Image Displacement and Deformation

Images can provide two types of information:

- **What?** Specific spatial patterns of image intensity. Eg. An image of Einstein.
- **Where?** Spatial (temporal) deformations of a known image, or deformations between several images. Eg. The deformation to an oblique image of a photograph of Einstein.

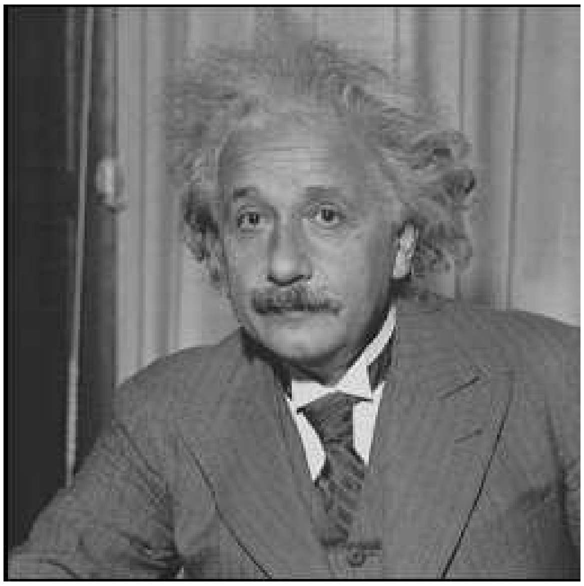

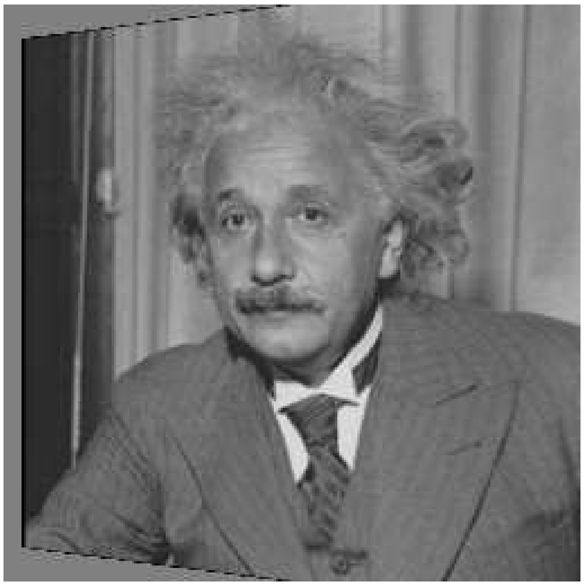

Rubber Sheet Deformation

Consider the image pair:

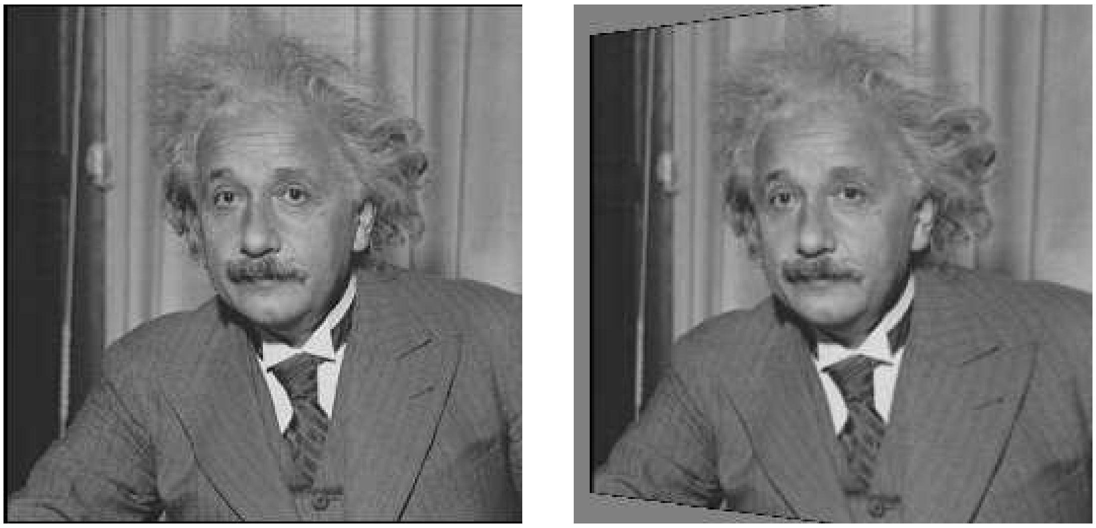

The deformation is a 2D displacement field, $\vec{v}(\vec{x})$:

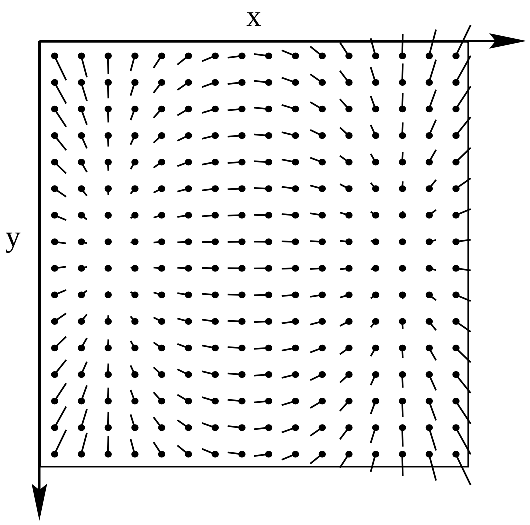

Here:

- $\vec{W}(\vec{x}) = \vec{x} + \vec{v}(\vec{x})$ is the image warp.
- $I_r(\vec{W}(\vec{x})) = I_l(\vec{x})$ describes the “rubber sheet” image deformation.

Representation of Smooth Deformation Fields

Pure Translation:

$$\vec{W}(\vec{x};\vec{a}) = \vec{x} + \vec{a}, \quad \text{with } \vec{a} \in \Re^2.$$

Affine Warp:

$$\vec{W}(\vec{x},\vec{a}) = \vec{a}_0 + A\vec{x}, \quad \text{with } \vec{a} \in \Re^6.$$

Provides a good approximation for the motion of objects with

- Small depth variation.
- Small angular extents within the image.

2D Homography (Rational Warp):

$$\vec{W}(\vec{x};\vec{a}) = \{\vec{a}_0 + A\vec{x}\} \,/\, \left[1 + \vec{b}\cdot\vec{x}\right] \quad \text{with } \vec{a} \in \Re^8.$$

Provides a good approximation for

- Thin structures (nearly planar surfaces).
- No restriction on angular extent.

Can also consider patchwise representations of smooth flow, or splines, RBFs, and so on.

Estimating Image Displacement

Given two signals, $q_1(x)$ and $q_2(x)$, the displacement d satisfies

$$q_2(x) = q_1(x - d).$$

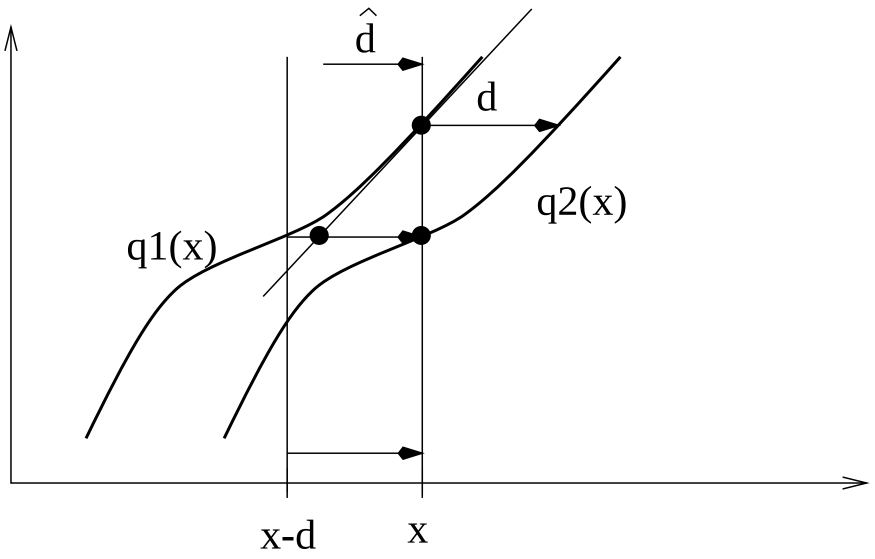

Given information about q_1 and q_2 at only one point x, we can use a **local linear model:**

$$q_1(x + r) = q_1(x) + r\frac{dq_1}{dx}(x) + O(r^2).$$

Neglecting $O(r^2)$ terms, d can then be approximated by

$$\hat{d} = [q_1(x) - q_2(x)] / \frac{dq_1}{dx}(x).$$

To leading order, the error is

$$|\hat{d} - d| \leq \frac{1}{2}\hat{d}^2 |\frac{d^2q_1}{dx^2}(x)| / |\frac{dq_1}{dx}(x)|.$$

Differential Formulation

For near-identity warps, that is $\vec{W}(\vec{x};\vec{a}) = \vec{x} + \vec{v}(\vec{x};\vec{a})$ with $\vec{v}(\vec{x};\vec{a}) \approx \vec{0}$, a computationally convenient objective function is

$$\tilde{E}(\vec{a}) = \sum_{\vec{x} \in \mathcal{R}} \rho(\nabla I_r(\vec{x}) \cdot \vec{v}(\vec{x};\vec{a}) + [I_r(\vec{x}) - I_l(\vec{x}]).$$

Here $\nabla I = (\frac{\partial I}{\partial x}, \frac{\partial I}{\partial y})$ is the spatial gradient of the image, and $\mathcal{R}$ is a selected image region.

Each individual differential constraint,

$$\nabla I_r(\vec{x}, t) \cdot \vec{v}(\vec{x};\vec{a}) + (I_r(\vec{x}) - I_l(\vec{x})) = 0$$

is a linearized brightness constancy constraint, analogous to the local linear model on the previous slide.

An iteratively reweighted least squares algorithm can be used to find local minima of $\tilde{E}(\vec{a})$. **Blackboard.**

Just as in the 1D case, this formulation does not take the second derivatives of image brightness into account, leading to errors in the estimated warp (of size $O(||\vec{v}||^2)$.

1D Iterative Displacement Matching

Given an estimate $\hat{d}$ for the displacement, consider estimating the update, u:

$$q_2(x) = q_1(x - \hat{d} - u).$$

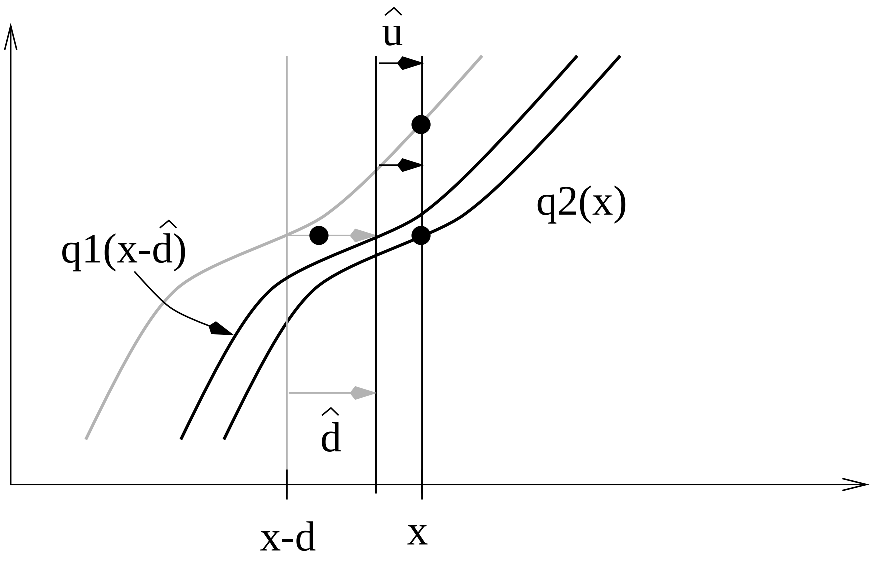

The new estimate for the displacement is $\hat{d} + \hat{u}$ where

$$\hat{u} = [q_1(x - \hat{d}) - q_2(x)] / \frac{dq_1}{dx}(x - \hat{d}).$$

The error in this new estimate is $O(u^2)$.

Further iteration gives Newton's method applied to

$$g_1(x - d) - g_2(x) = 0.$$

The approximation error $O(u^2)$ vanishes upon convergence (the displacement error will then be dominated by errors in $q_1(x)$ and $q_2(x)$).

Iterative Image Warping

For general warps $\vec{W}(\vec{x};\vec{a})$ between two image frames $I_l(\vec{x})$ and $I_r(\vec{x})$, we wish to minimize

$$E(\vec{a}) = \sum_{\vec{x}\in\mathcal{R}} \rho(I_r(\vec{W}(\vec{x};\vec{a})) - I_l(\vec{x})).$$

Given an current estimate $\vec{a}$ for the warp parameters, a convenient approximation for $E(\vec{a}+\vec{u})$ is to use the differential formulation $\tilde{E}(u)$ for the update.

It can be shown that $E(\vec{a}+\vec{u}) = \tilde{E}(\vec{u})+O(||\vec{u}||^2)$. In particular gradient directions wrt $\vec{u}$ agree at $\vec{u}=0$.

Iterative Rewarping: For accurate displacement estimates, the current warp parameters, say $\vec{a}_\mu$, are used to warp I_r,

$$I'_r(\vec{x}) = I_r(\vec{W}(\vec{x};\vec{a}_\mu)).$$

This warping compensates for changes of scale, orientation, and so on. Differential matching is performed on the warped pair $I'_r(\vec{W}(\vec{x};\vec{u}))$ and $I_l(\vec{x})$ to compute the update $\vec{u}_\mu$. The warp parameters are then updated so that $\vec{W}(\vec{x},\vec{a}_{\mu+1}) = \vec{W}(\vec{W}(\vec{x},\vec{a}_\mu),\vec{u}_\mu)$. The process is repeated until the updates $\vec{u}_\mu$ are sufficiently small.