graphics and vision gravis

University of Basel

Case-study: Non-rigid Registration

Marcel Lüthi, 03. Juni 2019

Graphics and Vision Research Group
Department of Mathematics and Computer Science
University of Basel

The registration problem

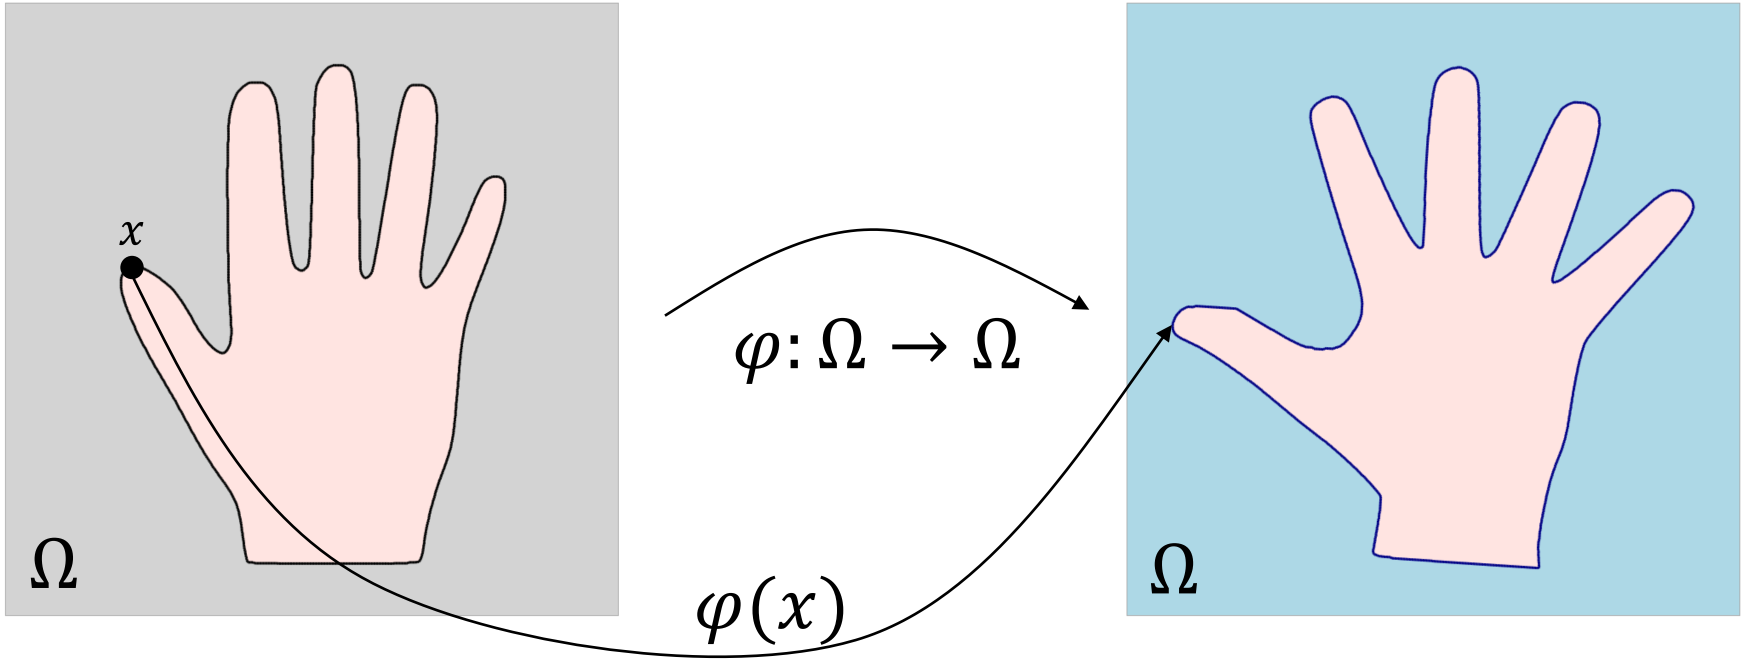

Reference:

$I_R: \Omega \to \mathbb{R}$

Target:

$I_T: \Omega \to \mathbb{R}$

Why is it important?

- Do automatic measurements
- Compare shapes
  - Statistics
  - Build statistical models
- Transfer labels and annotations
  - Atlas based segmentation

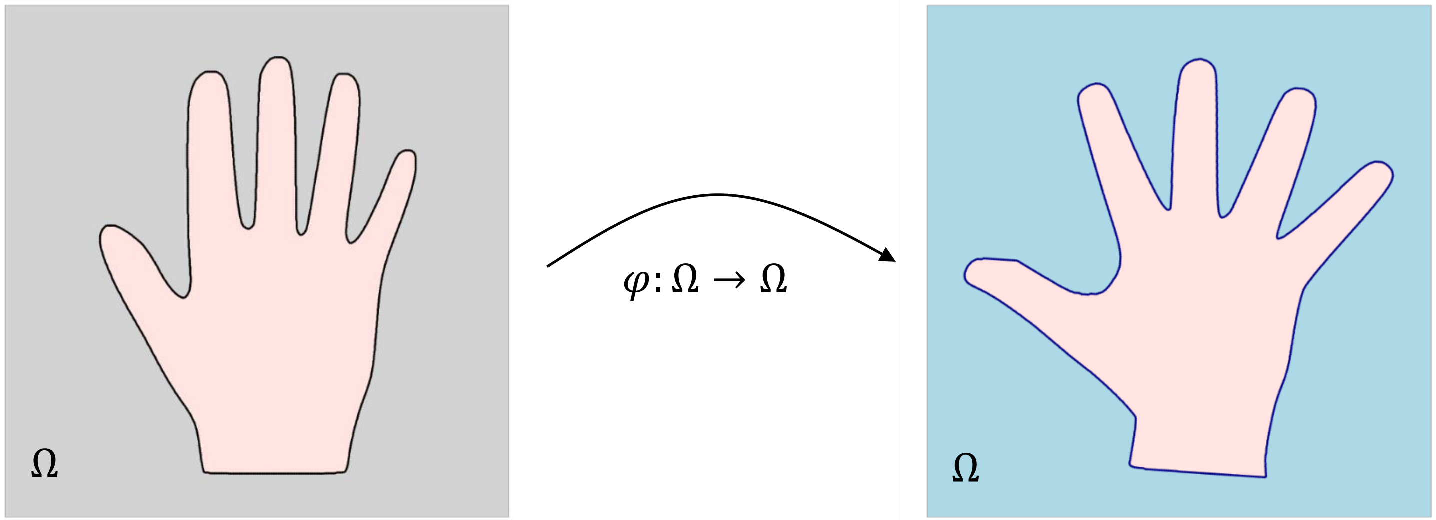

Maybe the most important problem in computer vision and medical image analysis

Registration as analysis by synthesis

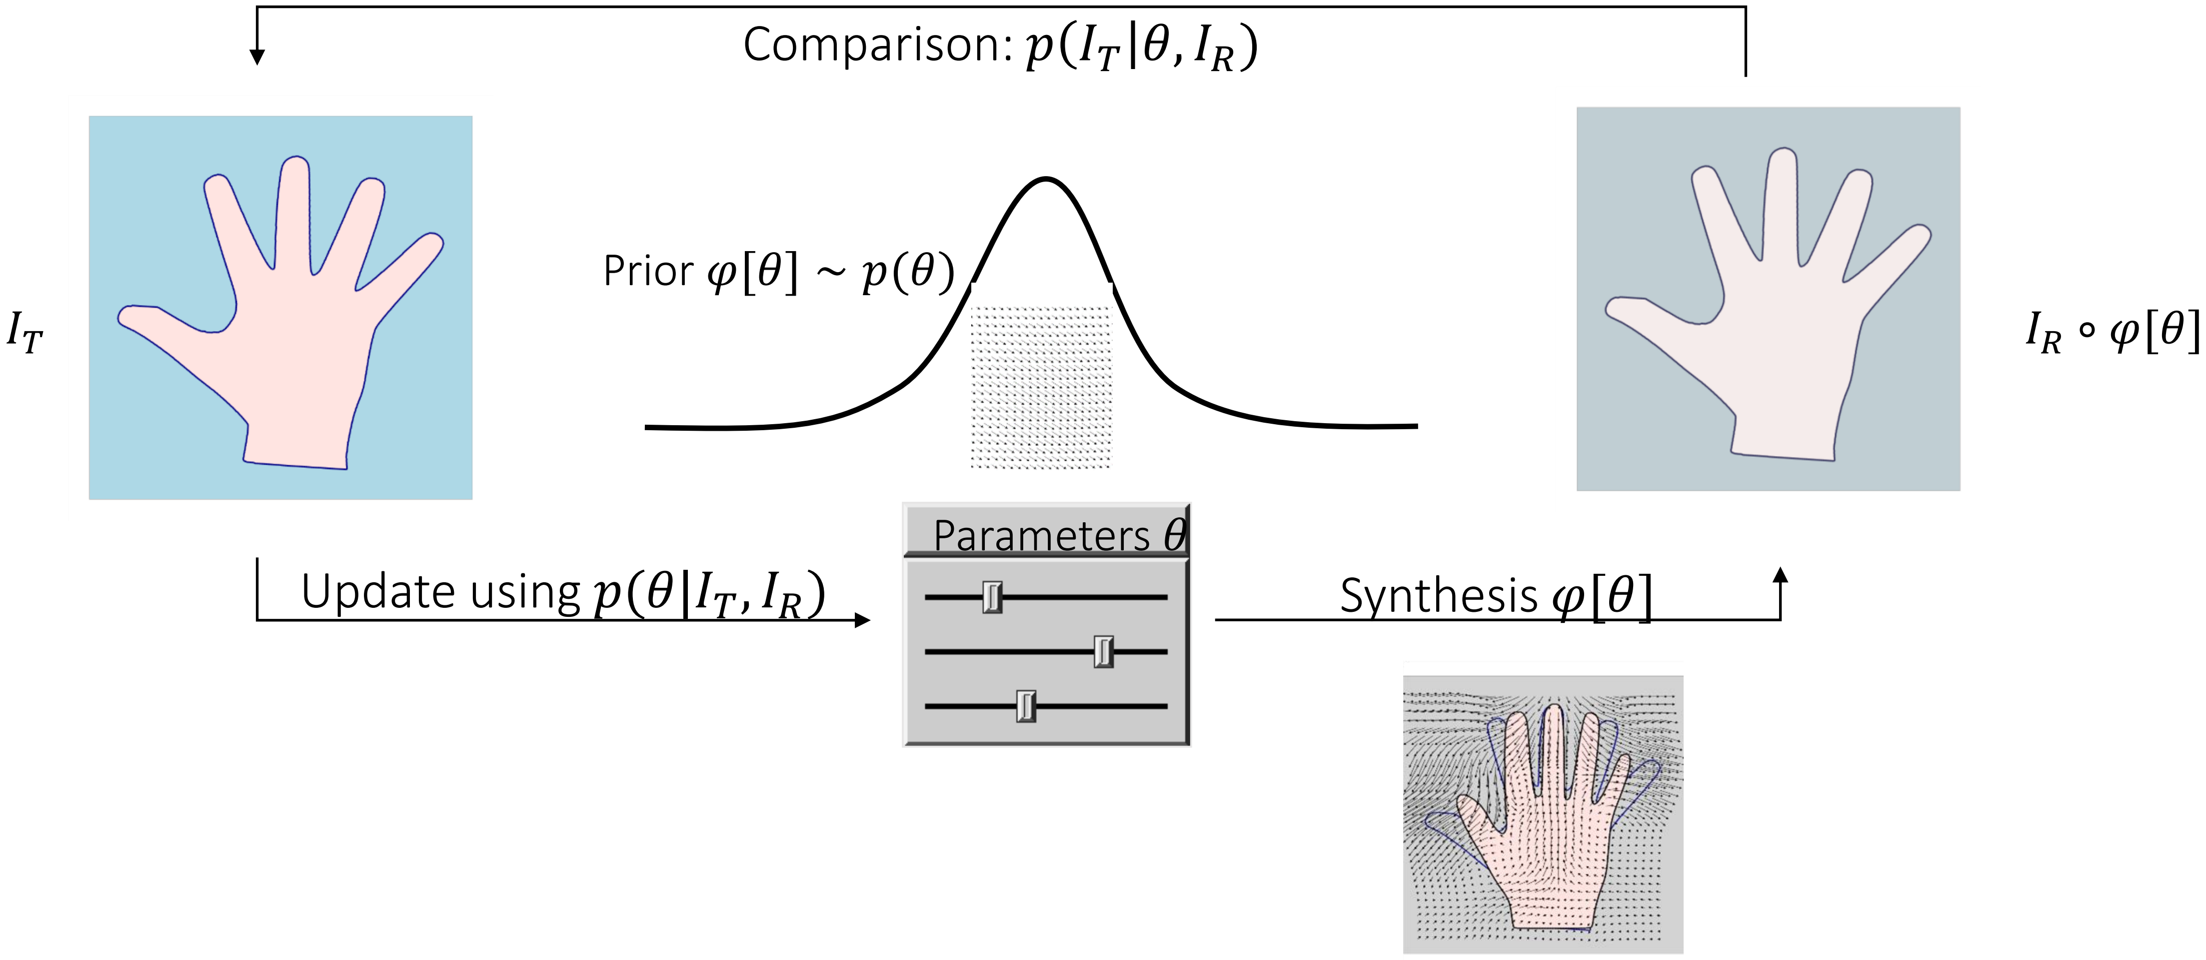

The registration problem

MAP-Estimation

$$\theta^* = \arg\max_\theta p(\theta|I_T, I_R) = \arg\max_\theta p(\theta)p(I_T|\theta, I_R)$$

Mapping $\varphi[\theta^*]$ is trade-off that

- how well does the mapping explain the target image (likelihood function)
- matches the prior assumptions (prior distribution)

The registration problem

$$\theta^* = \arg\max_\theta p(\theta | I_T, I_R) = \arg\max_\theta \cancel{p(\theta)} p(I_T | \theta, I_R)$$

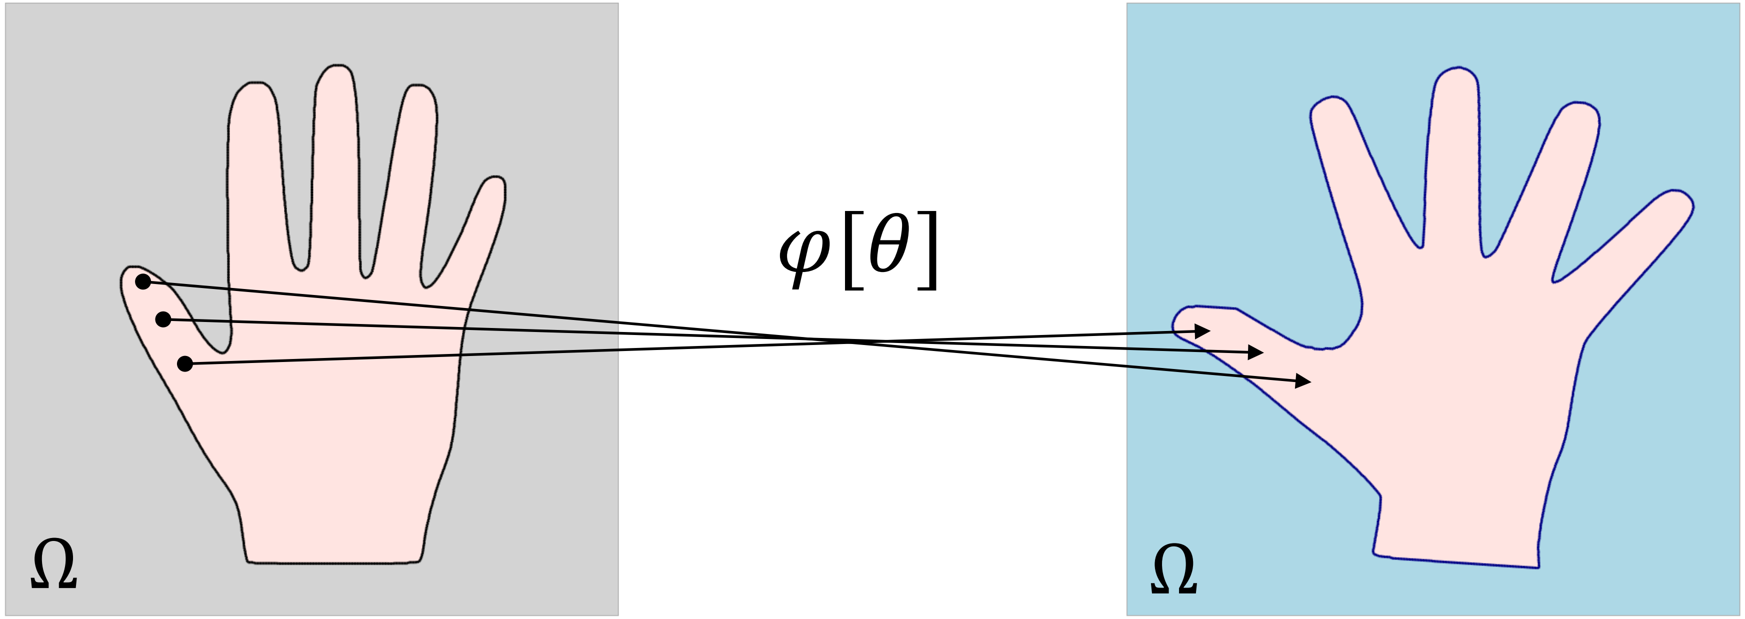

The registration problem

$$\theta^* = \arg\max_{\theta} p(\theta | I_T, I_R) = \arg\max_{\theta} p(\theta) p(I_T \cancel{|\theta}, I_R)$$

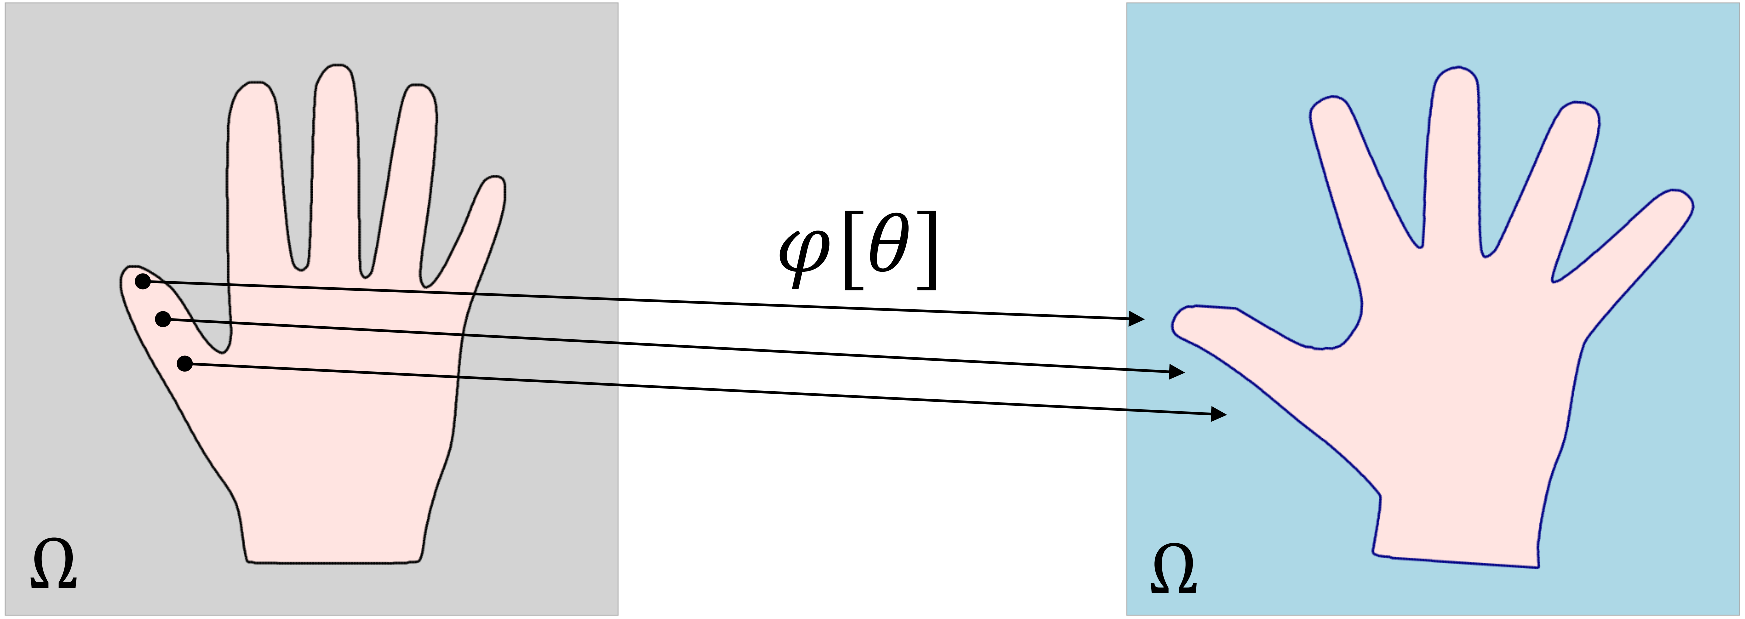

The registration problem

$$\theta^* = \arg\max_\theta p(\theta|I_T, I_R) = \arg\max_\theta p(\theta)p(I_T|\theta, I_R)$$

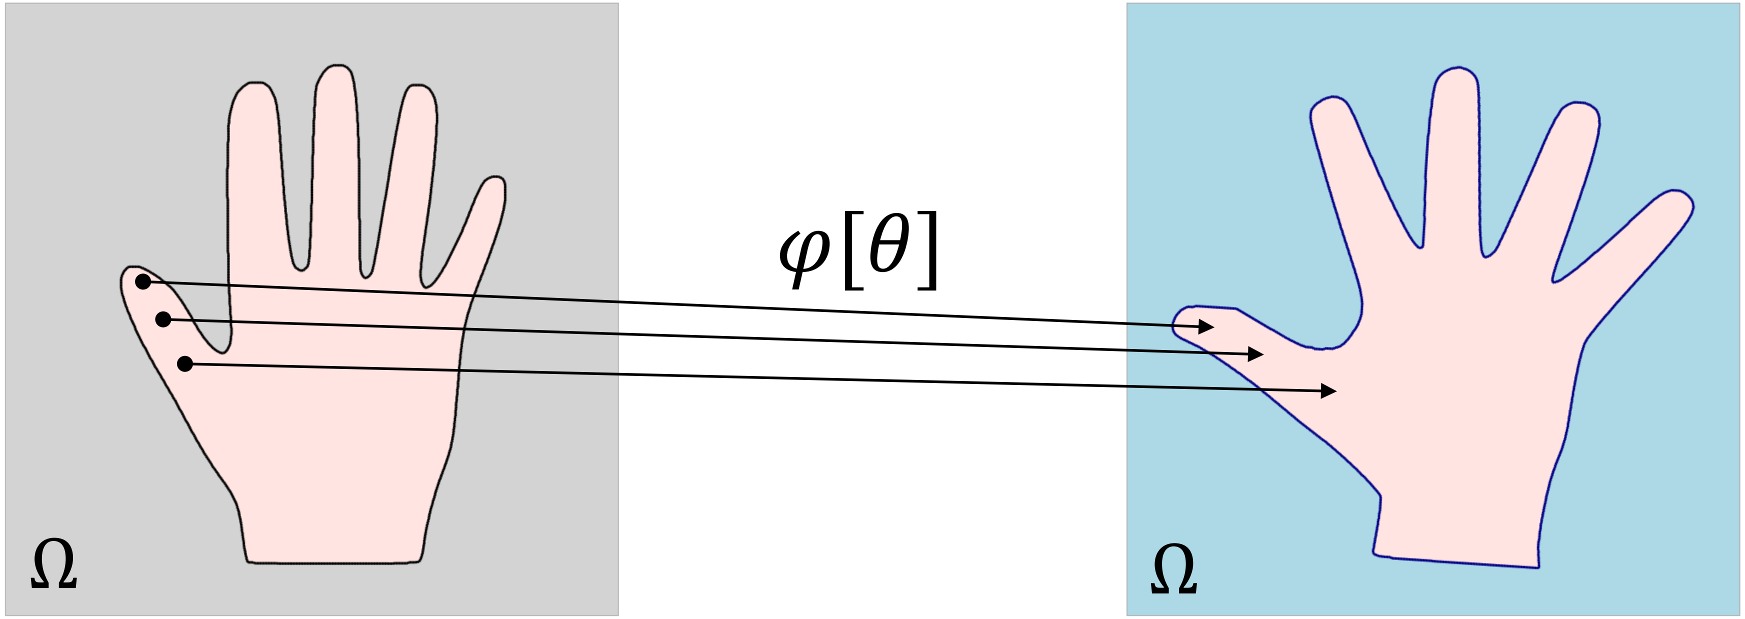

The registration problem

Probabilistic formulation

$$\varphi^* = \arg\max_{\varphi} p(\varphi | I_T, I_R) = \arg\max_{\varphi} p(\varphi) p(I_T | \varphi, I_R)$$

Main questions:

- How do we represent the mapping?
- How do we define the prior?
- What is the likelihood function?
- How can we solve the optimization problem?

Representation of the mapping φ

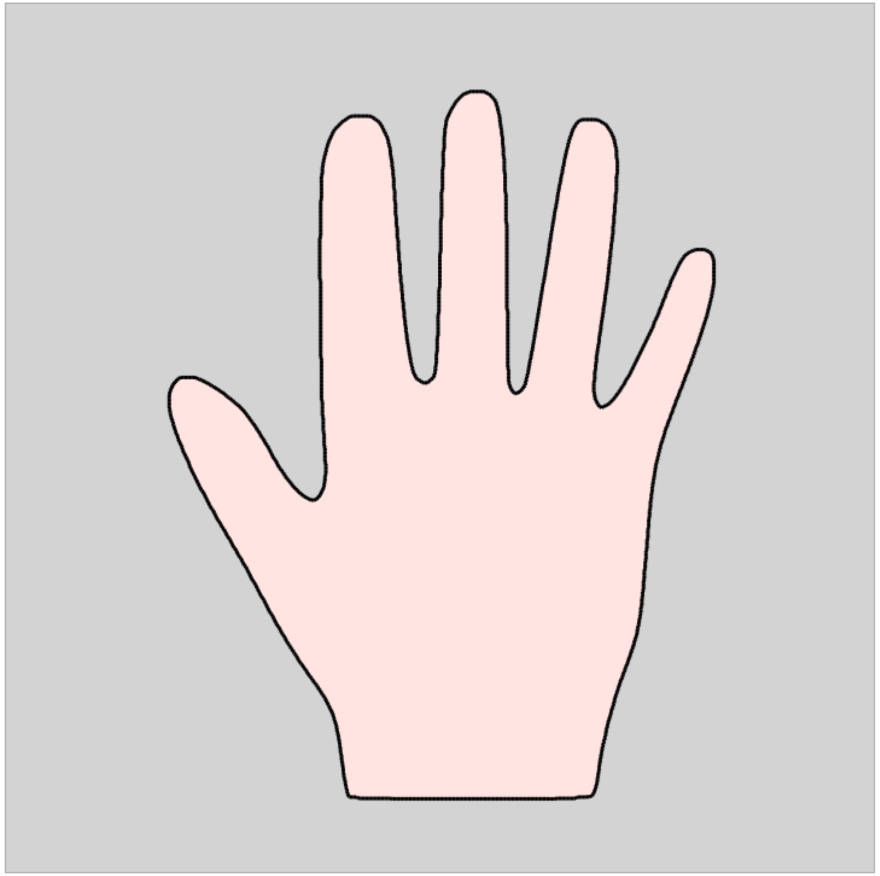

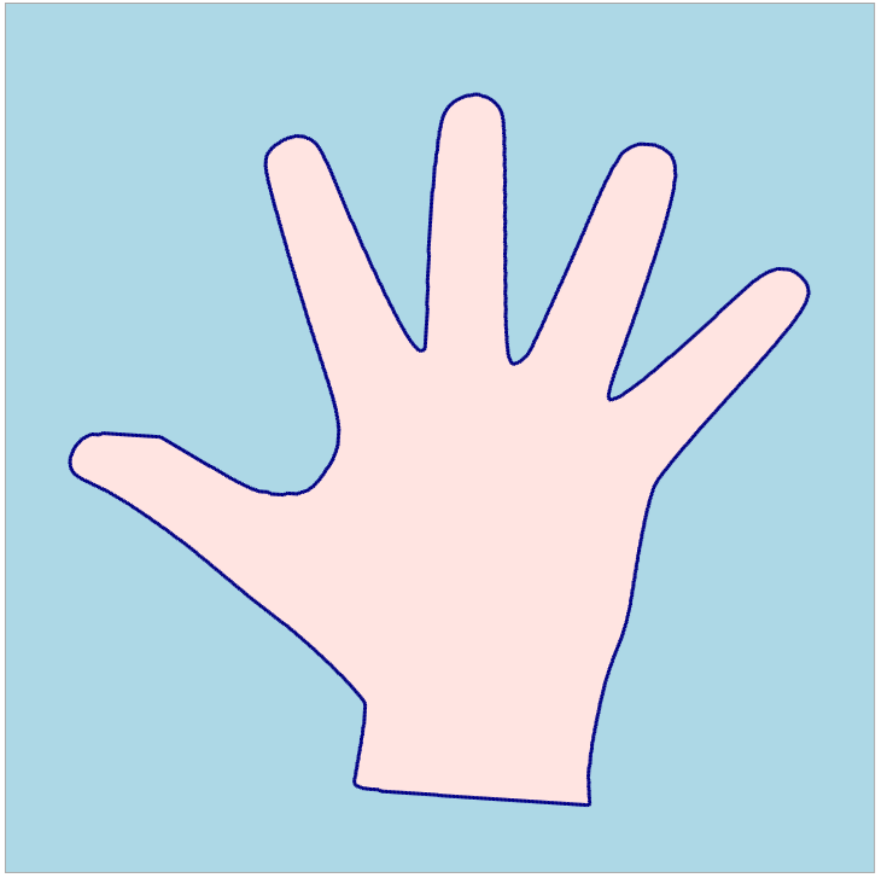

Representation of the mapping φ

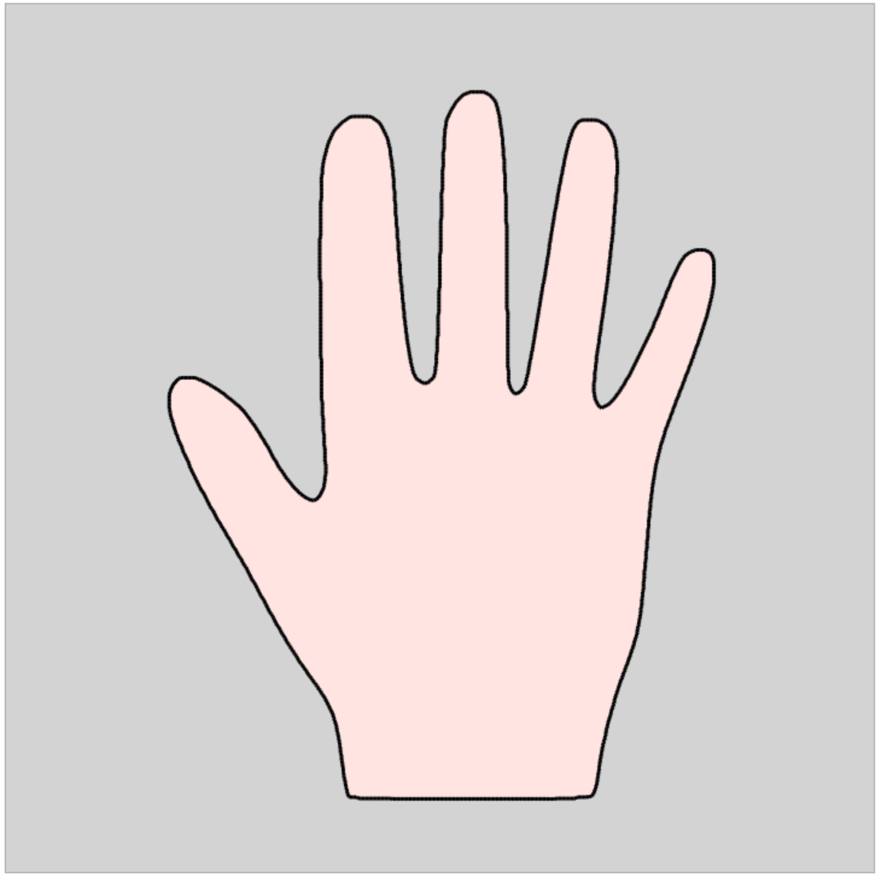

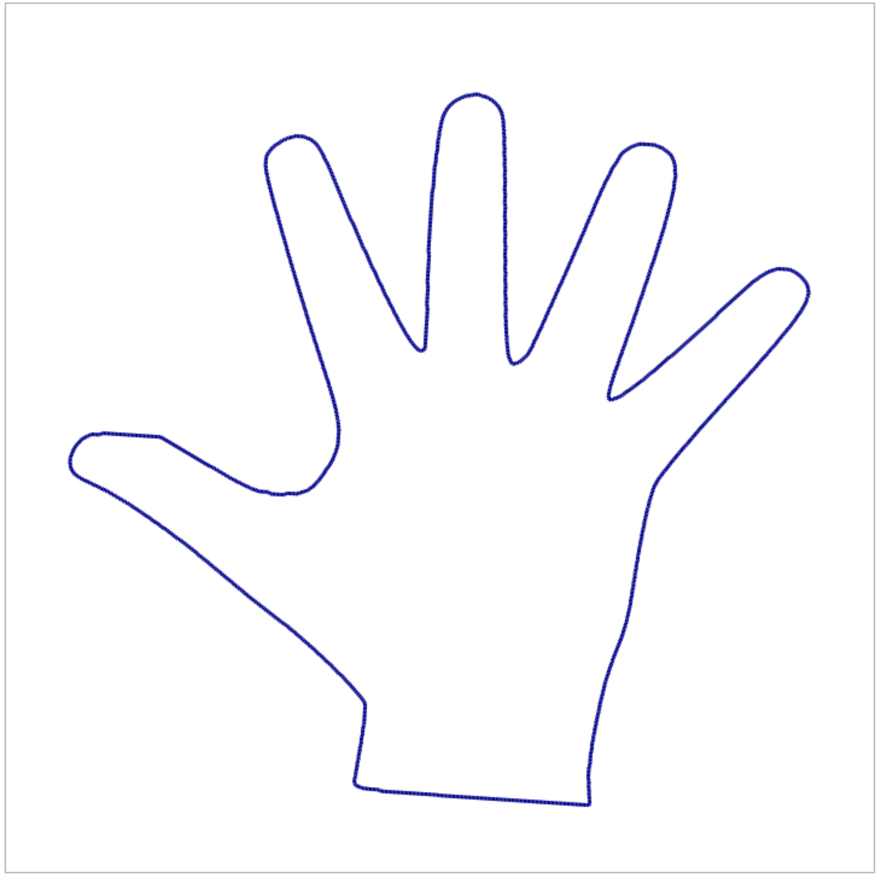

Representation of the mapping φ

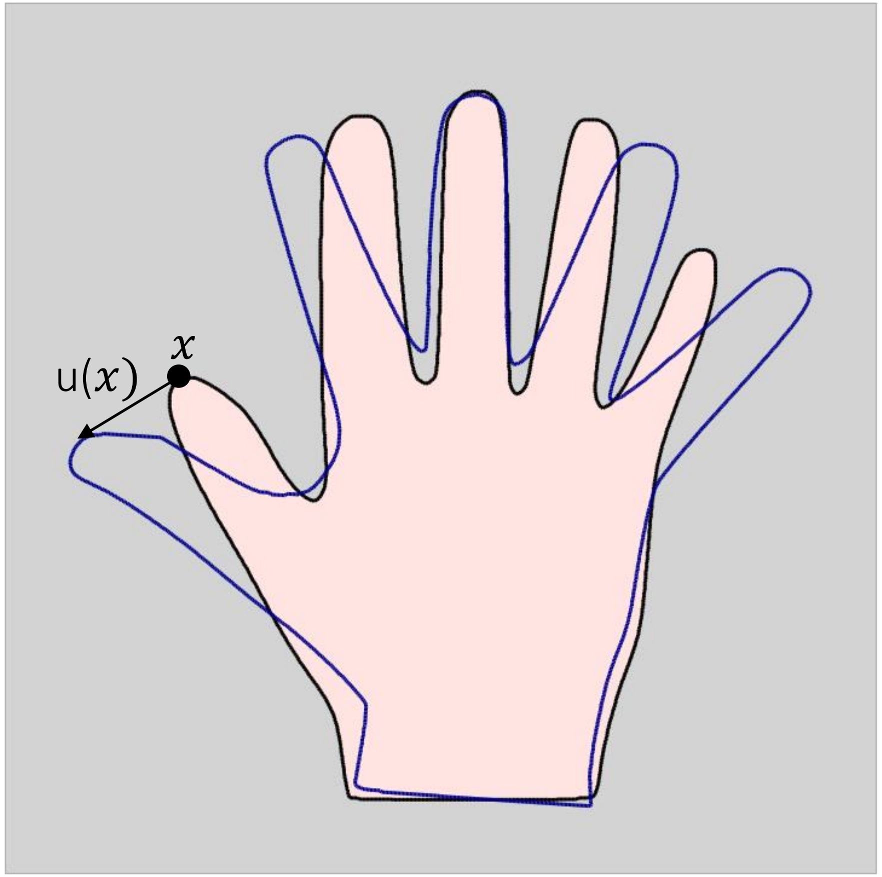

Assumption:
Images are rigidly aligned

- Mapping can be represented as a displacement vector field:

$$\varphi(x) = x + u(x)$$
$$u : \Omega \to \mathbb{R}^d$$

Representation of the mapping φ

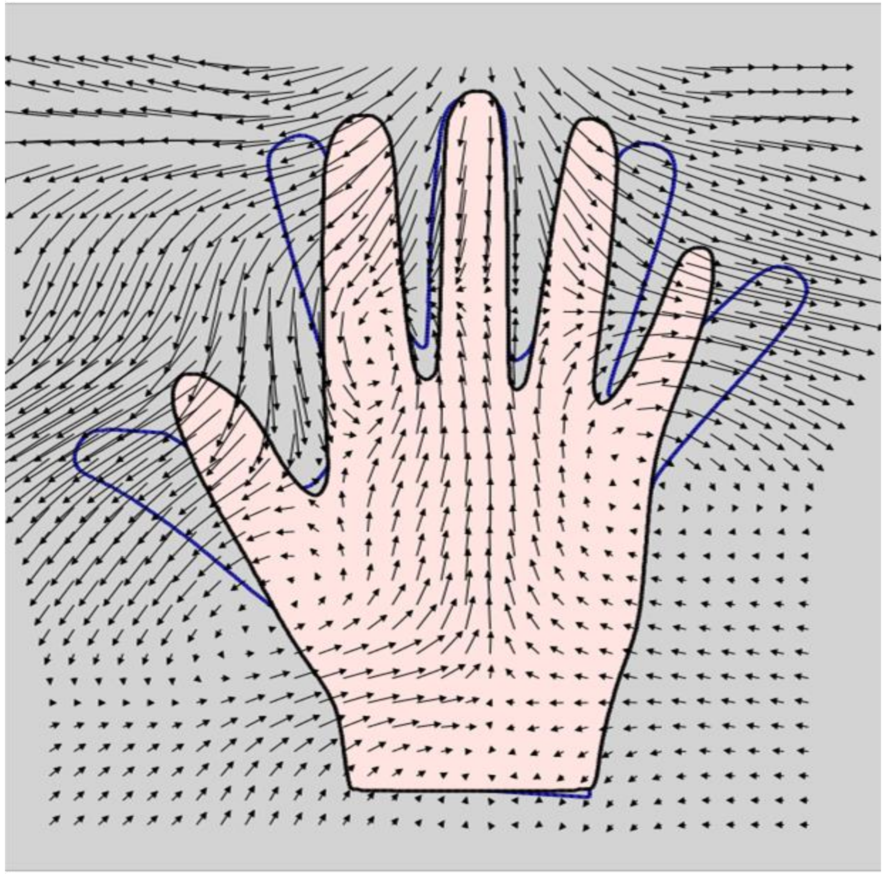

Assumption:
Images are rigidly aligned

- Mapping can be represented as a displacement vector field:

$$\varphi(x) = x + u(x)$$
$$u : \Omega \to \mathbb{R}^d$$

Observation:
Knowledge of u and I_R allows us to synthesize target image I_T

Registration as analysis by synthesis

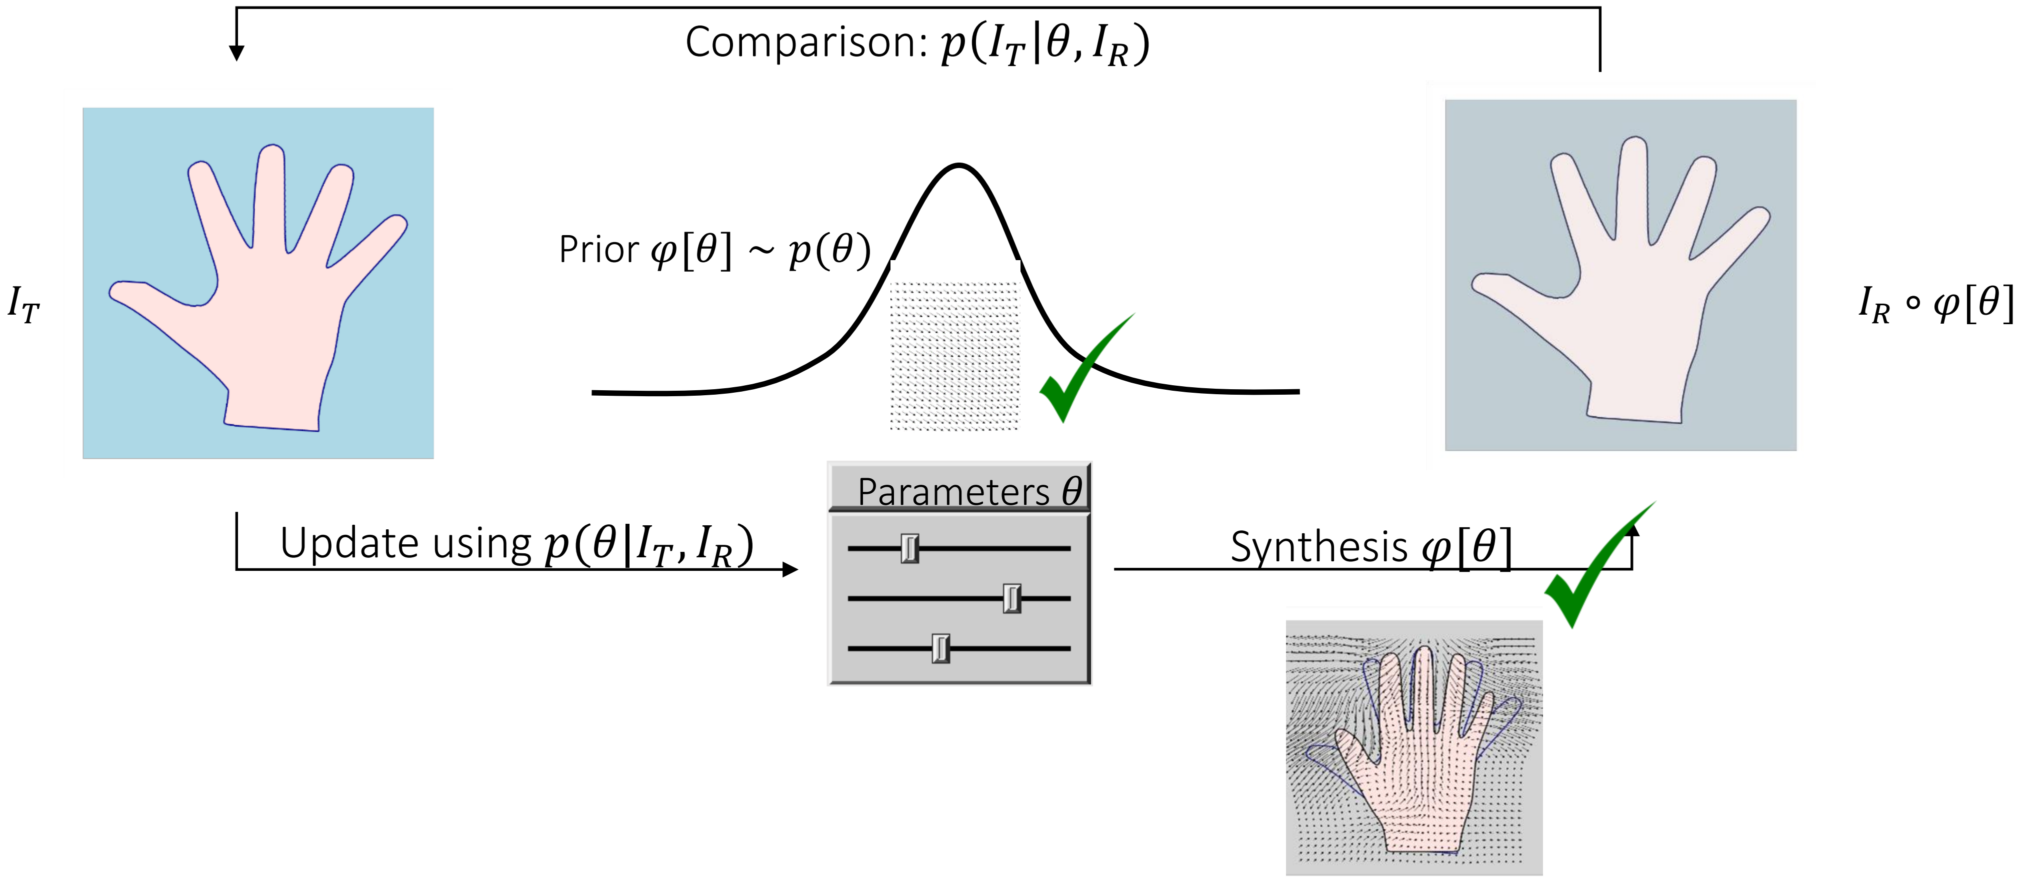

Priors

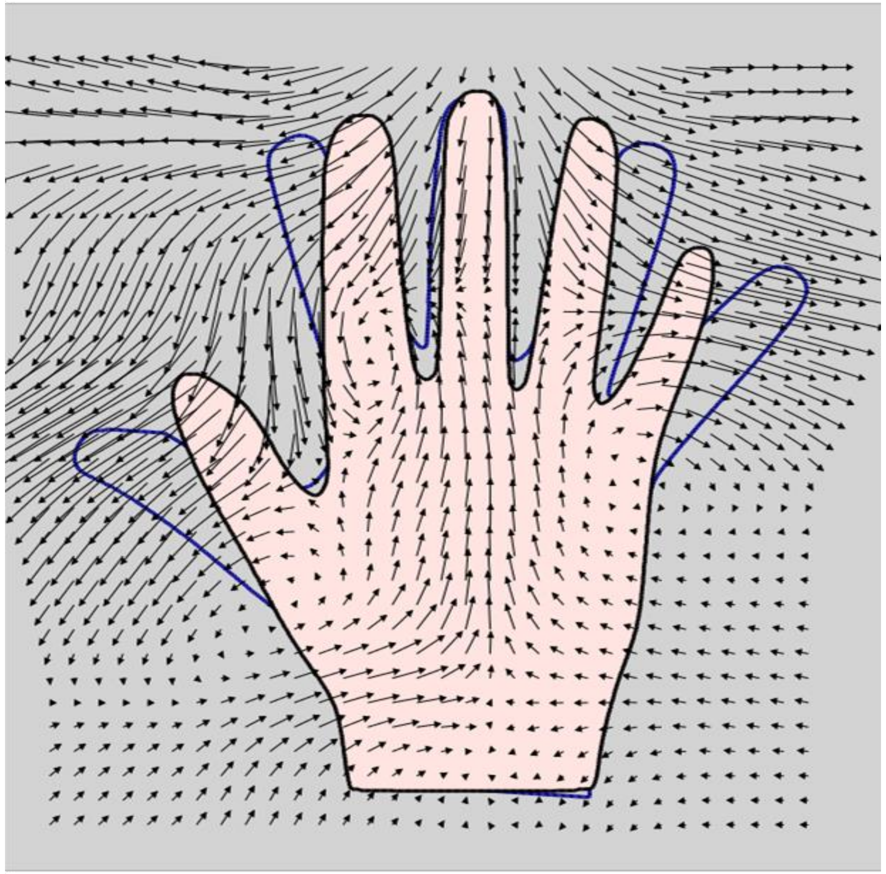

Define the Gaussian process

$$u \sim GP(\mu, k)$$

with mean function

$$\mu: \Omega \to \mathbb{R}^2$$

and covariance function

$$k: \Omega \times \Omega \to \mathbb{R}^{2\times 2} \, .$$

Example prior: Smooth 2D deformations

Zero mean:

$$\mu(x) = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

Squared exponential covariance function (Gaussian kernel)

$$k(x, x') = \begin{pmatrix} s_1 \exp\left(-\frac{\|x - x'\|^2}{\sigma_1^2}\right) & 0 \\ 0 & s_2 \exp\left(-\frac{\|x - x'\|^2}{\sigma_2^2}\right) \end{pmatrix}$$

Example prior: Smooth 2D deformations

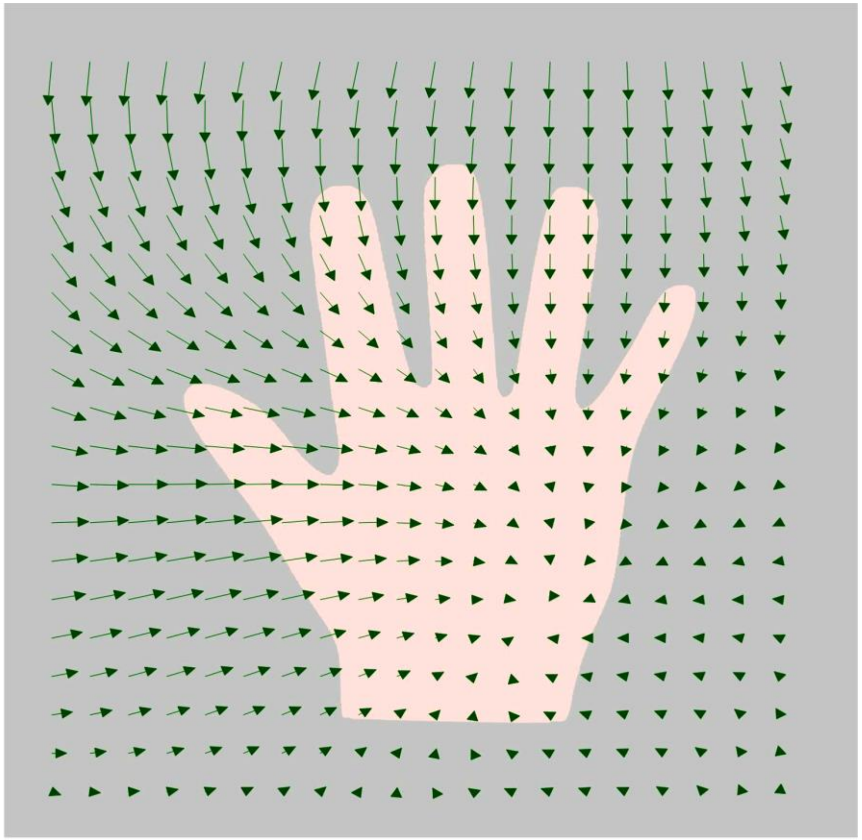

$s_1 = s_2$ small, $\sigma_1 = \sigma_2$ large

Example prior: Smooth 2D deformations

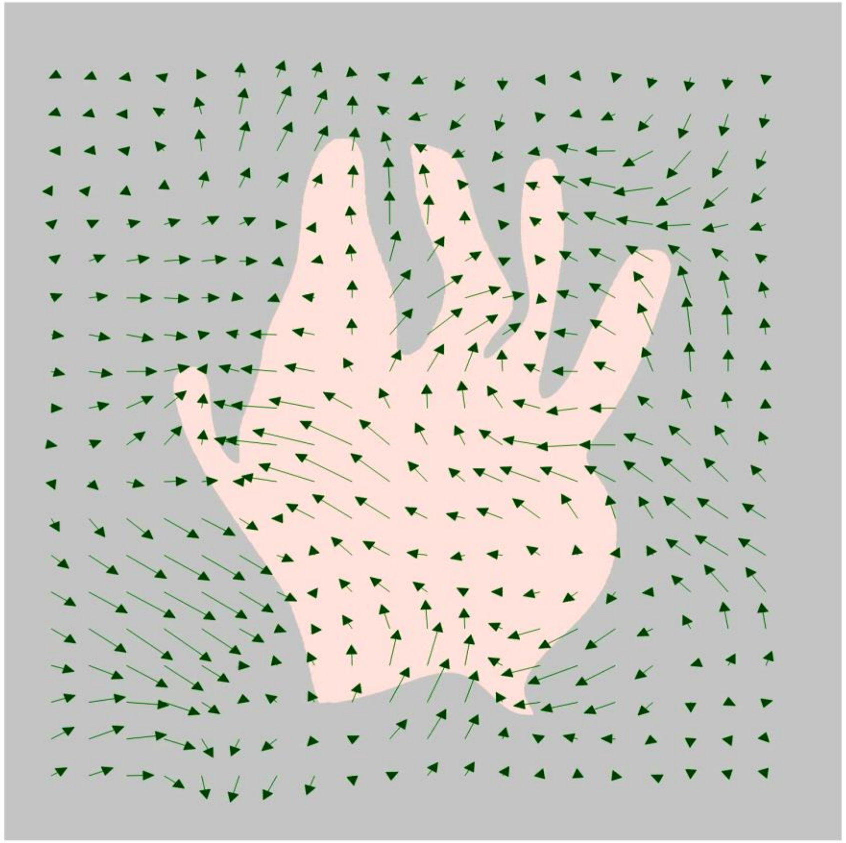

$s_1 = s_2$ small, $\sigma_1 = \sigma_2$ small

Example prior: Smooth 2D deformations

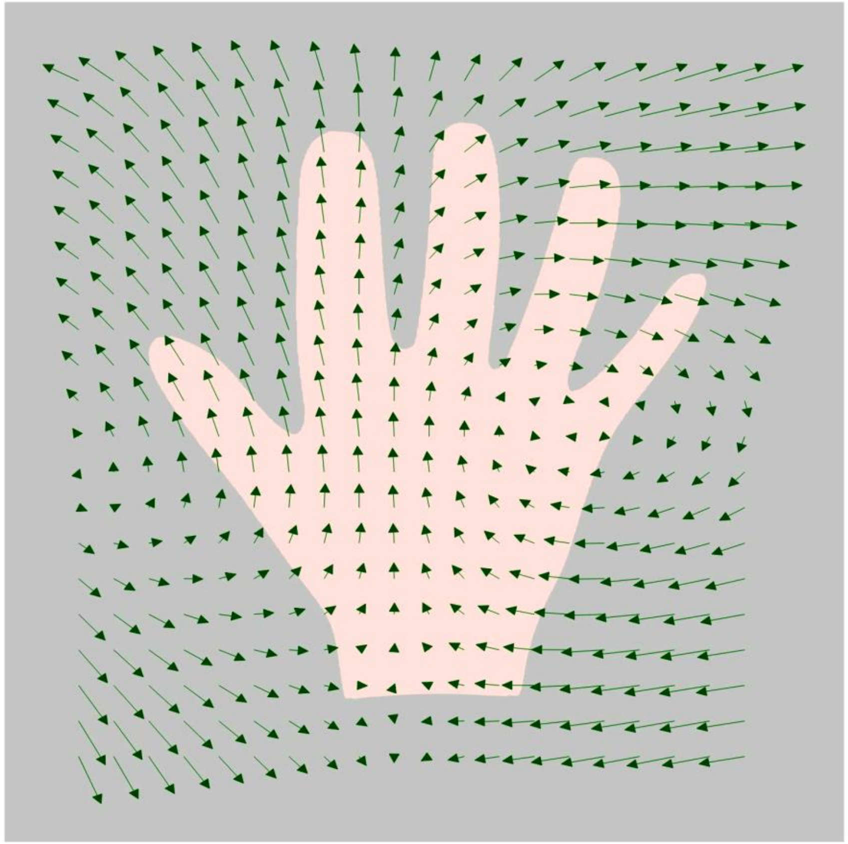

$s_1 = s_2$ large, $\sigma_1 = \sigma_2$ large

Parametric representation of Gaussian process

Represent $GP(\mu, k)$ using only the first r components of its KL-Expansion

$$u = \mu + \sum_{i=1}^{r} \alpha_i \sqrt{\lambda_i}\, \phi_i, \qquad \alpha_i \sim N(0,1)$$

- We have a finite, parametric representation of the process.
- We know the pdf for a deformation u

$$p(u[\alpha]) = p(\alpha) = \prod_{i=1}^{r} \frac{1}{\sqrt{2\pi}} \exp(-\alpha_i^2/2) = \frac{1}{Z} \exp(-\frac{1}{2} \|\alpha\|^2)$$

Registration as analysis by synthesis

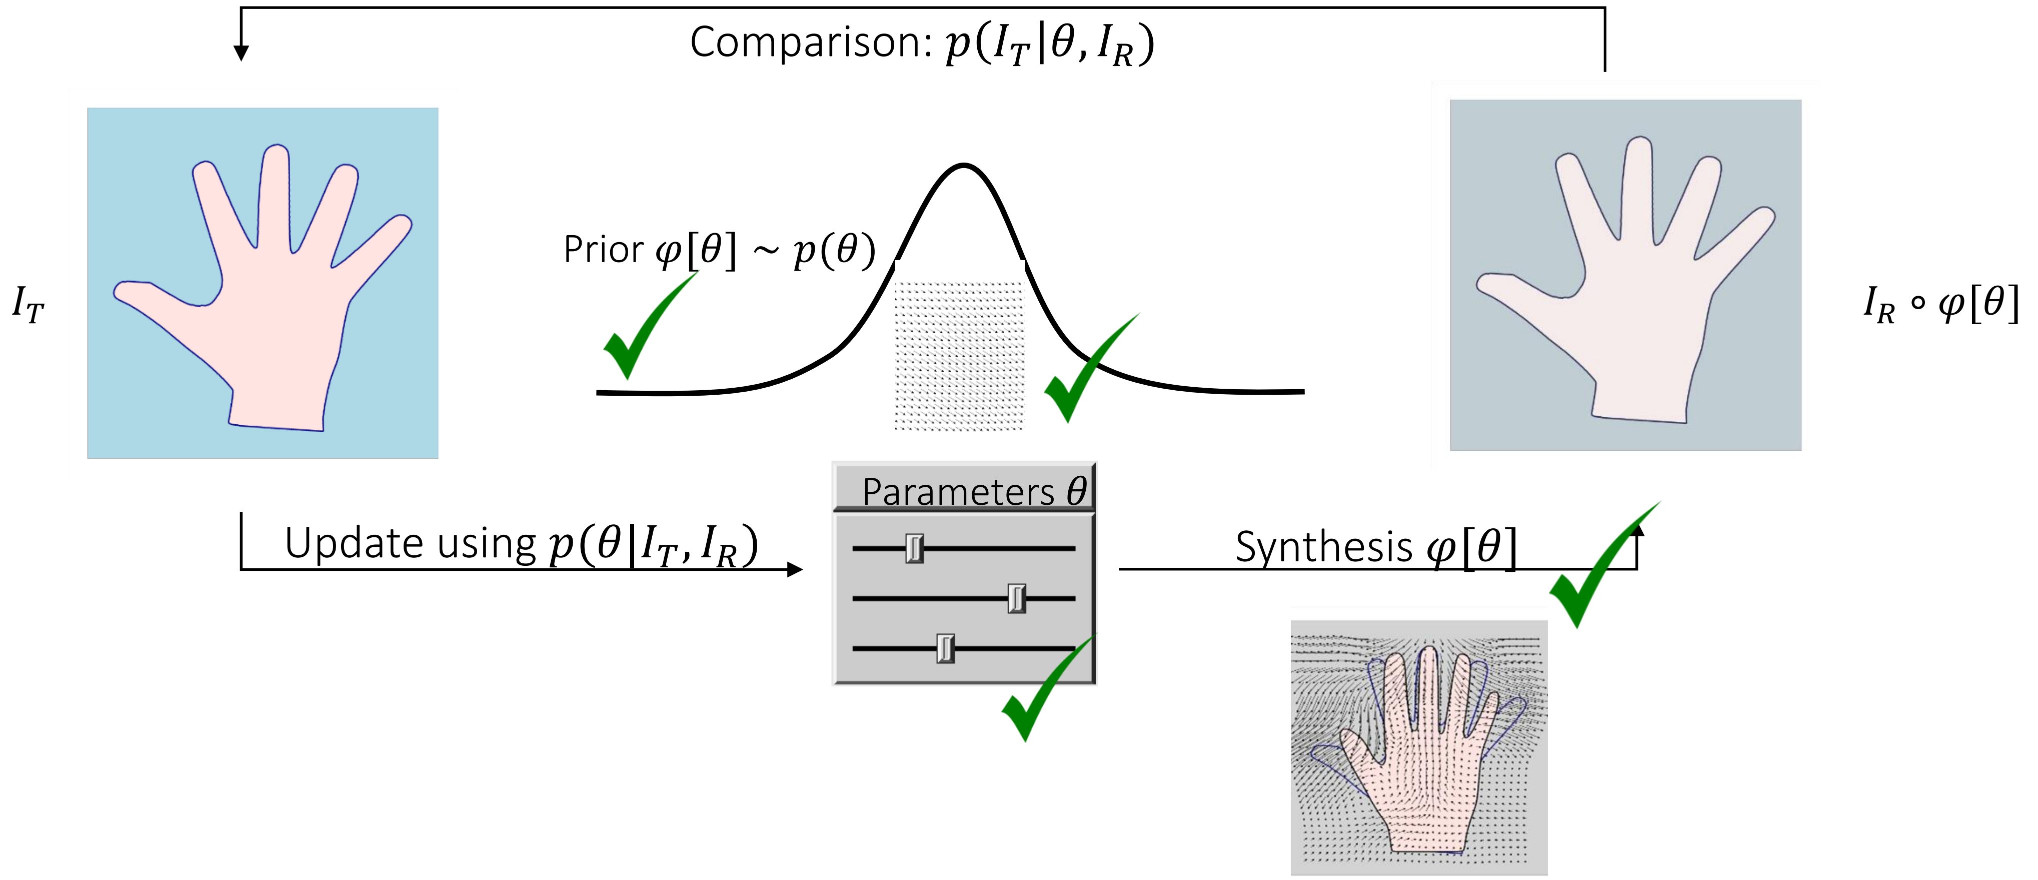

Likelihood function: Image registration

Images are similar when the intensities match

Assumptions:

- Corresponding points have the same image intensity (up to i.i.d. noise)

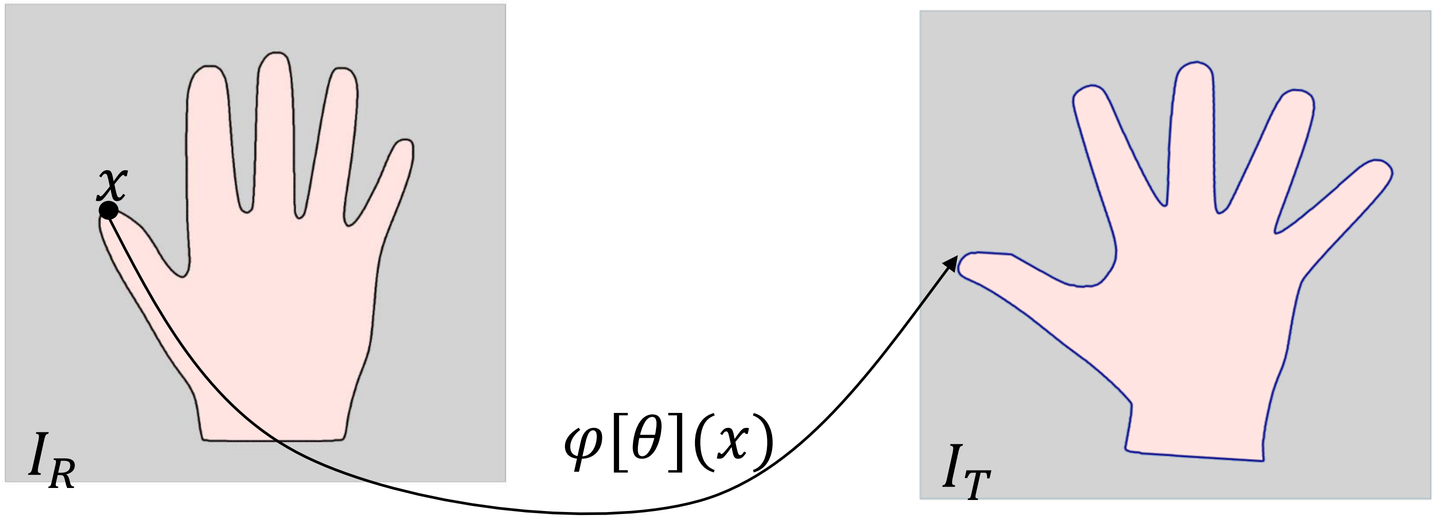

$$p(I_T(\varphi[\theta](x))|I_R, \theta, x) \sim N(I_R(x), \sigma^2)$$

Likelihood function: Image registration

Images are similar when the intensities match

Assumptions:

- Corresponding points have the same image intensity (up to i.i.d. noise)

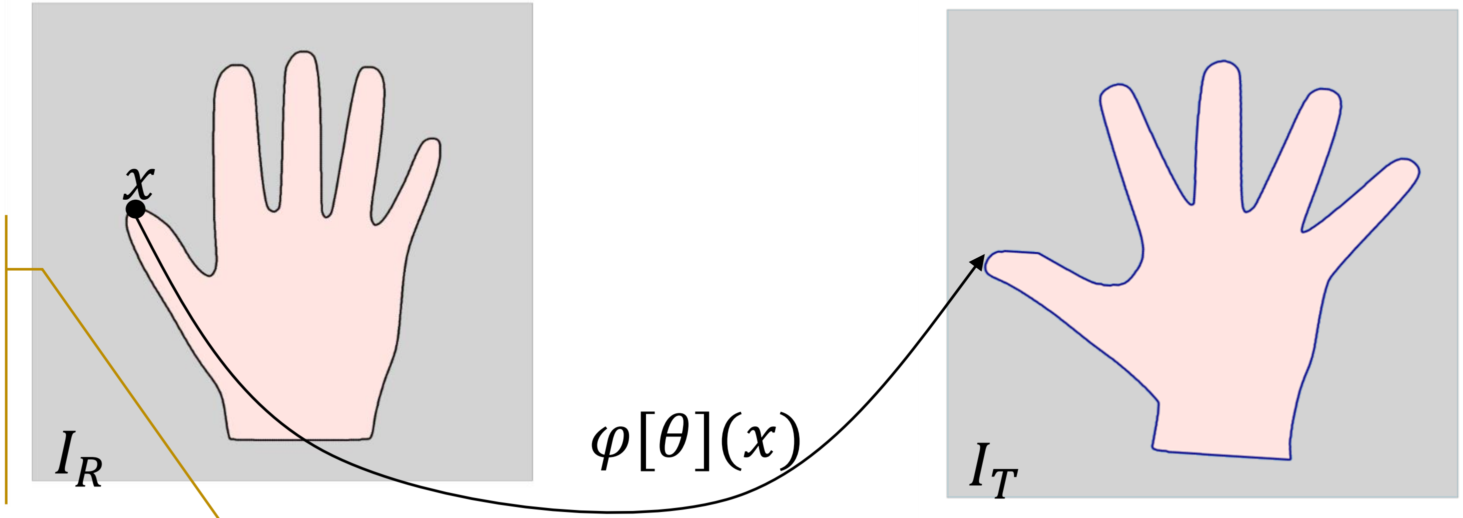

Image term outside mapping function. Makes problem really difficult

$$p(I_T|I_R,\theta) = \prod_{x\in\Omega} p(I_T(\varphi[\theta](x))|I_R,\theta,x) = \prod_{x\in\Omega} \frac{1}{Z}\exp\left(-\frac{\left(I_T\big(\varphi(x)\big) - I_R(x)\right)^2}{\sigma^2}\right)$$

Registration as analysis by synthesis

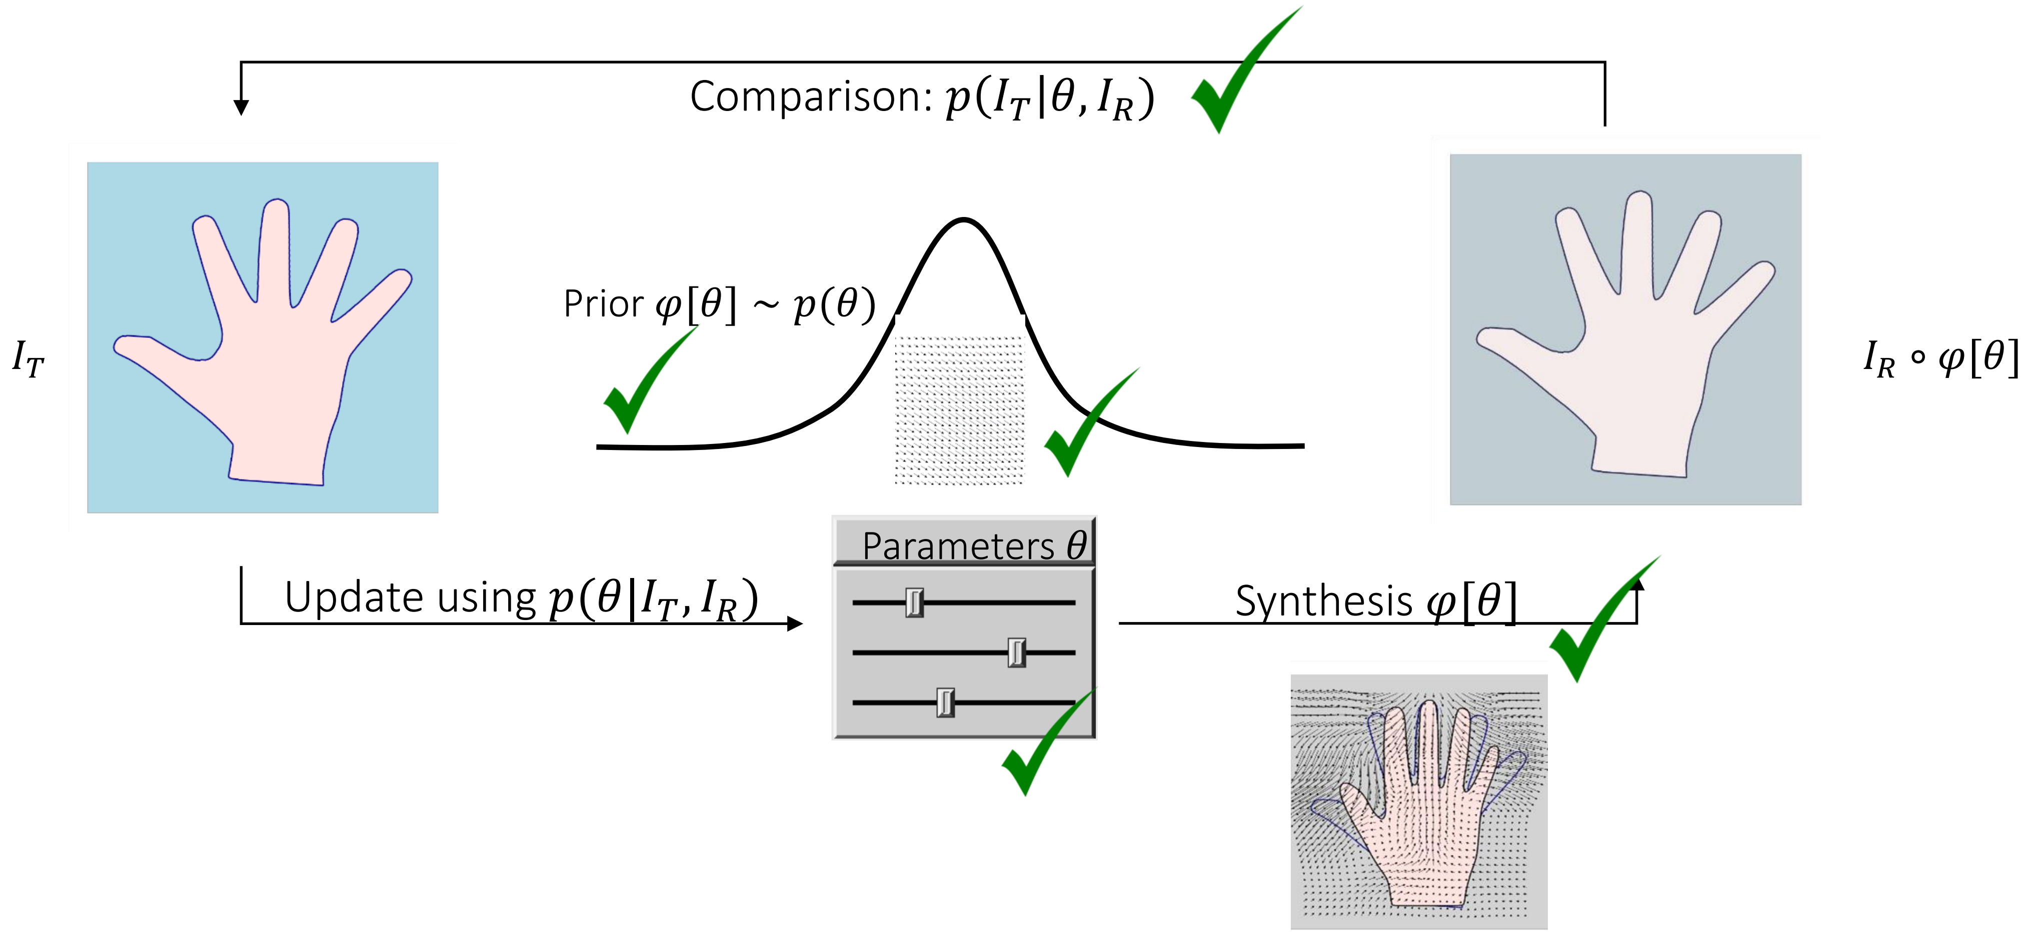

Registration problem

$$\theta^* = \arg\max_\theta p(\varphi[\theta])p(I_T|\varphi[\theta], I_R)$$

$$= \arg\max_\theta \frac{1}{Z_1}\exp\left(-\frac{1}{2}\|\theta\|^2\right)\frac{1}{Z_2}\prod_x \exp\left(-\frac{\left(I_T(\varphi[\theta](x)) - I_R(x)\right)^2}{\sigma^2}\right)$$

- Parametric problem, since:

$$\varphi[\theta](x) = x + \mu(x) + \sum_{i=1}^{r} \theta_i \sqrt{\lambda_i}\, \phi_i(x)$$

- Can be optimized using gradient descent

Variational formulation

$$\arg\max_{\theta} \frac{1}{Z_1} \exp\left(-\frac{1}{2}\|\theta\|^2\right) \frac{1}{Z_2} \prod_x \exp\left(-\frac{\left(I_T\left(\varphi[\theta](x)\right) - I_R(x)\right)^2}{\sigma^2}\right)$$

$$= \arg\max_{\theta} \ \ln \frac{1}{Z_1} \exp\left(-\frac{1}{2}\|\theta\|^2\right) + \ln \frac{1}{Z_2} \prod_x \exp\left(-\frac{\left(I_T\left(\varphi[\theta](x)\right) - I_R(x)\right)^2}{\sigma^2}\right)$$

$$= \arg\max_{\theta} \ \ln \frac{1}{Z_1} - \frac{1}{2}\|\theta\|^2 + \ln \frac{1}{Z_2} - \sum_{x\in\Omega} \frac{\left(I_T\left(\varphi[\theta](x)\right) - I_R(x)\right)^2}{\sigma^2}$$

$$= \arg\min_{\theta} \sum_{x\in\Omega} \frac{\left(I_T\left(\varphi[\theta](x)\right) - I_R(x)\right)^2}{\sigma^2} + \frac{1}{2}\|\theta\|^2$$

Image metric

Regularizer

The registration problem

Probabilistic formulation

$$\theta^* = \arg\min_{\theta} -\ln\left(p(I_T | I_R, \varphi[\theta])\right) - \ln p(\varphi[\theta])$$

Variational formulation

$$\theta^* = \arg\min_{\theta} D[I_T, I_R, \varphi[\theta]] + \lambda R[\varphi[\theta]]$$

Related to noise assumption in metric

GP-Registration in Scalismo

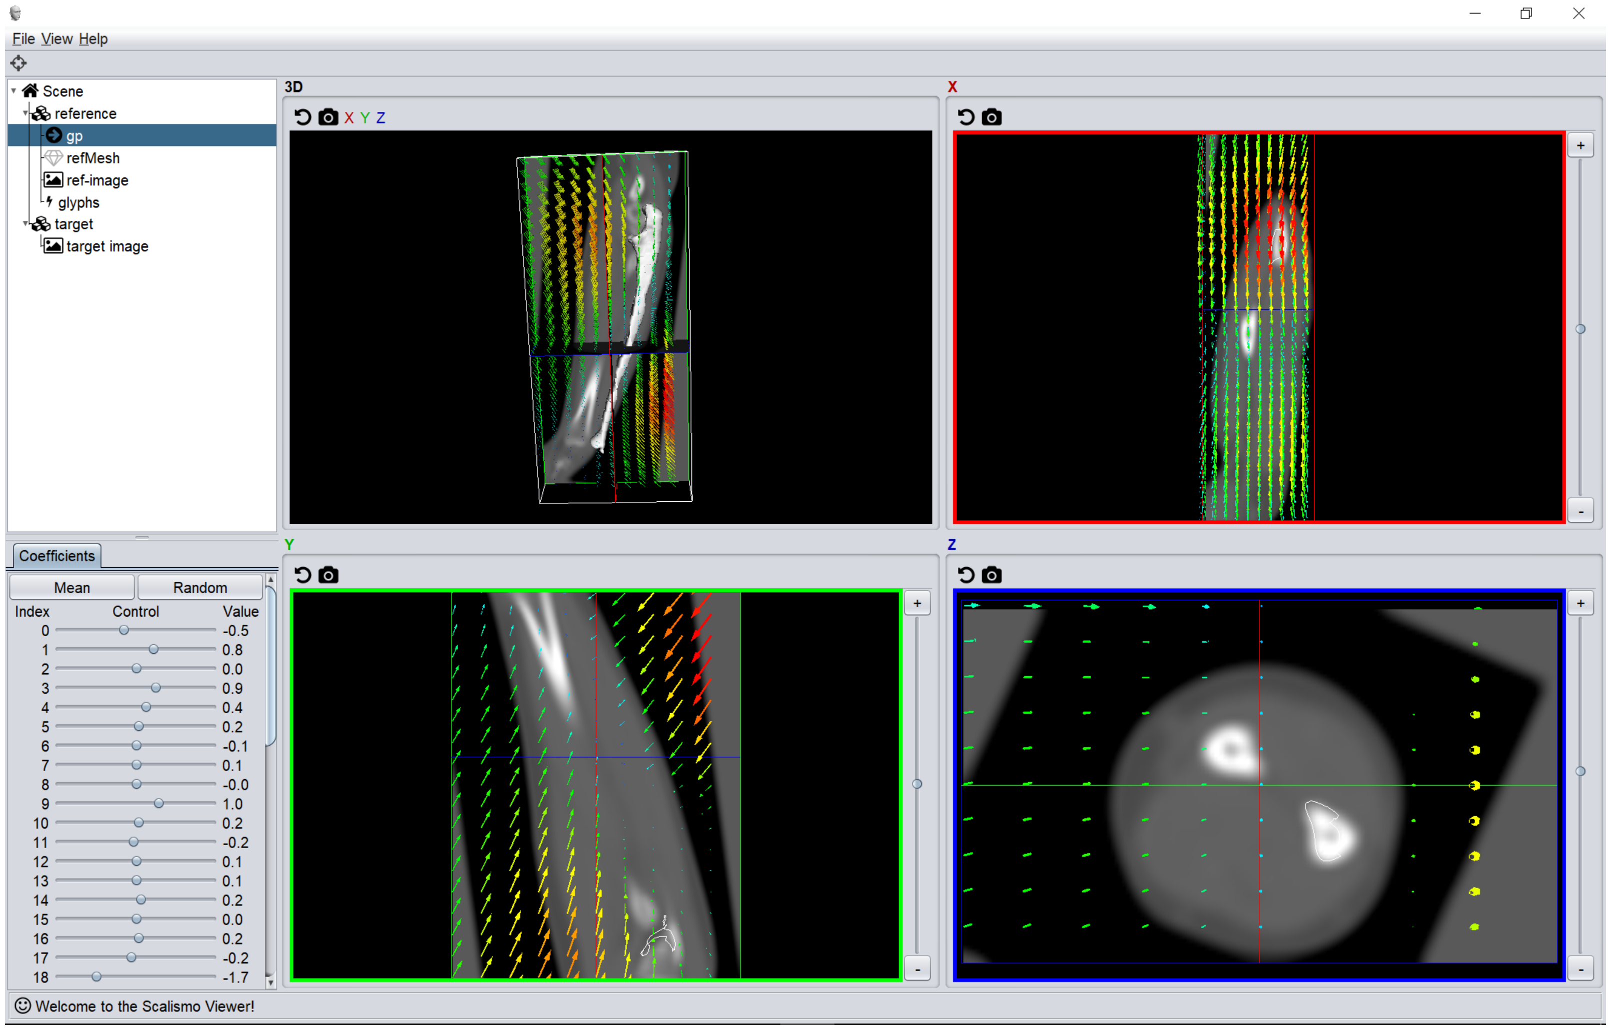

A selection of useful likelihood functions

Landmark likelihood

For one landmark pair (l_R, l_T):

$$p(l_T|\theta, l_R) = N(\varphi[\theta](l_R), I_{2x2}\sigma^2)$$

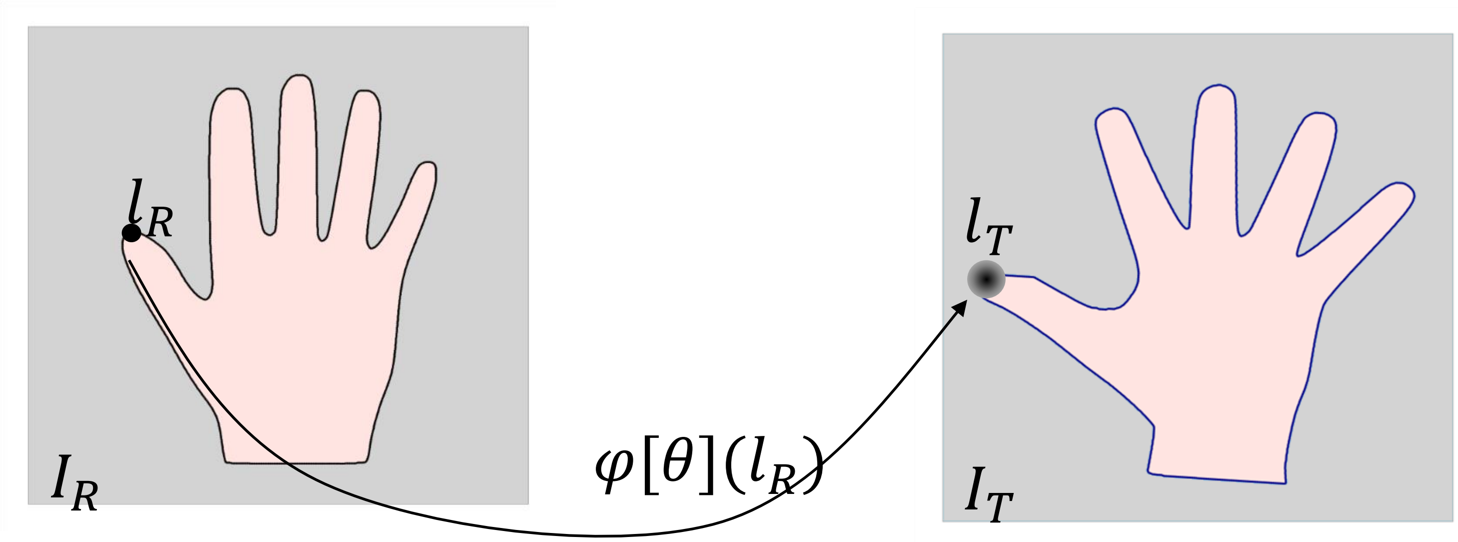

Landmark likelihood

For many landmarks: $L = ((l_R^1, l_T^1), \ldots, (l_R^n, l_T^n))$

$$p(l_1^T, \ldots, l_n^T \mid \theta, l_R^1, \ldots, l_R^n) = \prod_i N(\varphi[\theta](l_R), I_{2x2}\sigma^2)$$

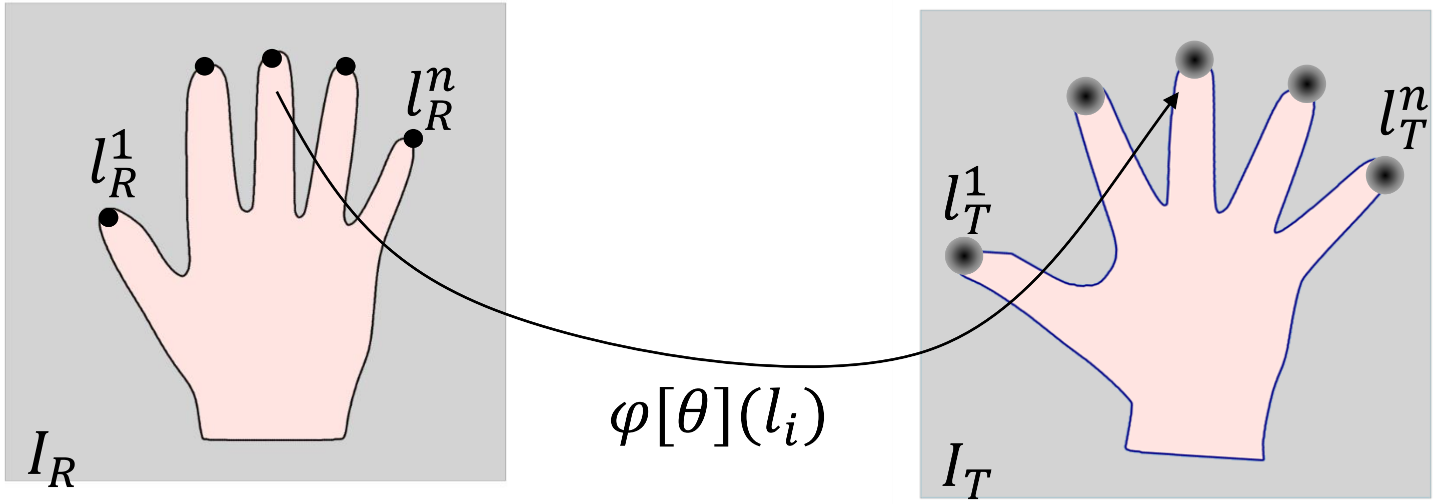

Likelihood function: Image registration

Images are similar when the intensities match

Assumptions:

- Corresponding points have the same image intensity (up to i.i.d. noise)

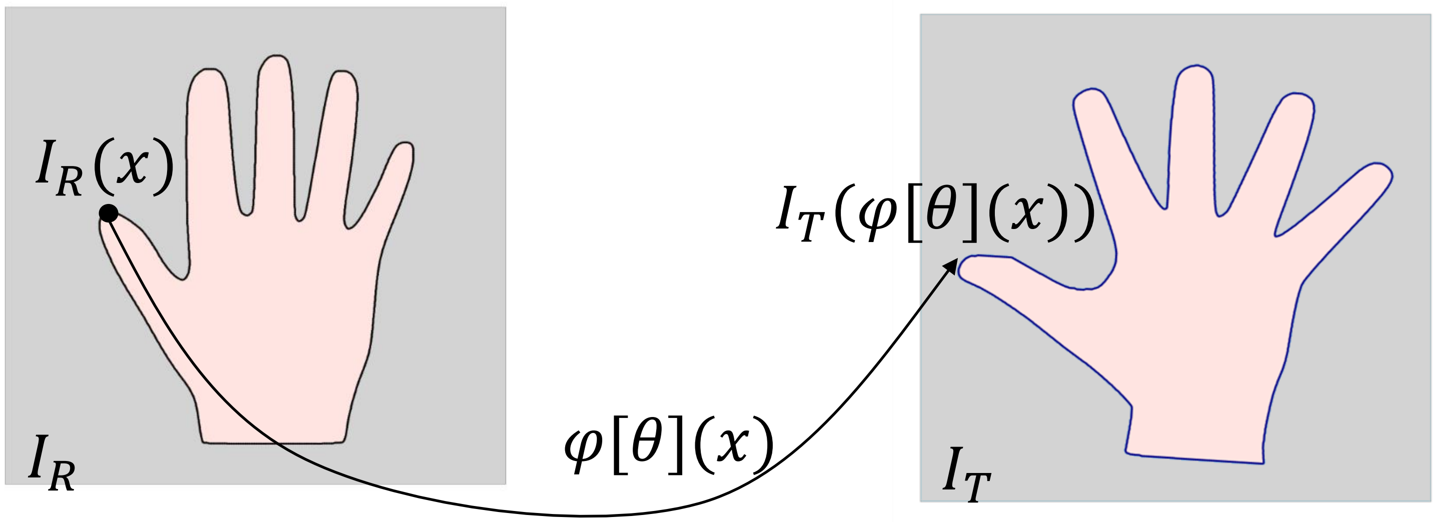

$$p(I_T(\varphi[\theta](x))|I_R,\theta,x) \sim N(I_R(x),\sigma^2)$$

Likelihood function: Image registration

Images are similar when the intensities match

Assumptions:

- Corresponding points have the same image intensity (up to i.i.d. noise)

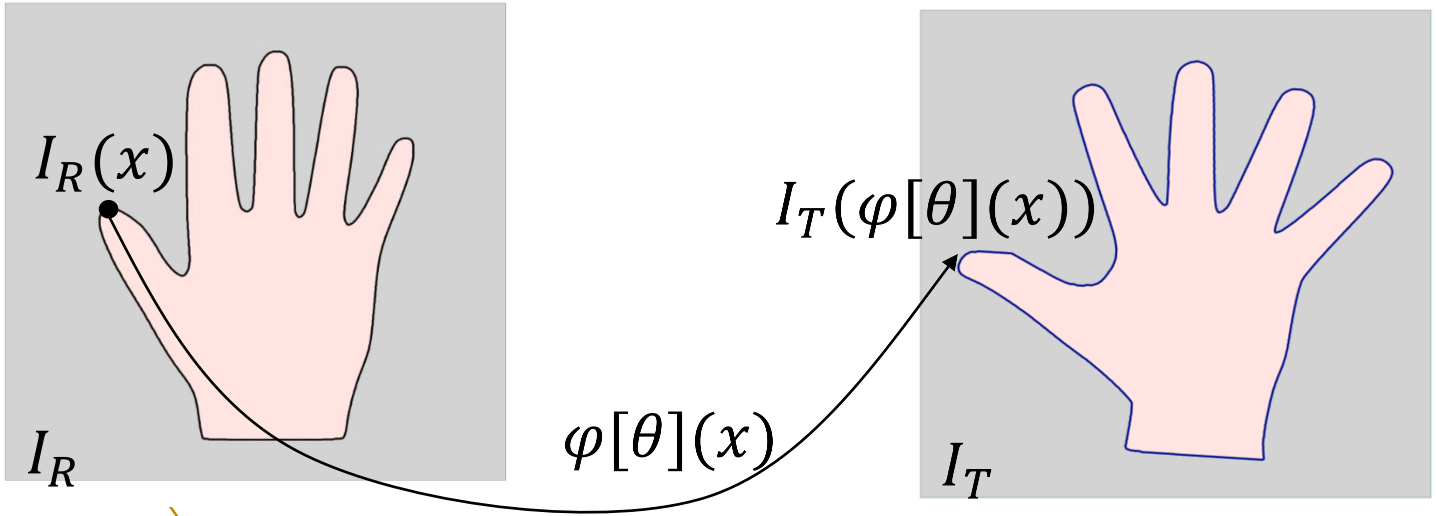

Image term outside mapping function. Makes problem really difficult

$$p(I_T|I_R,\theta) = \prod_{x\in\Omega} p(I_T(\varphi[\theta](x))|I_R,\theta,x)$$

Image vs. Landmark registration

- Landmark registration is easy
  - All components are Gaussian
  - Closed form solution using Gaussian process regression

- Image registration is hard
  - Image destroys Gaussian assumption
  - Likelihood function is not Gaussian
  - Problem with many local minima

What about surface registration?

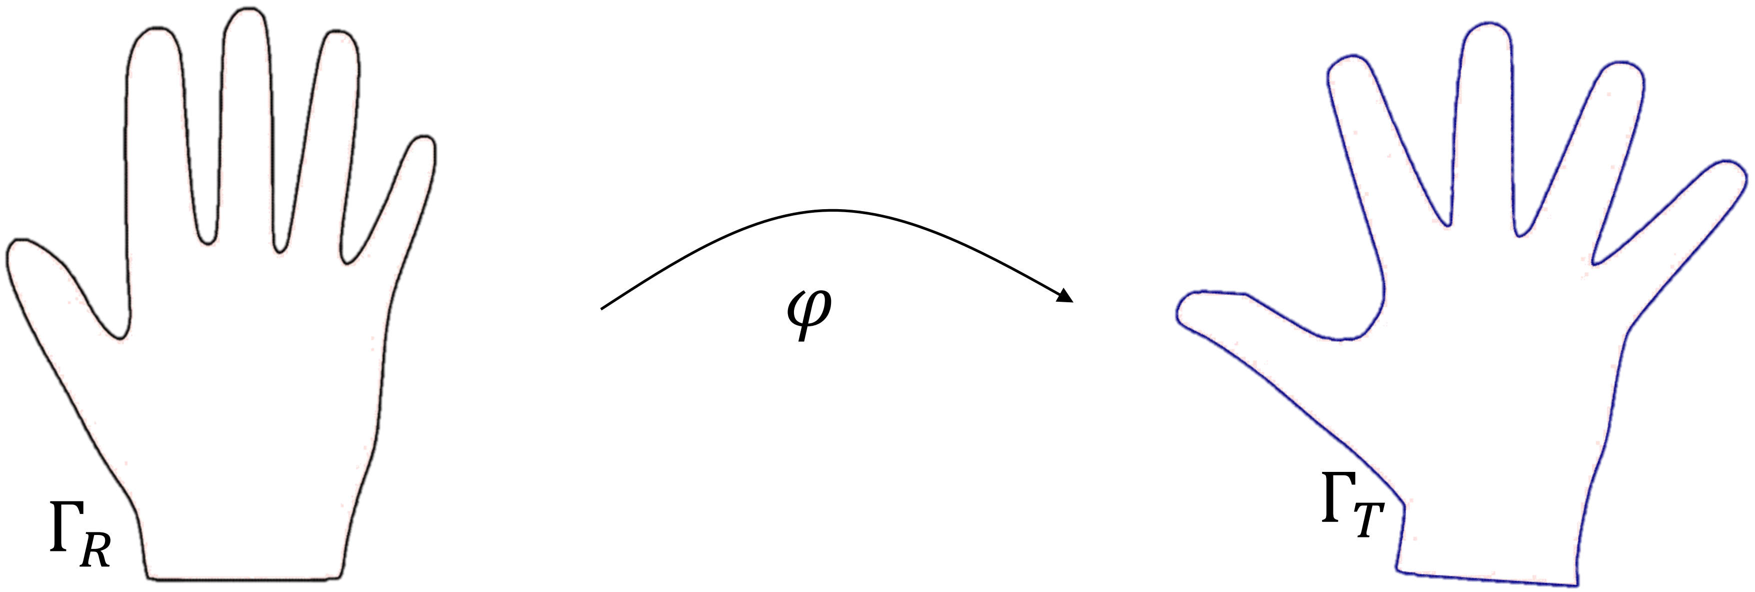

Reference (surface): Γ_R

Target (surface): Γ_T

A trick: Implicit definition of a surface

- Surface Γ can be represented as the zero level set of a distance function defined as

$$D_\Gamma(x) = \|\mathrm{ClosestPoint}_\Gamma(x) - x\|$$

with

$$\mathrm{ClosestPoint}_\Gamma(x) = \arg\min_{x' \in \Gamma} \|x - x'\|$$

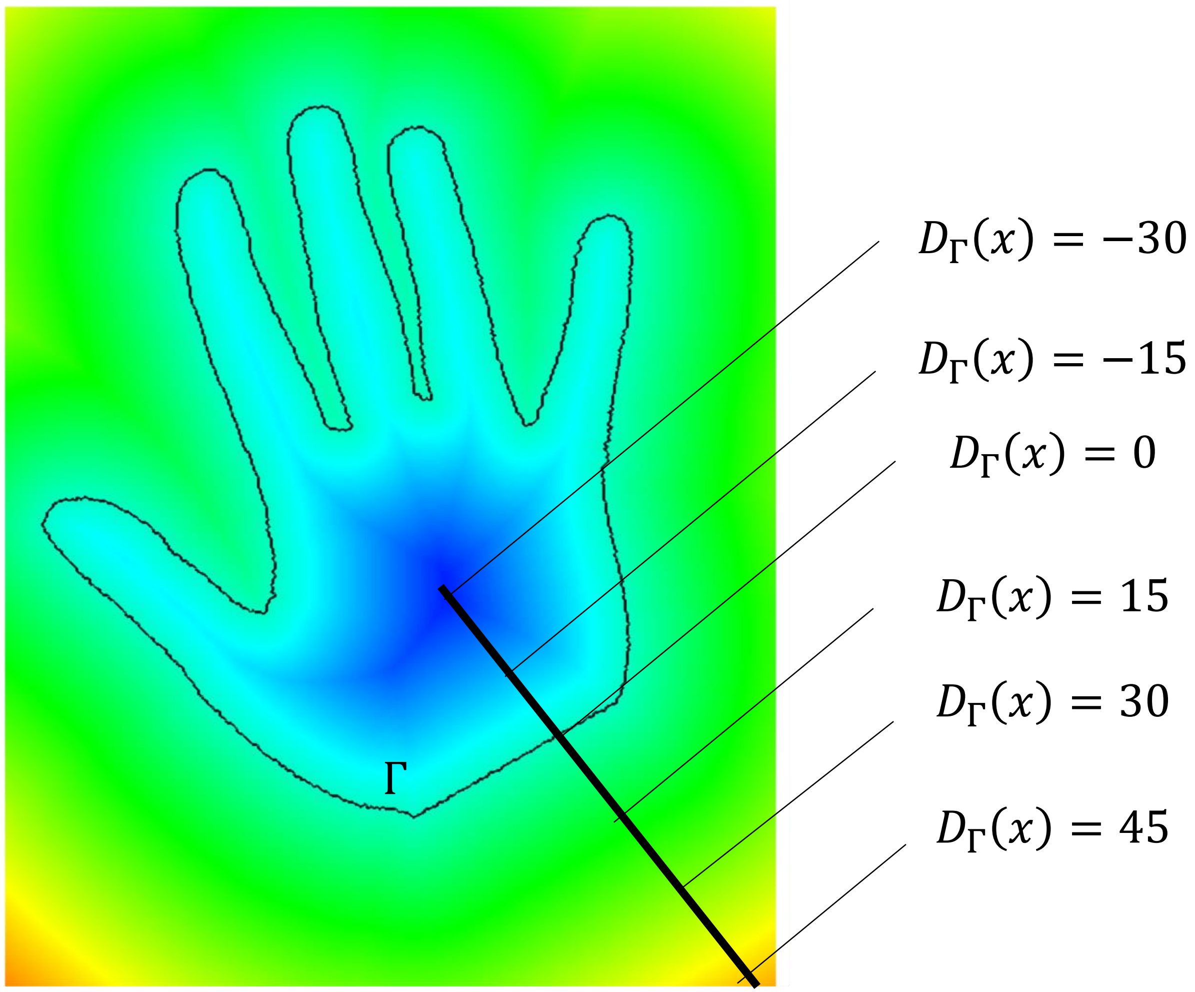

Likelihood function: Surface registration

Surface registration becomes image registration of distance images:

$$p(D_T(\varphi[\theta](x))|\theta, D_R, x) \sim N(D_R(x), \sigma^2)$$

- Most likely solution: Points with same distances are mapped to each other
- σ^2 has now geometric interpretation

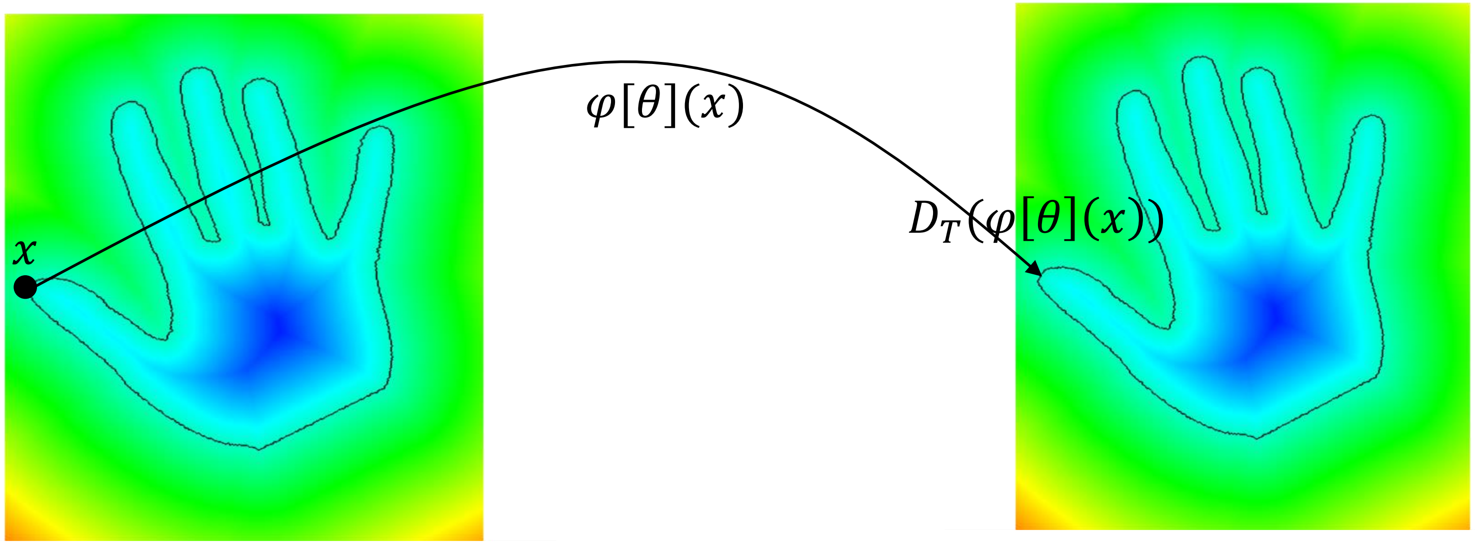

Reference $D_R: \Omega_R \to \mathbb{R}$ Target $D_T : \Omega_T \to \mathbb{R}$

Likelihood function: Active shape models

- ASM models profile $\rho(x_i)$ as a normal distribution $p(\rho(x_i)) = N(\mu_i, \Sigma_i)$

- Single profile point x_i:

$$p(\rho(\varphi[\theta](x_i))|\theta, x_i) = N(\mu_i, \Sigma_i)$$

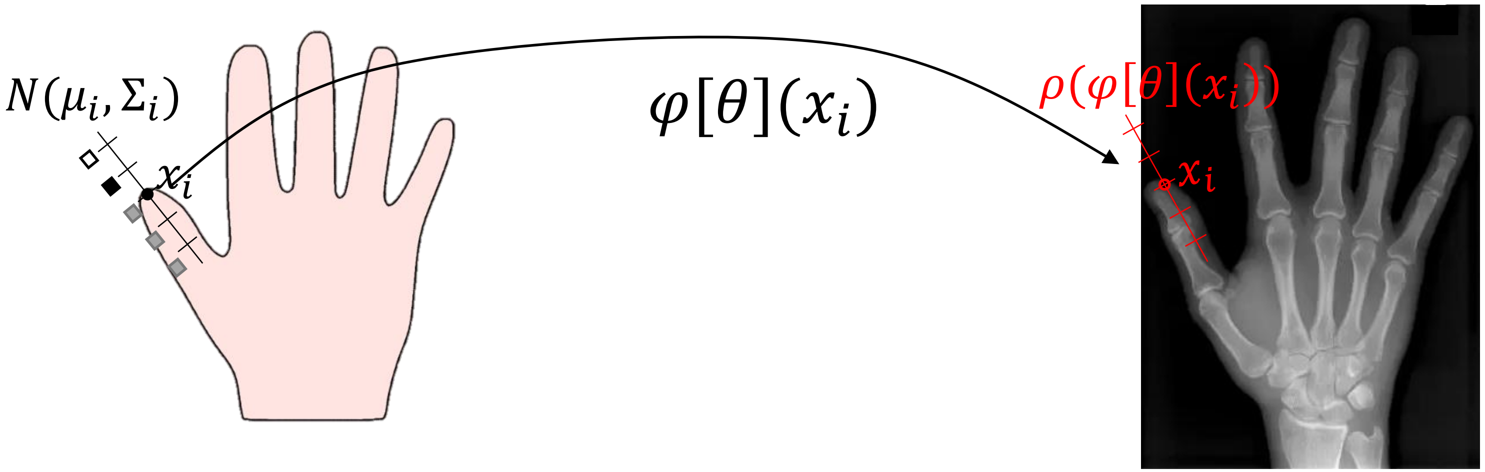

Shape is well matched if environment around profile points is likely under trained model.

Likelihood function: Active shape models

- ASM models profile $\rho(x_i)$ as a normal distribution $\quad p(\rho(x_i)) = N(\mu_i, \Sigma_i)$

- Multiple points

$$p(\rho(\varphi[\theta](x_1)), \ldots, \rho(\varphi[\theta](x_n)) | \theta) = \prod_i N(\mu_i, \Sigma_i)$$

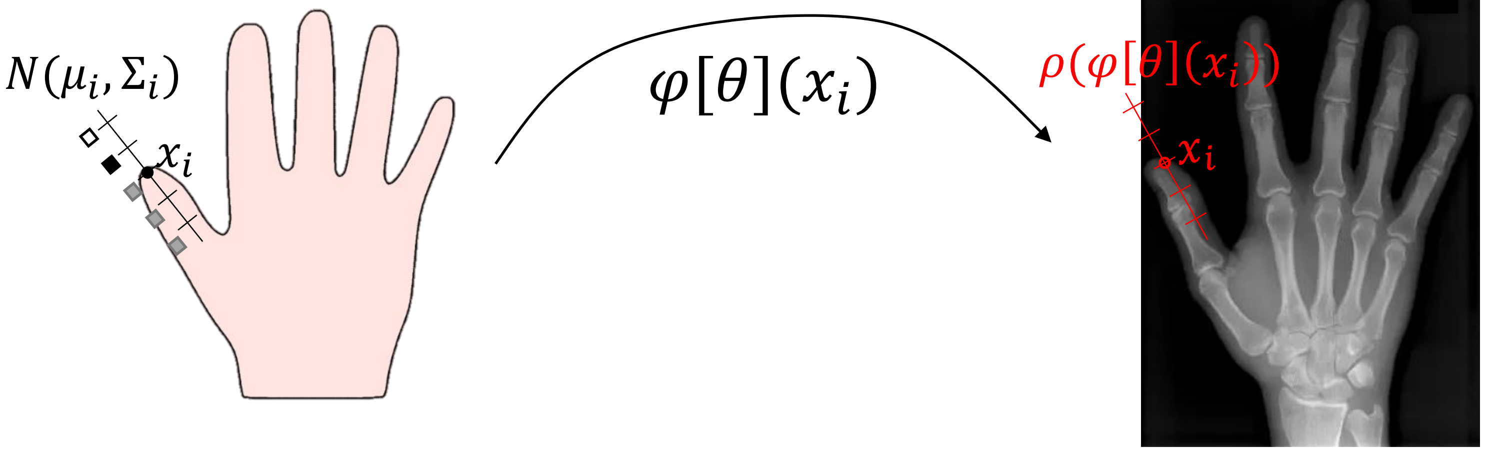

Conclusion

- Registration is analysis by synthesis problem
  - Synthesis: Warp reference
  - Likelihood: Compare properties defined on (warped) reference with those found in target
- Typical likelihood functions model:
  - Point positions (landmark registration)
  - Intensities distributions (image registration and Active Shape models)
  - Distance (surface registration)

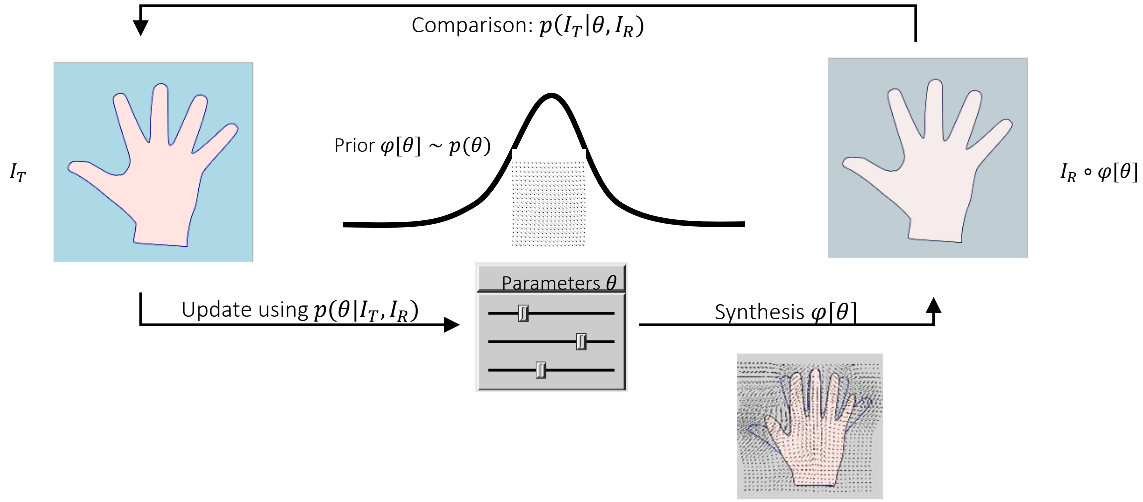

- Usually only MAP solution is sought.
  - Justified if uncertainty in correspondence can be ignored