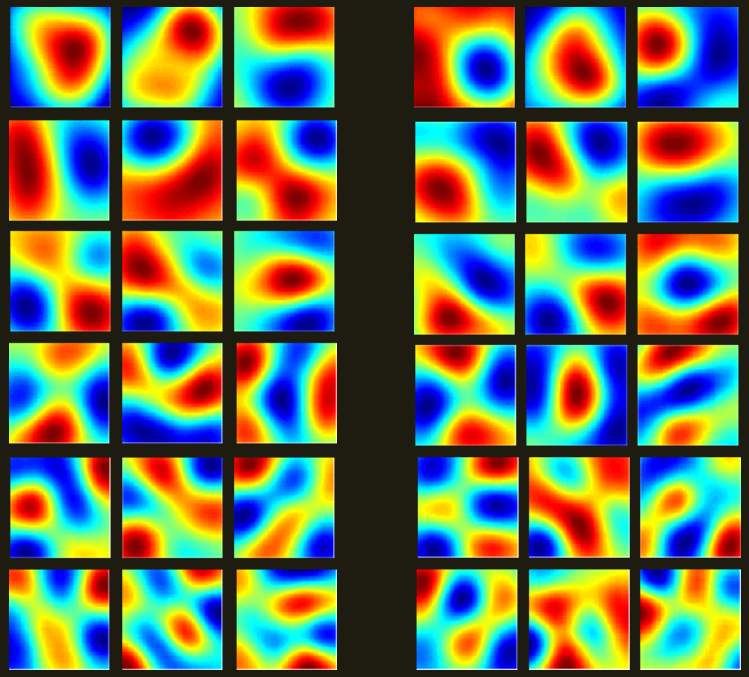

Locally smooth random Bases

Cloth tracking

Original video

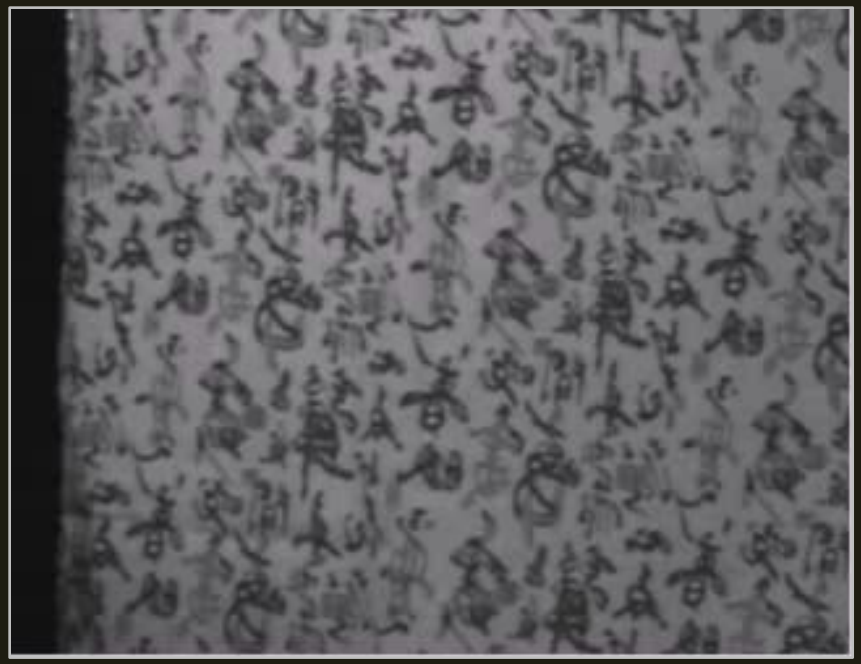

Tracking result

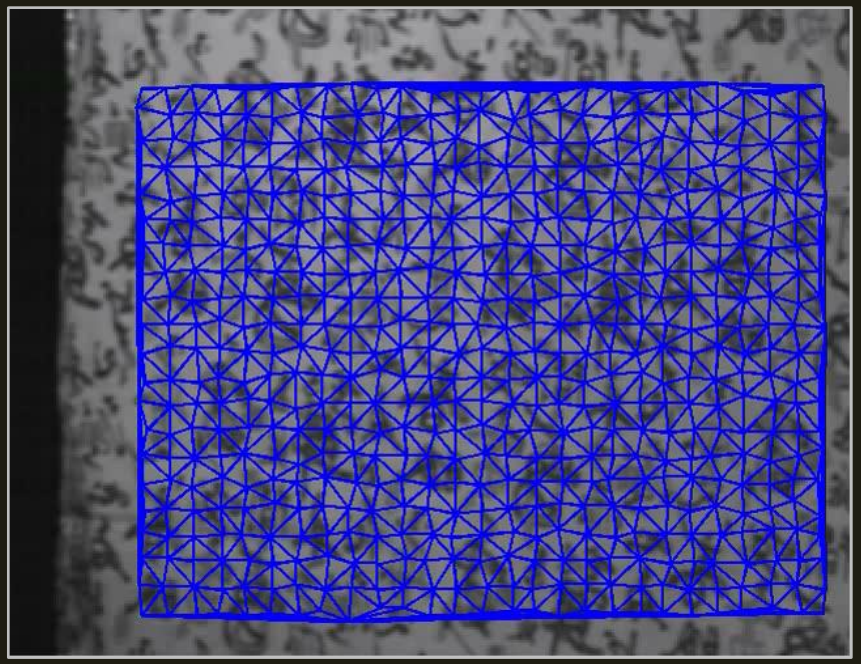

[Data from Taylor et.al, CVPR 2010]

Paper bending

Original video

Tracking result

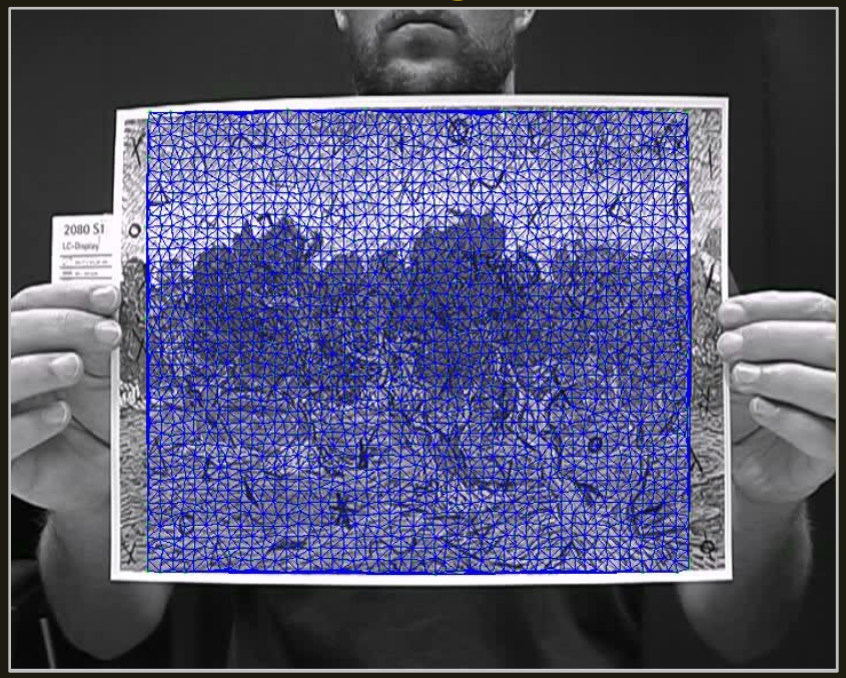

[Data from M. Salzmann et al. 2007]

Summary

- An iterative algorithm converges to the global optimum with much fewer training samples.

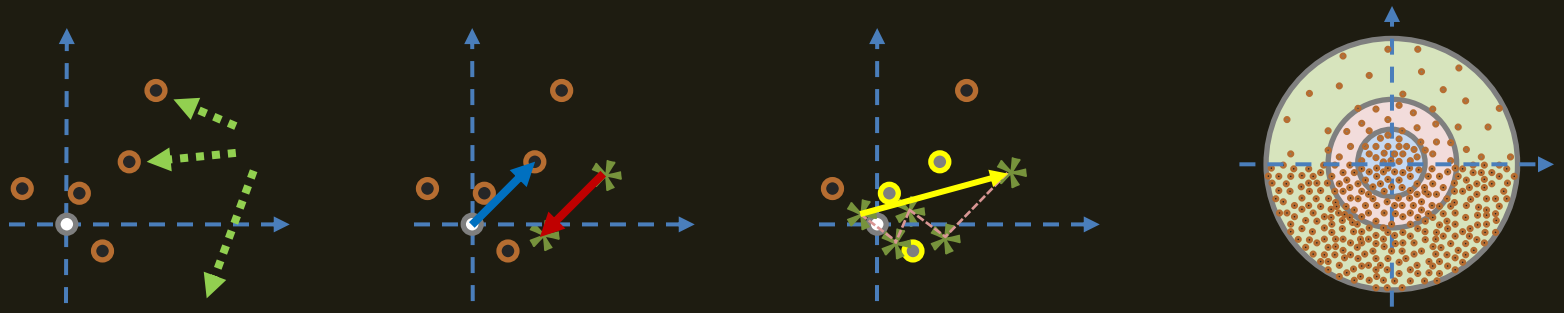

- Drift-free tracking, image & surface reconstruction

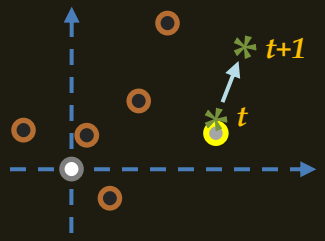

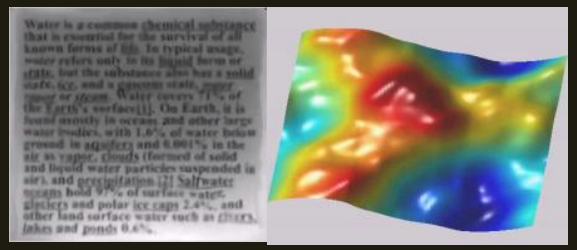

Thank you!

http://www.cs.cmu.edu/~ILIM
http://www.cs.cmu.edu/~yuandong