Video Stabilization

Graphics, Vision and Video - Interdisciplinary Topics in Visual Computing

Nadia Robertini

21 June 2012

Structure

- The Problem
- Solution techniques
- Conclusions

- Discussion

The Problem

Solution

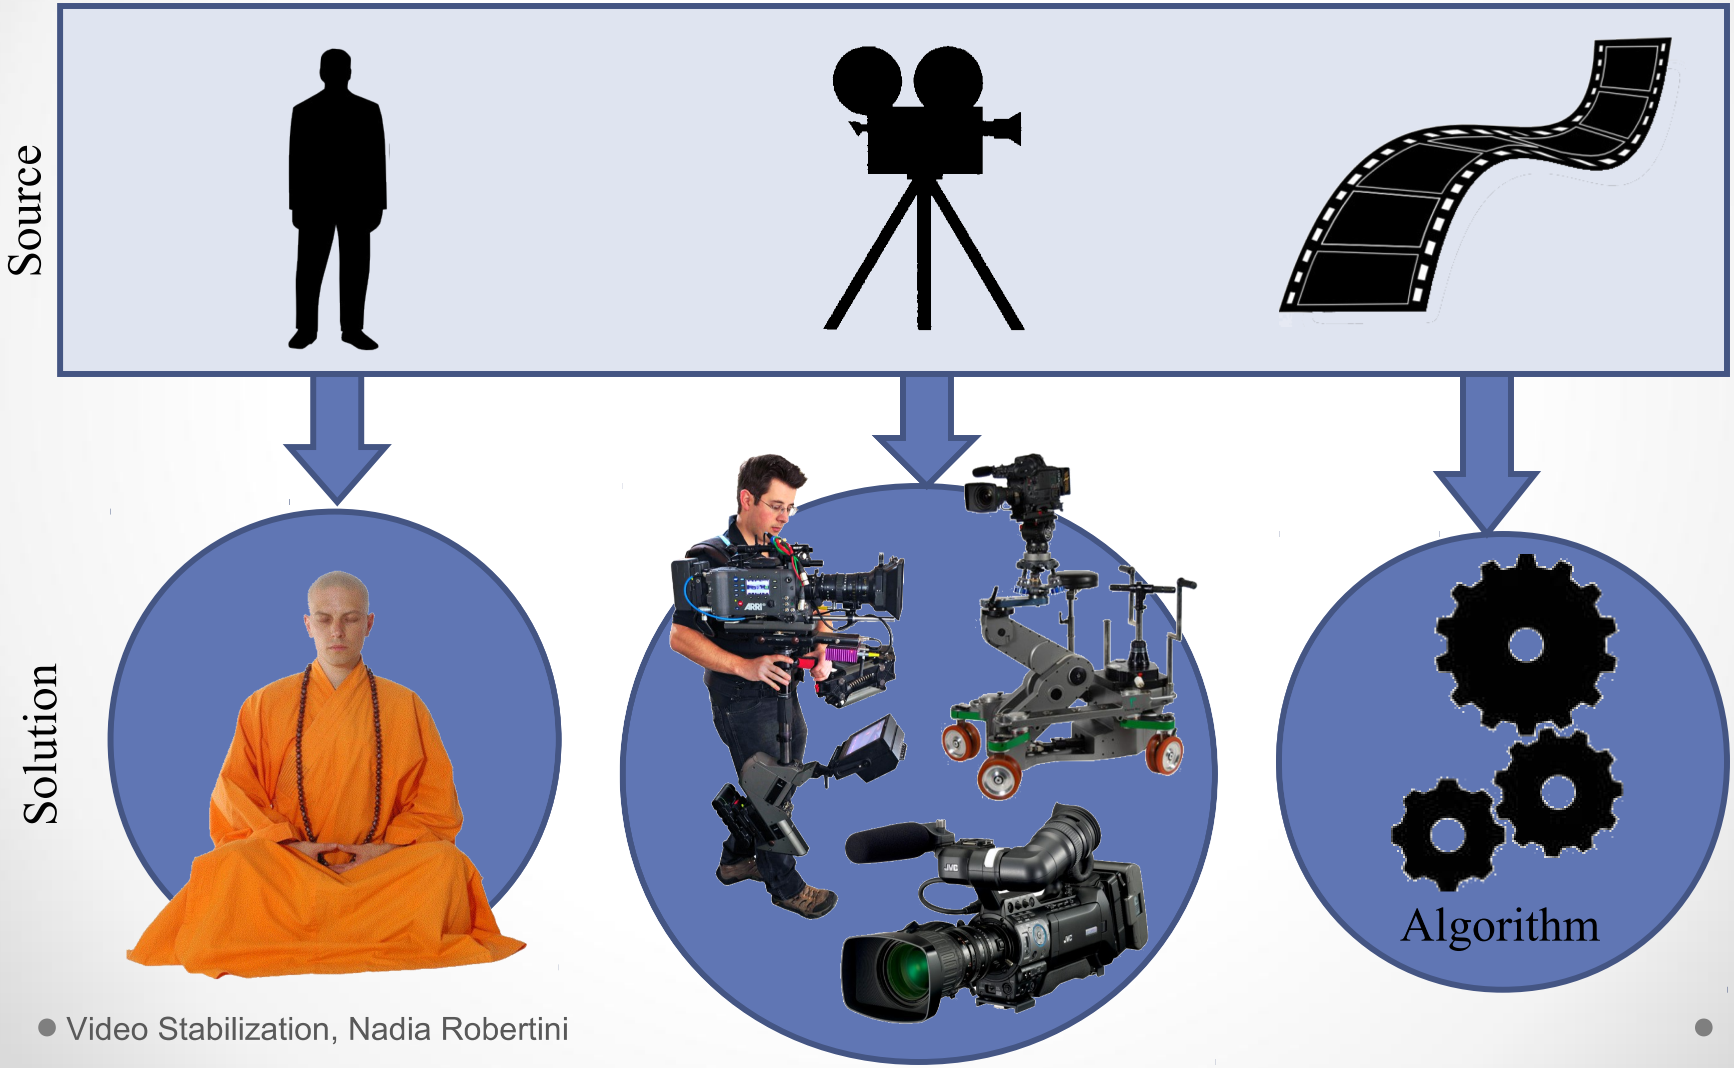

What do we want

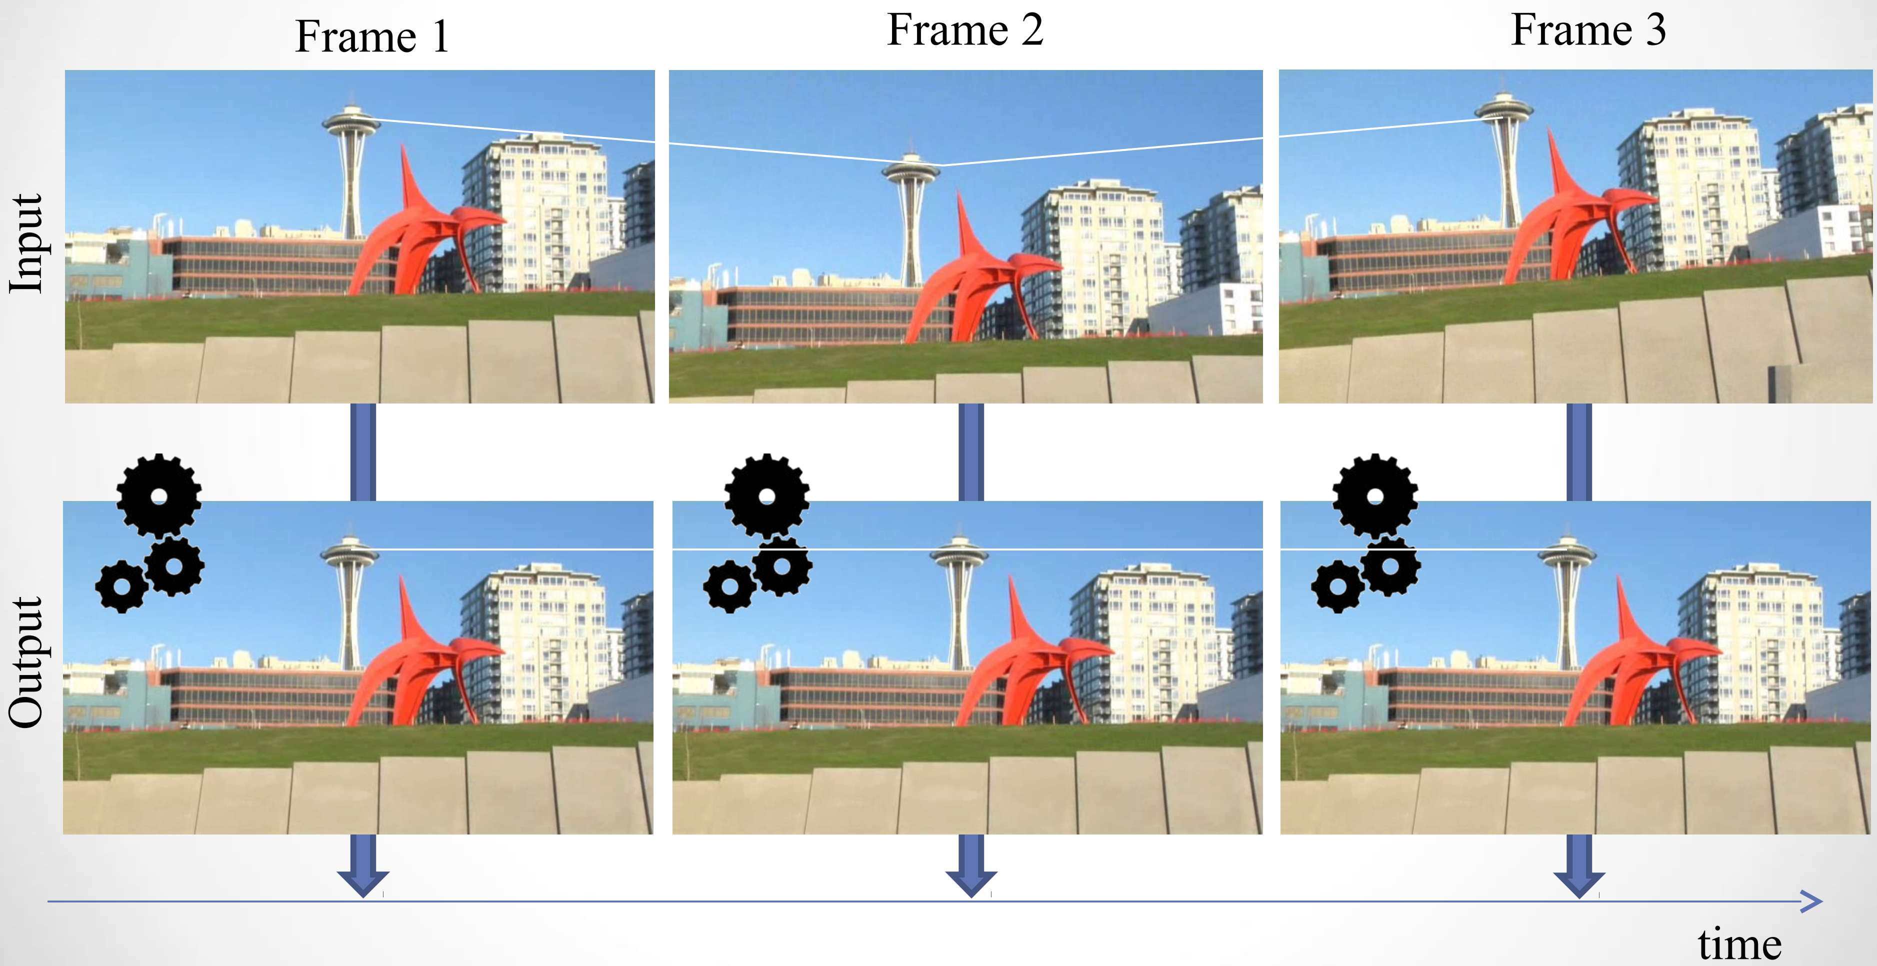

What is needed

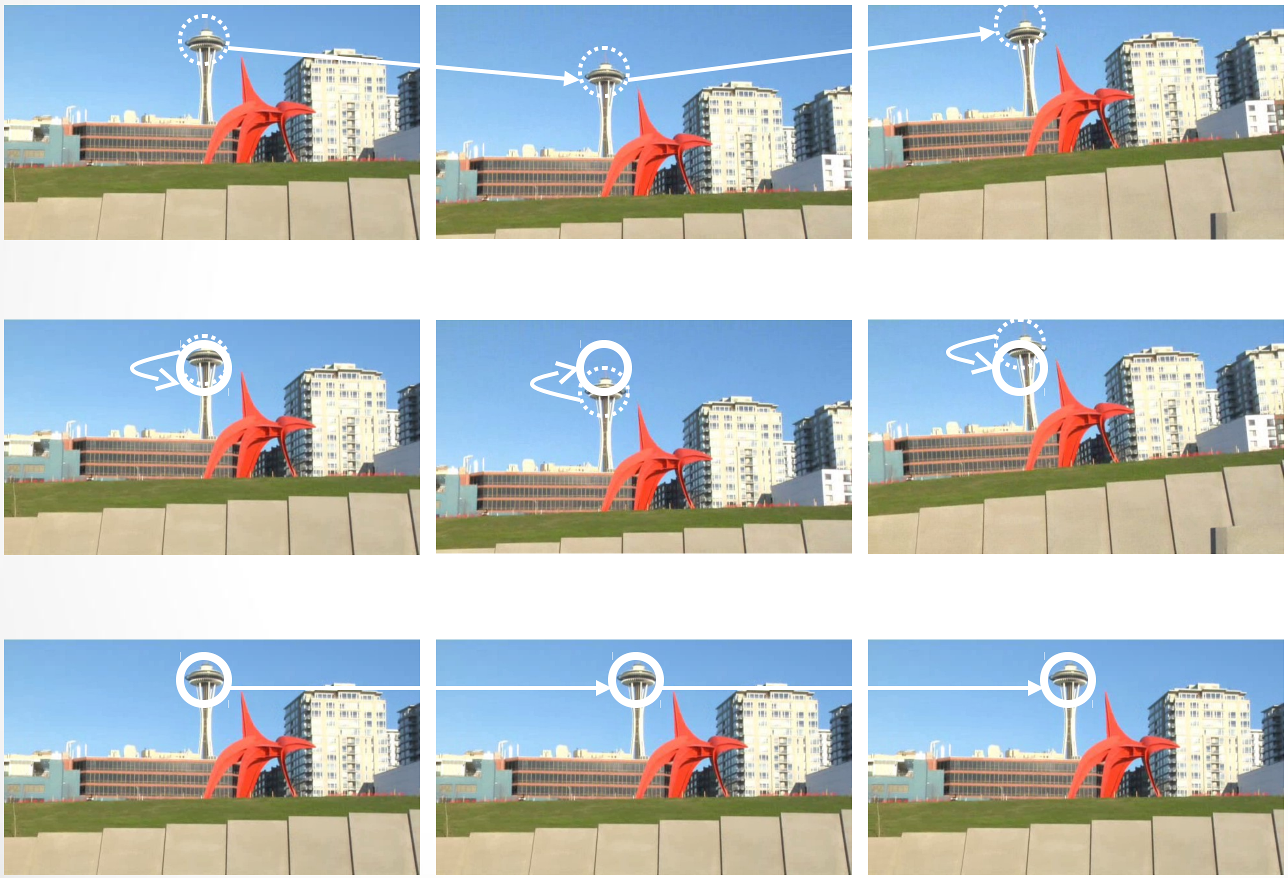

1) Motion Estimation

2) Motion Compensation

3) Image Composition

Motion Estimation

1) 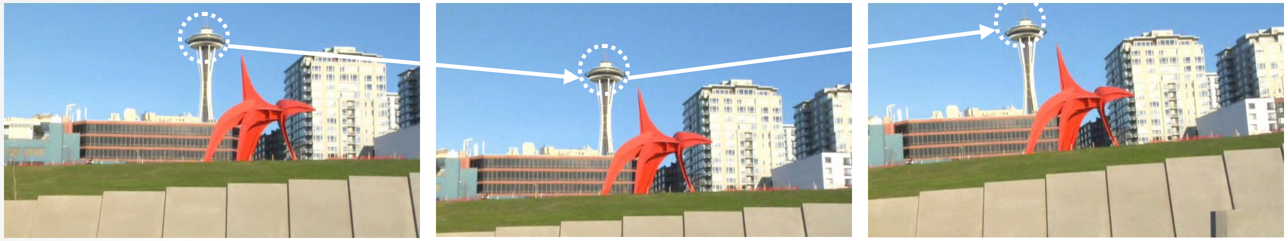

Motion Estimation

- Feature-based approach
  - KLT Tracker
  - SIFT matching
  - ...

Motion Compensation

2)

Motion Compensation

- 2D approach
- 3D approach
  - Structure-from-Motion (SFM)
- Subspace approach *[Liu et al. 2011]*

Image Composition

3) 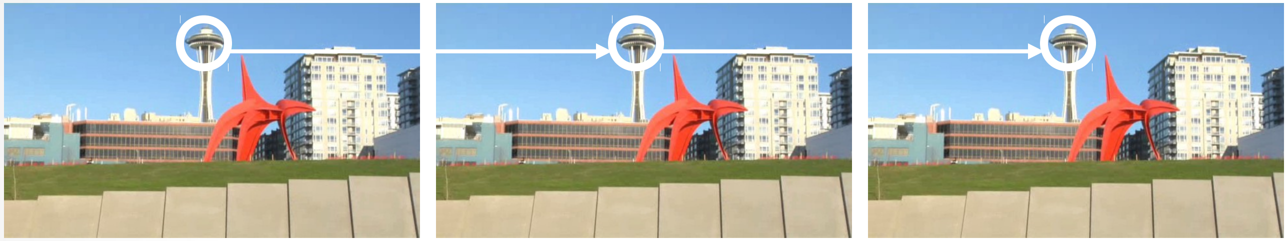

Image Composition

- Novel view interpolation (3D)
  - The Lumigraph
  - View-dependent texture mapping (VDTM)
  - …
- Full-frame Warp (2D)
- Content-preserving Warp *[Liu et al. 2009]*

The complete flow

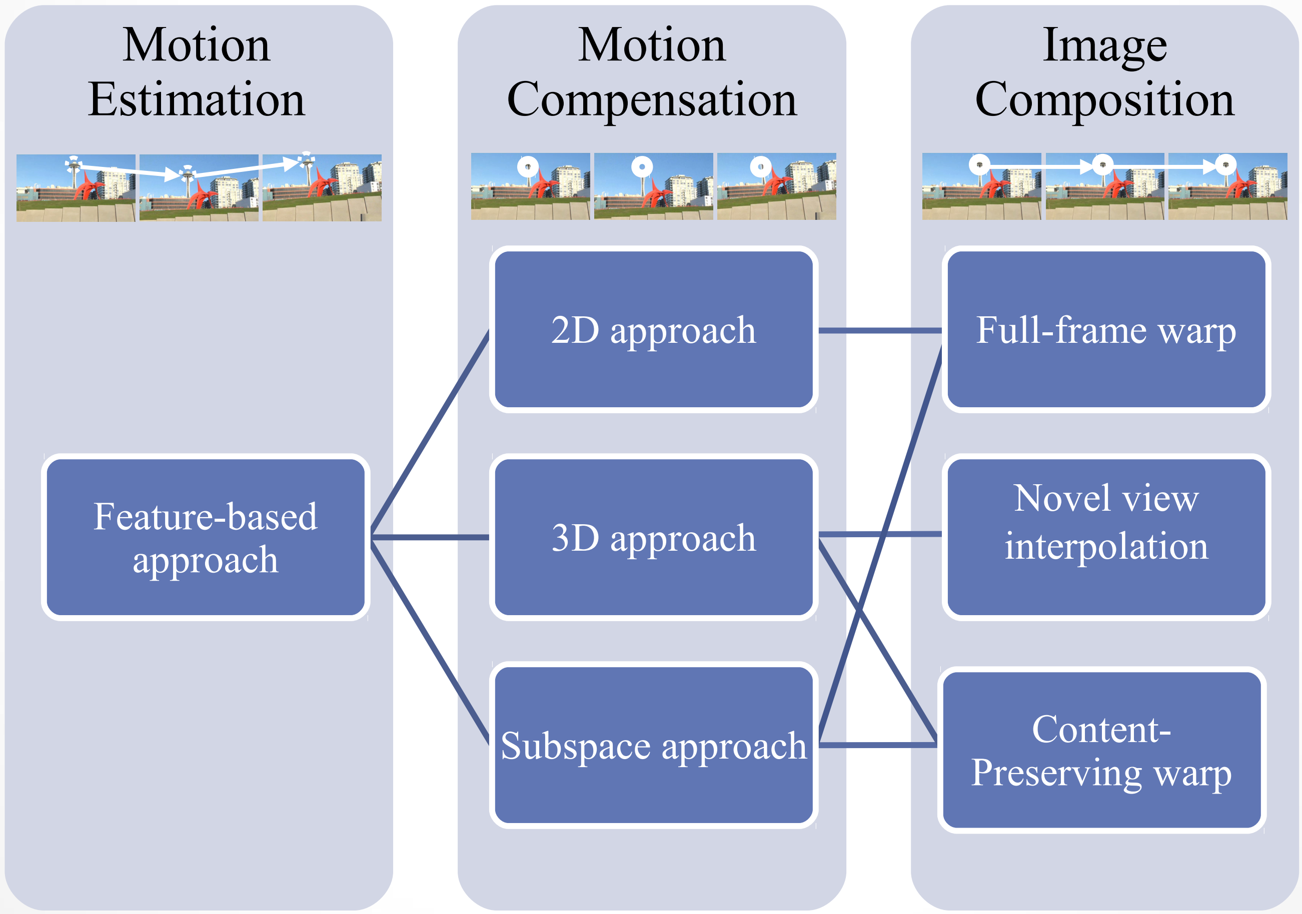

2D Stabilization

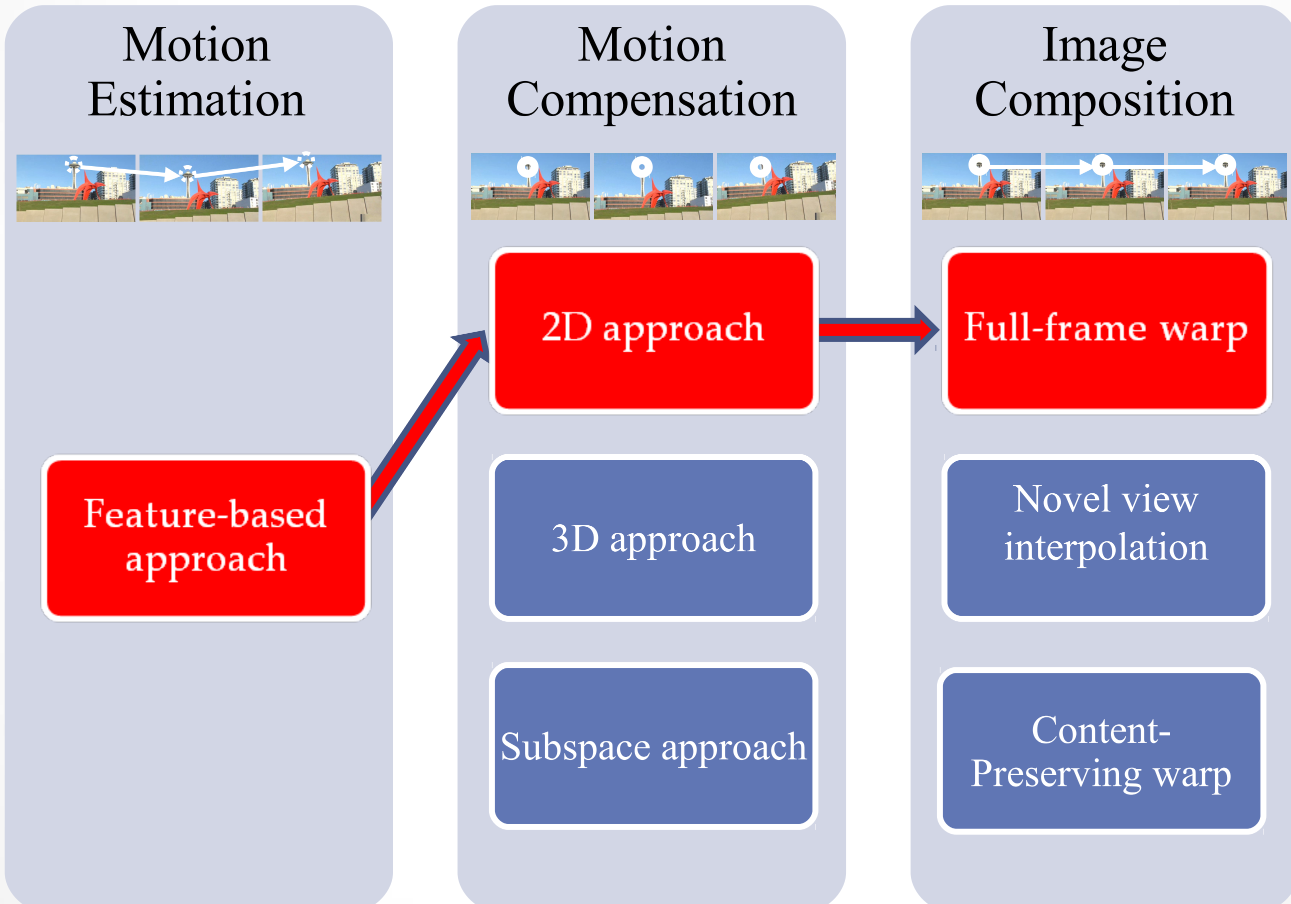

2D Stabilization

Motion Estimation

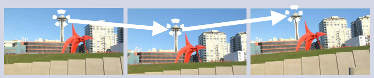

Feature-based approach

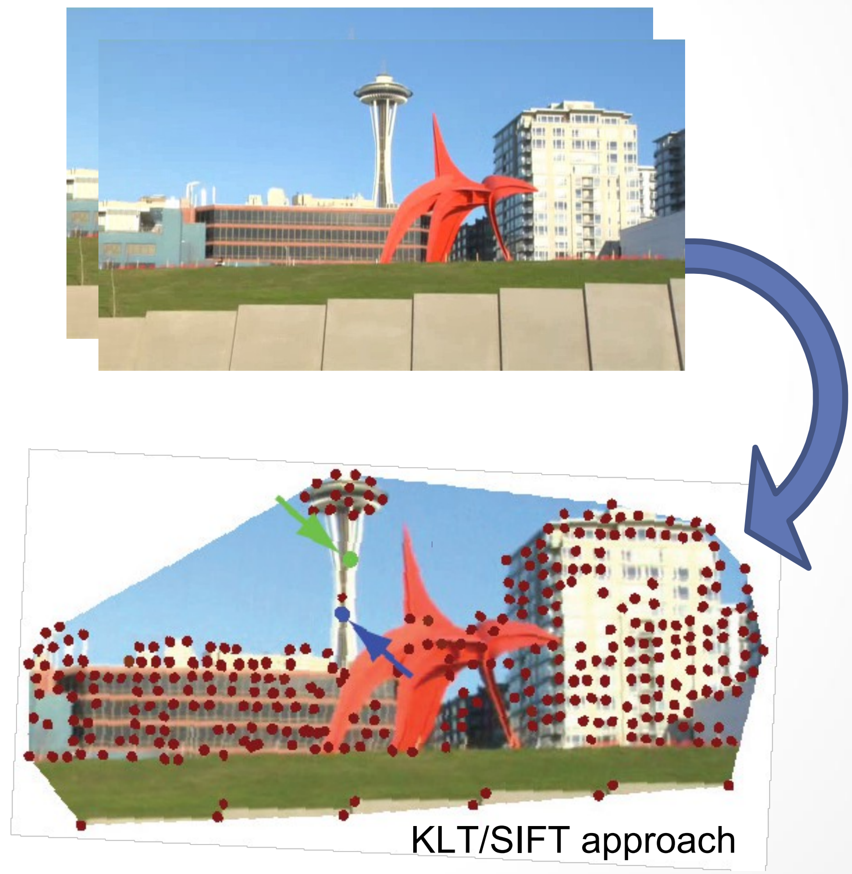

2D Stabilization

Motion Compensation

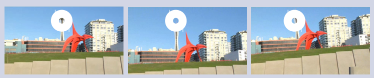

2D approach

3D approach

Subspace approach

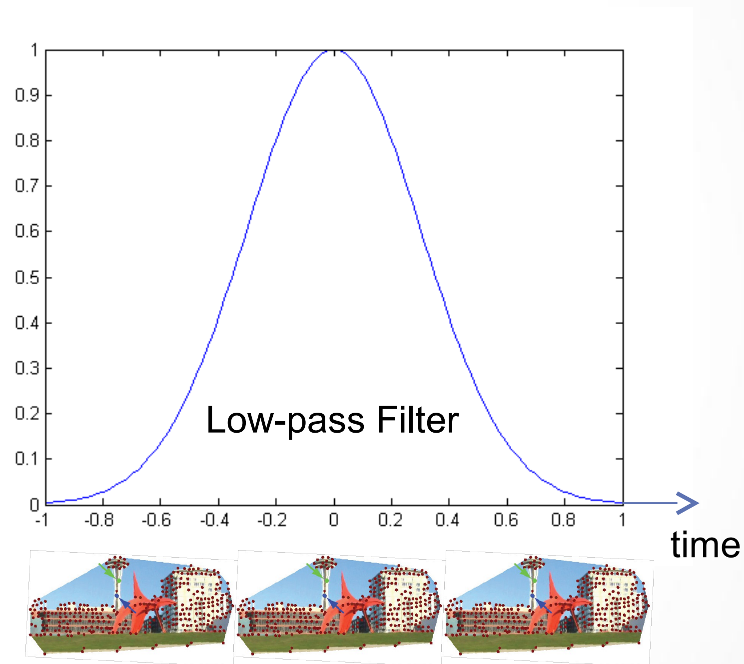

2D Stabilization

Image Composition

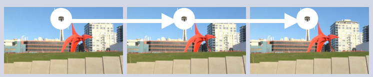

Full-frame warp

Novel view interpolation

Content-Preserving warp

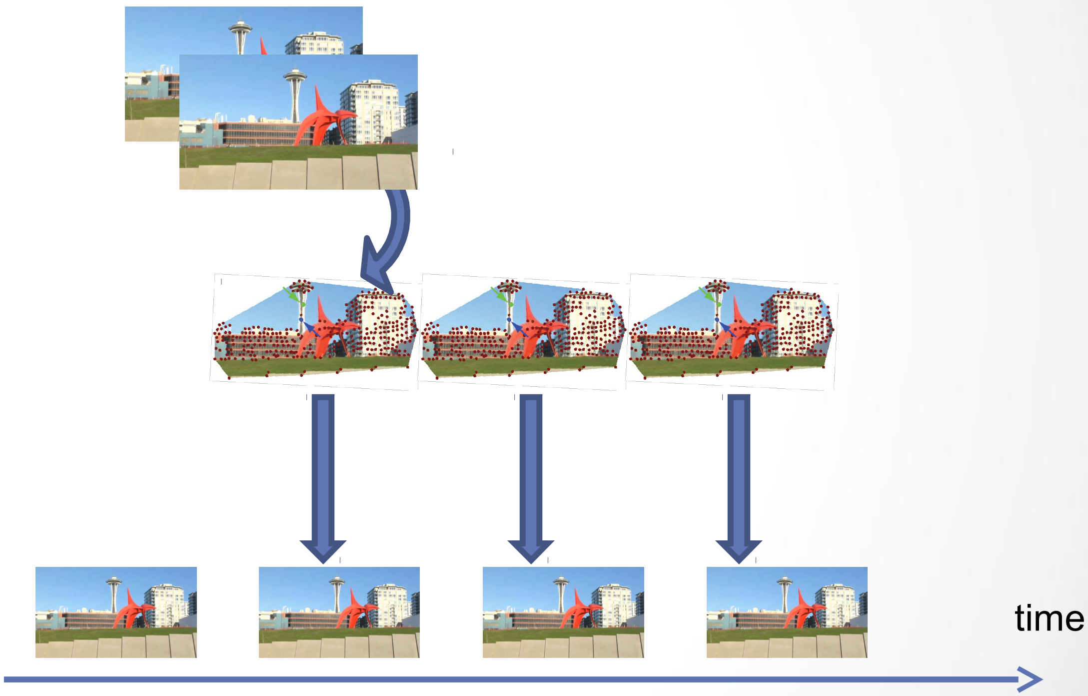

2D Stabilization

Image Composition

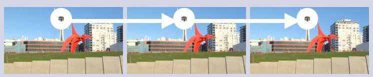

Full-frame warp

Novel view interpolation

Content-Preserving warp

Image Warping

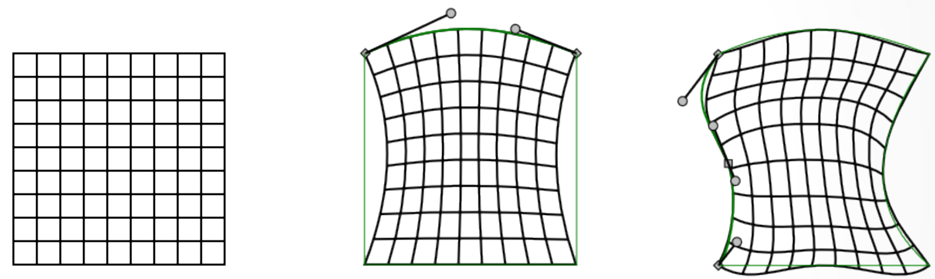

Original

Warp example

Warp example

2D Stabilization

Original shaky video

2D Stabilization

2D Stabilization

- V Significantly reduces camera-shake
- V Efficient
- X Not consistent with the geometry of the scene

→ 3D Stabilization

3D Stabilization

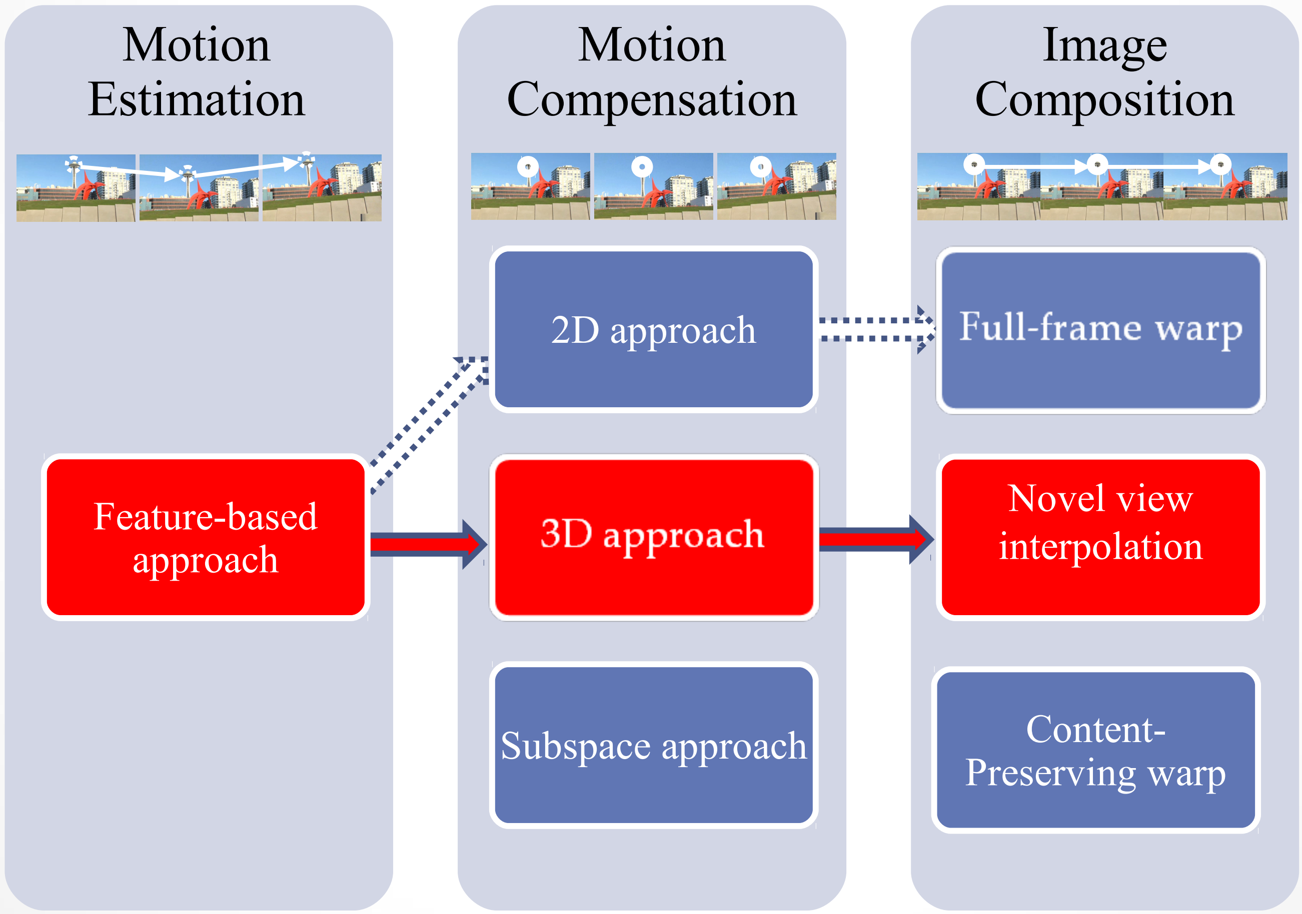

3D Stabilization

Motion Compensation

2D approach

3D approach

Subspace approach

- Structure-from-Motion (SFM)

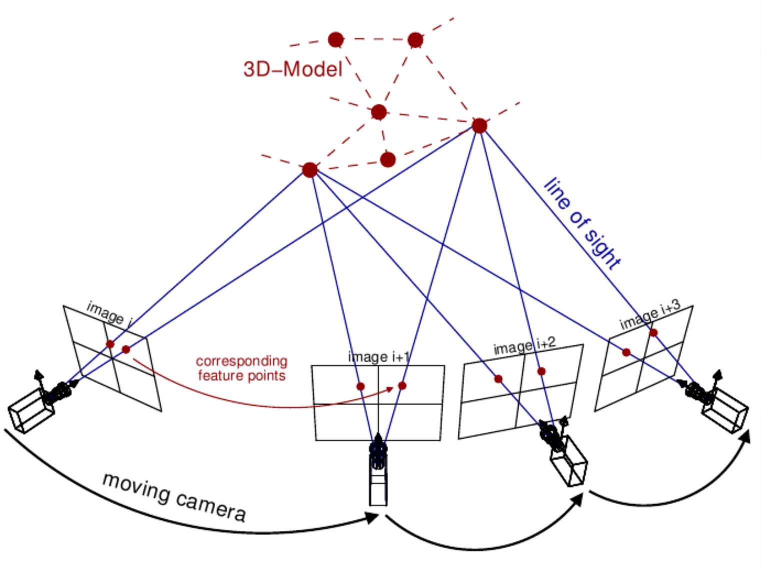

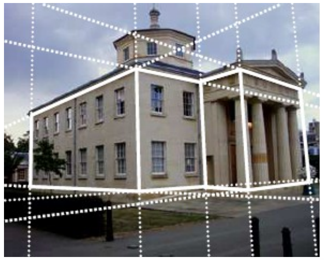

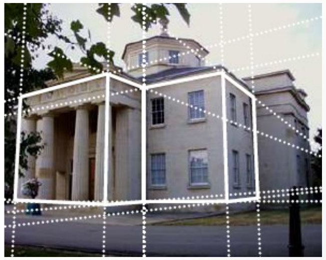

3D Stabilization

Motion Compensation

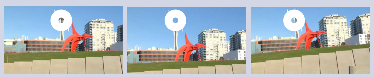

2D approach

3D approach

Subspace approach

- Motion Planning

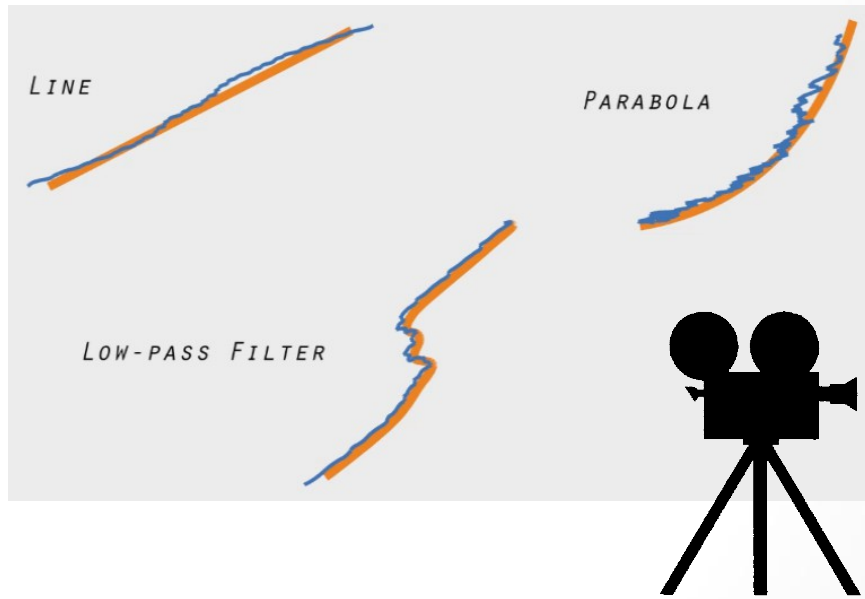

- Open question: How far can the output camera diverge from the input one before artefacts occur?

3D Stabilization

Image Composition

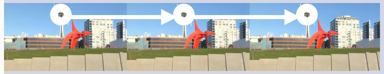

Full-frame warp

Novel view interpolation

Content-Preserving warp

- Image-Based-Rendering

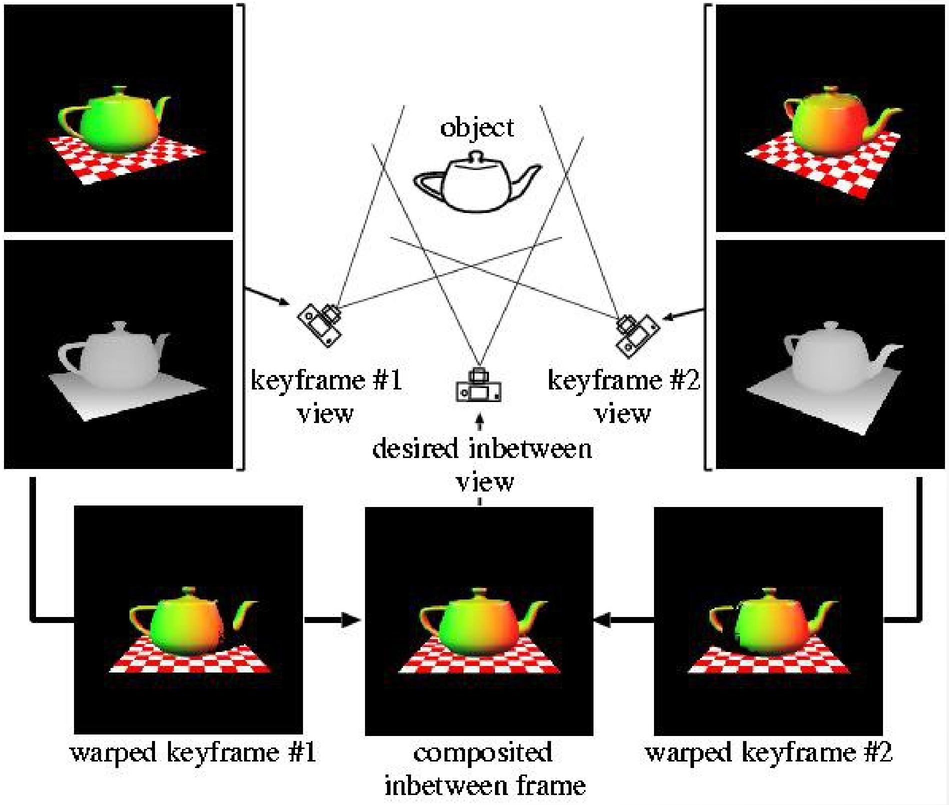

3D Stabilization

Stabilized video

3D Stabilization

- V Powerful
- V Can simulate 3D motions
- X Ghosting arising due to averaging of neighbor frames

→ **Content-Preserving warp** *[Liu et al. 2009]*

Content-Preserving warp

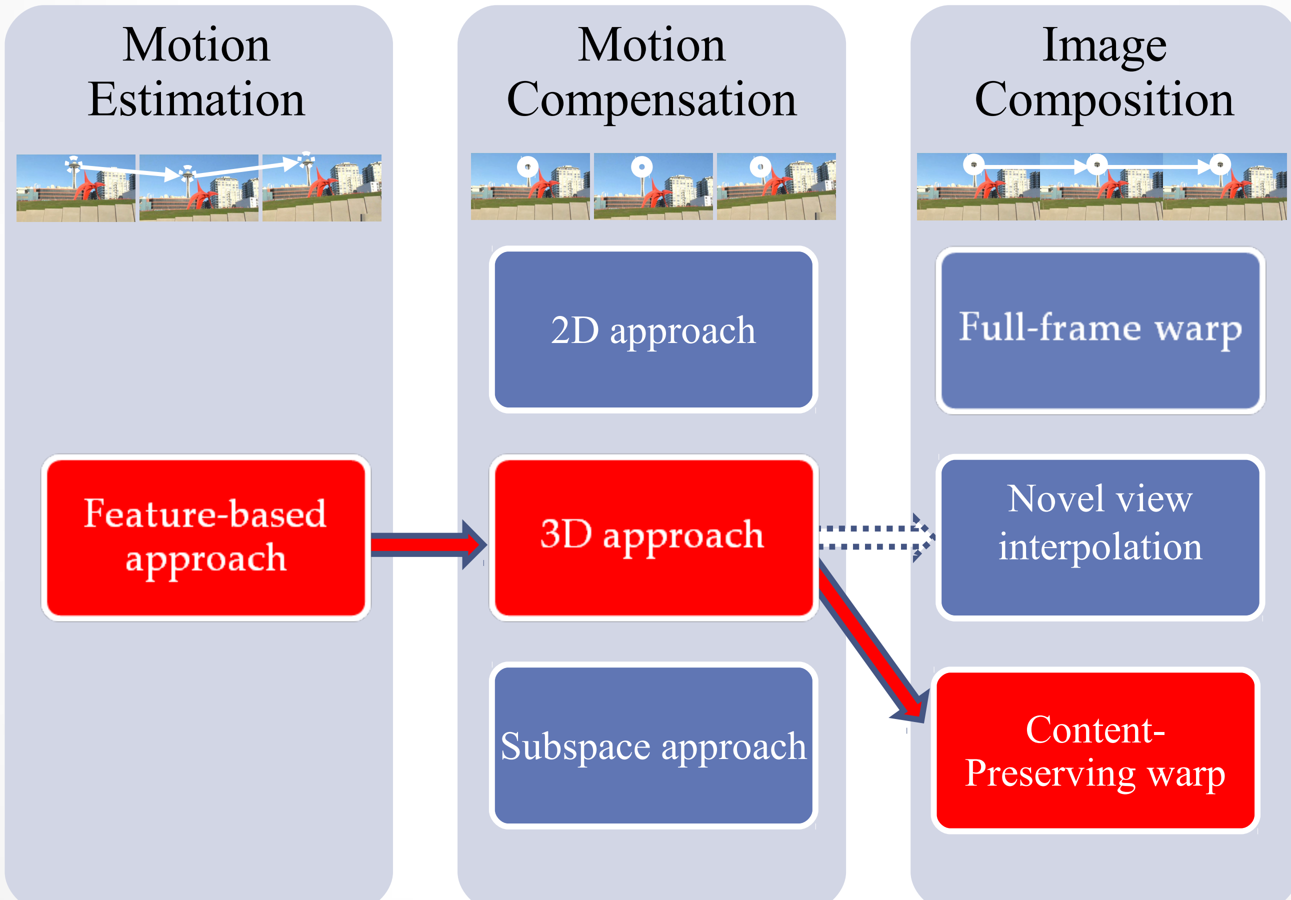

Content-Preserving warp

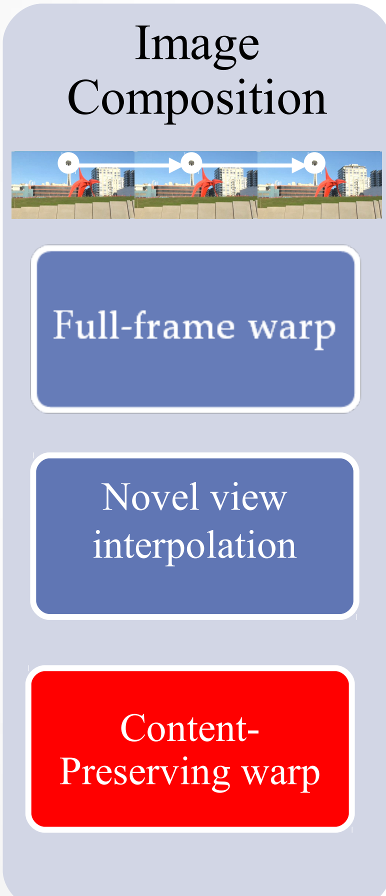

- Idea: warp each video frame independently to preserve content

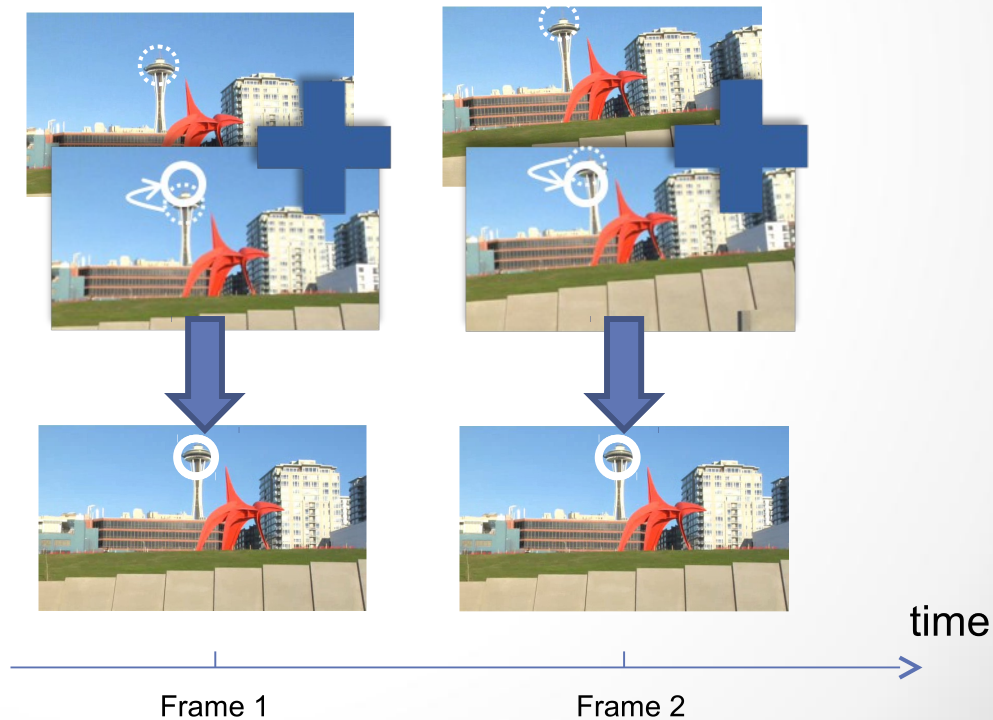

Content-Preserving warp

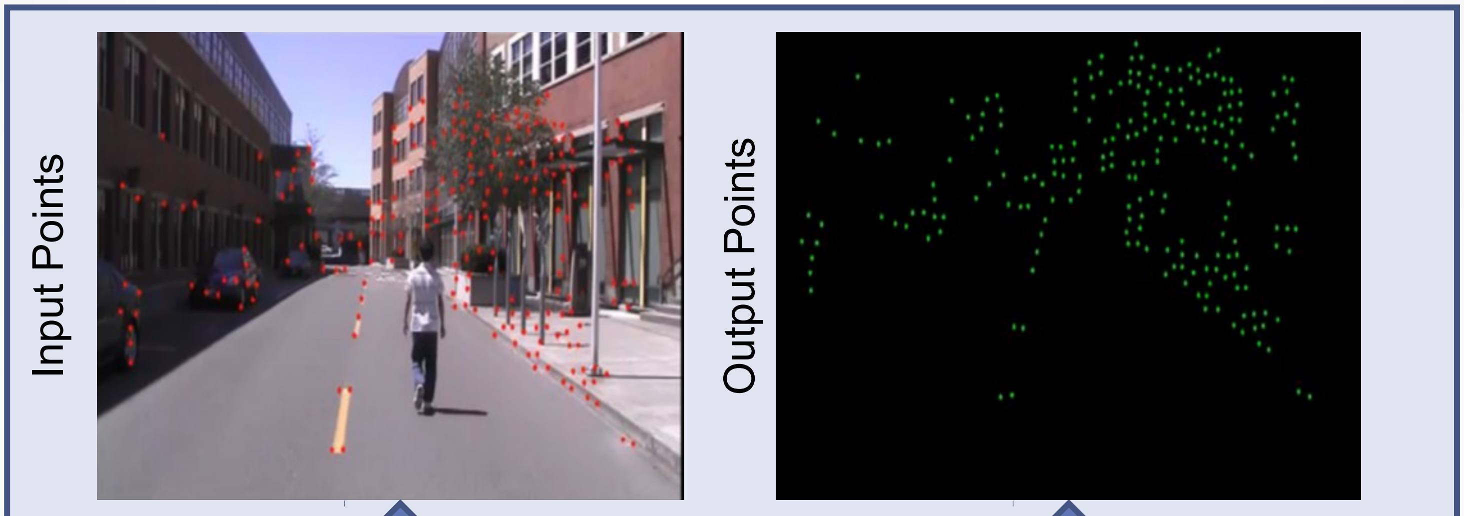

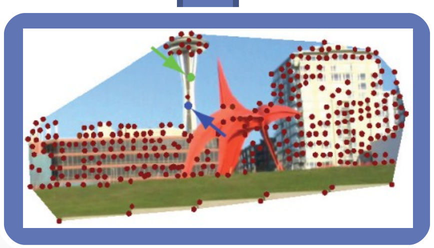

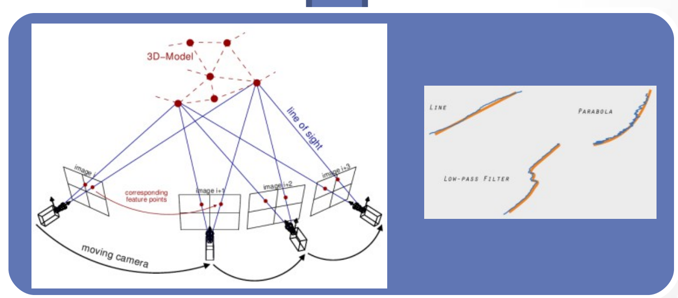

Content-Preserving warp

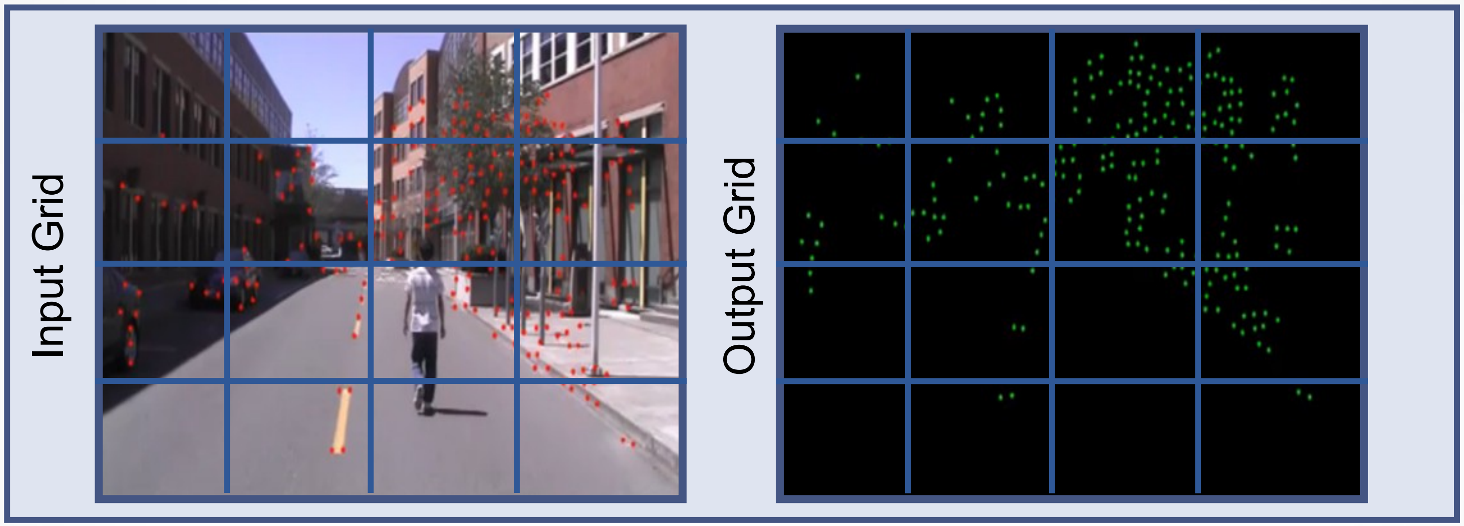

- $\widehat{P_k} := w_k^T \widehat{V_k}$

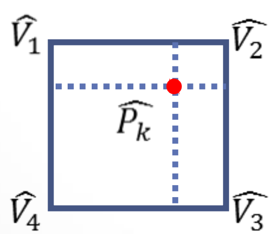

- $P_k := \widetilde{w_k^T} V_k$

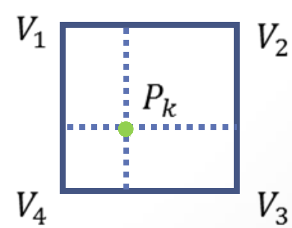

Content-Preserving warp

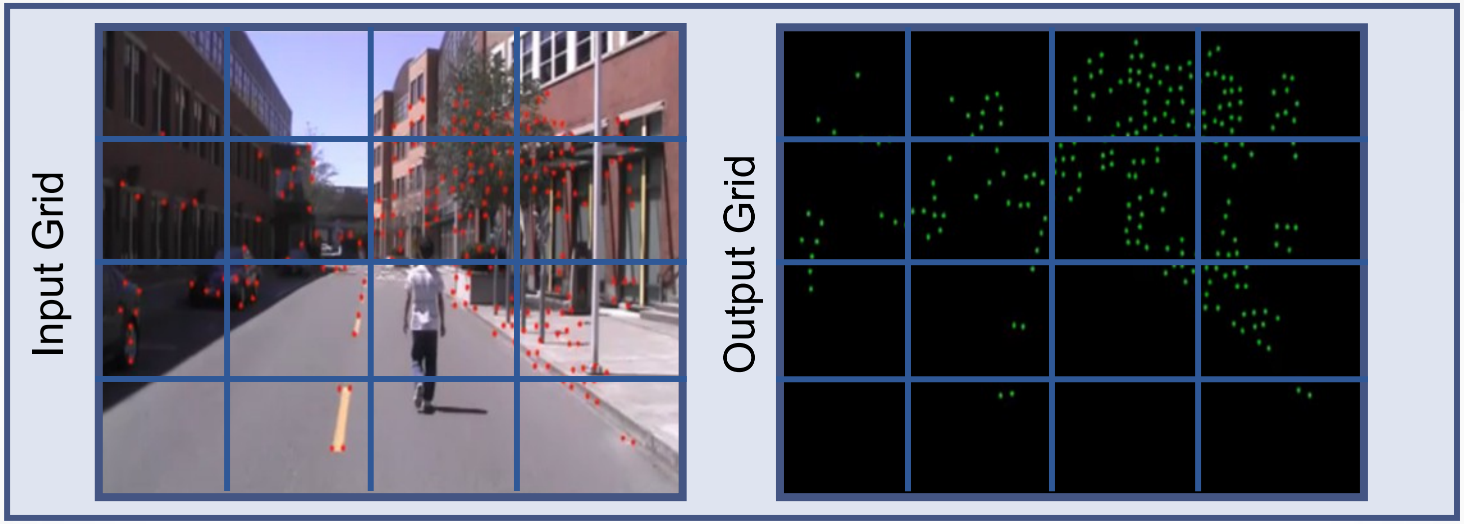

- $\widehat{P_k} := w_k^T \widehat{V_k}$
- $P_k := \widetilde{w_k^T} V_k$

$$E_d = \sum_k \left\| w_k^T V_k - P_k \right\|^2$$

Content-Preserving warp

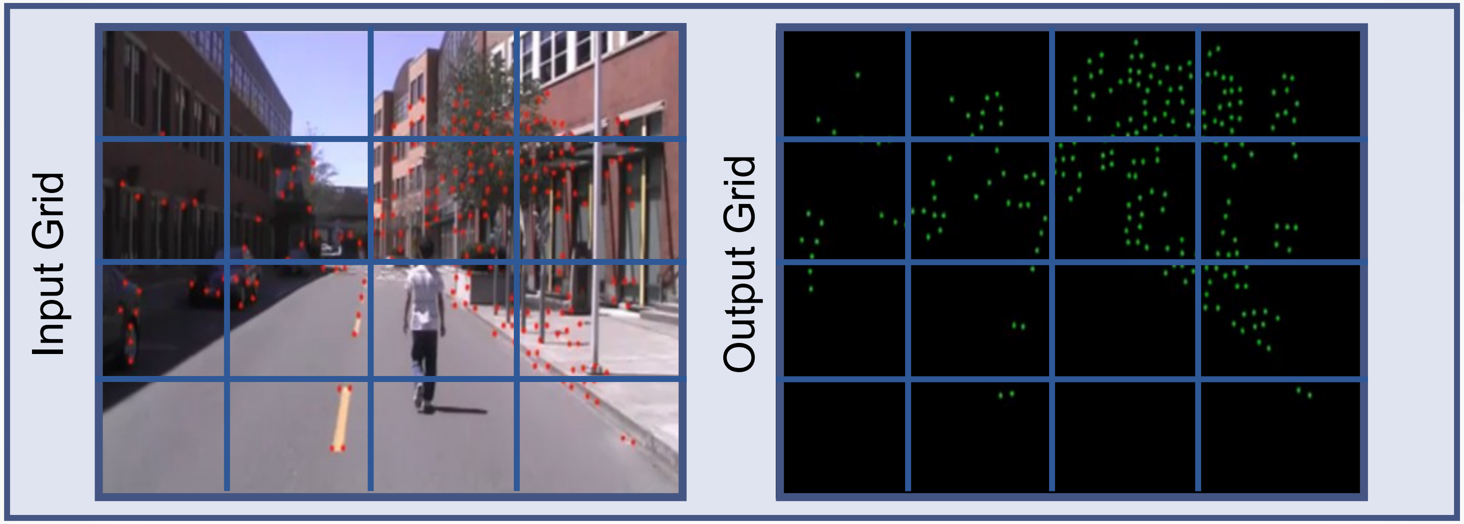

$$E_s(V_1) = w_s\|V_1 - \underbrace{(V_2 + u(V_3 - V_2) + vR_{90^\circ}(V_3 - V_2))}_{\widetilde{V_1}}\|^2$$

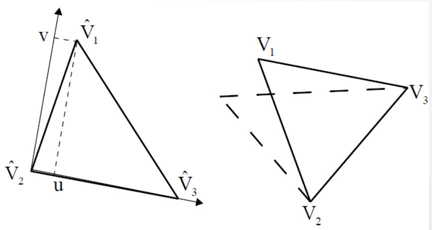

Content-Preserving warp

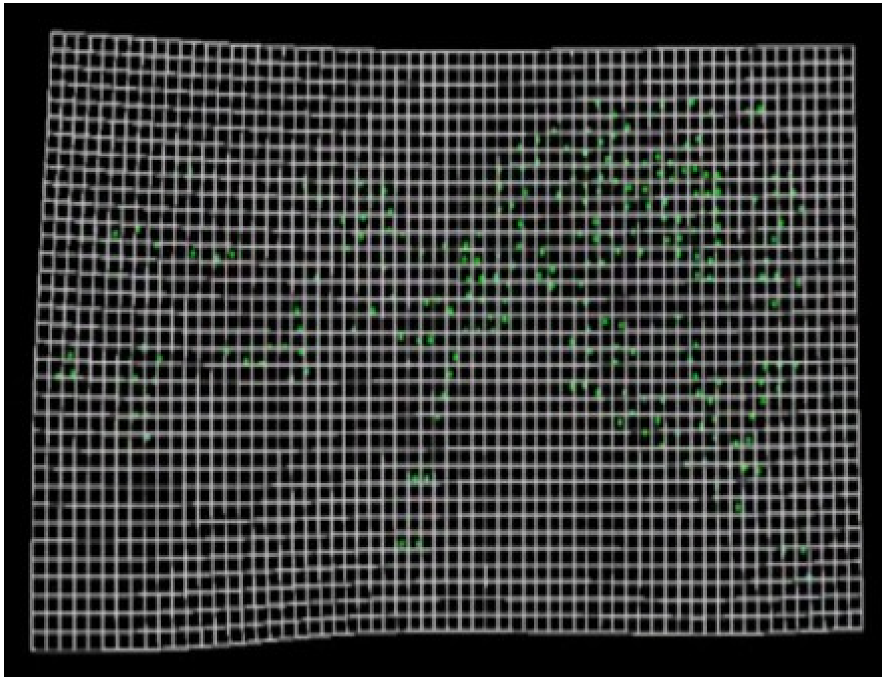

•Open question: Is it completely correct to warp moving objects in the scene just considering the background displacement? Would they ever appear distorted?

Content-Preserving warp

Content-Preserving warp

- V No ghosting
- V Stable
- X Structure-from-Motion not always successful
  - Lack of parallax
  - Camera zooming or just rotating
  - In-camera stabilization
  - Rolling shutter

→ Subspace Approach *[Liu et al. 2011]*

- Open question: What if we use a video-depth camera (with motion sensor) instead of a usual one?

Subspace Stabilization

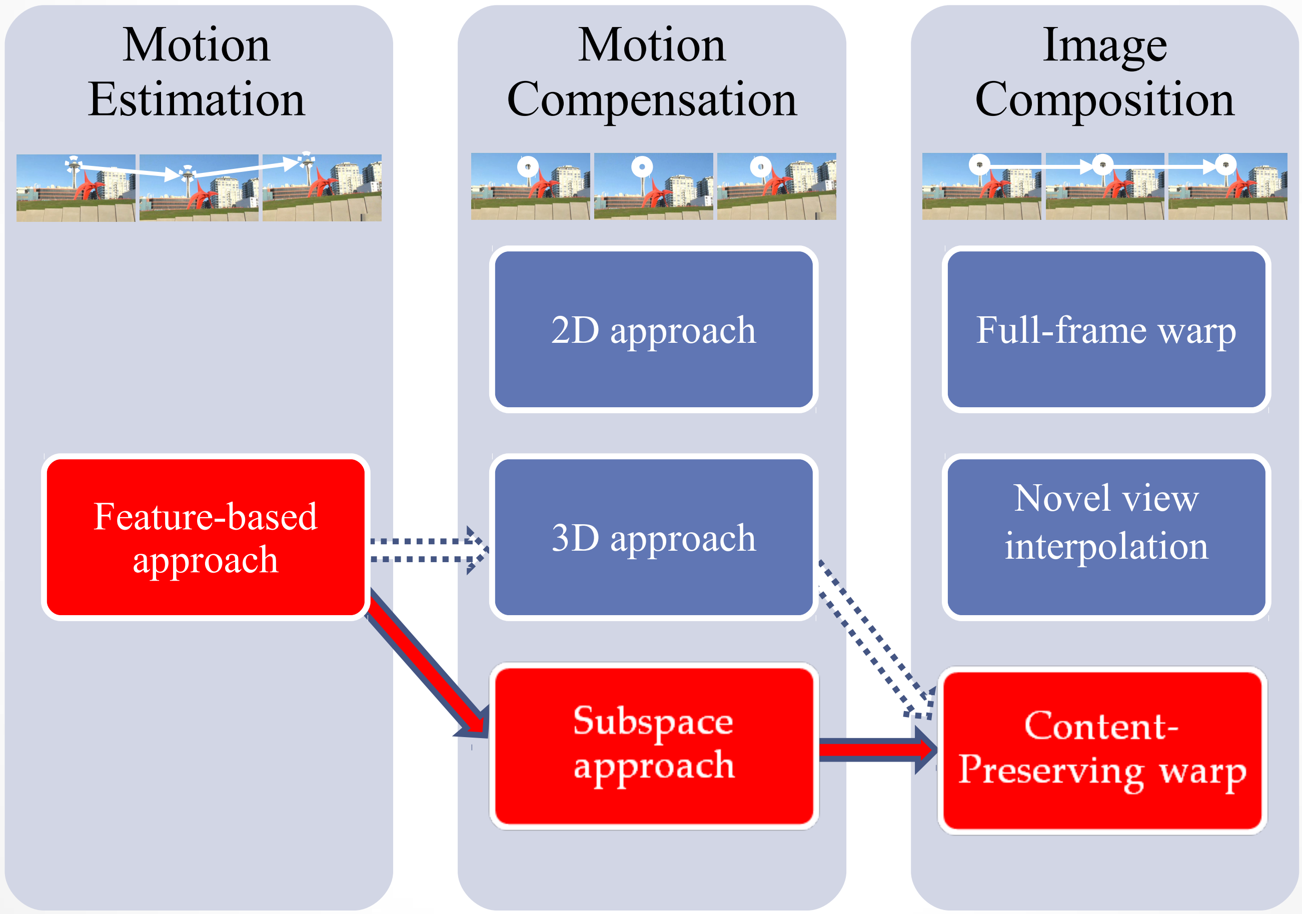

Subspace Stabilization

Motion Compensation

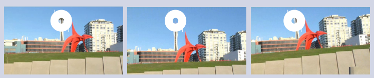

2D approach

3D approach

Subspace approach

When a rigid 3D scene is imaged by a moving affine camera, the observed motion trajectories should reside in a low-dimensional subspace. [Tomasi, Kanade 1992]

Motion trajectories from a perspective camera will lie on a non-linear manifold (!) instead of a linear subspace [Goh, Vidal 2007]

Idea: Approximate the manifold locally with a linear subspace.

Subspace Stabilization

- The trajectory matrix

$$M_{2N\times F} = \begin{bmatrix} x_1^1 & x_2^1 & \cdots & x_F^1 \\ y_1^1 & y_2^1 & \cdots & y_F^1 \\ & & \vdots & \\ x_1^N & x_2^N & \cdots & x_F^N \\ y_1^N & y_2^N & \cdots & y_F^N \end{bmatrix}$$

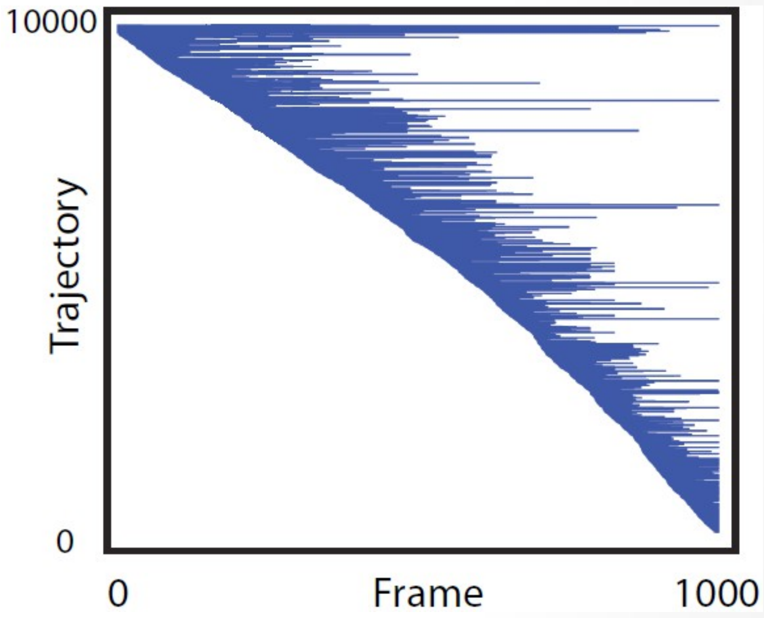

Low-rank constraint: a trajectory matrix for instantaneous motions should have at most rank 9 [Irani, 2002]

Subspace Stabilization

- Moving Factorization

$$M^i_{2m\times k} = C_{2m\times r}E_{r\times k}$$

$$M^0 = \begin{bmatrix} A^{00} & A^{01} \\ A^{10} & \boxed{A^{11}} \end{bmatrix} = \begin{bmatrix} C^0 \\ C^1 \end{bmatrix} \begin{bmatrix} E^0 & E^1 \end{bmatrix}$$

$$M^1 = \begin{bmatrix} \boxed{A^{11}} & A^{12} \\ A^{21} & A^{22} \end{bmatrix} = \begin{bmatrix} C^1 \\ C^2 \end{bmatrix} \begin{bmatrix} E^1 & E^2 \end{bmatrix}$$

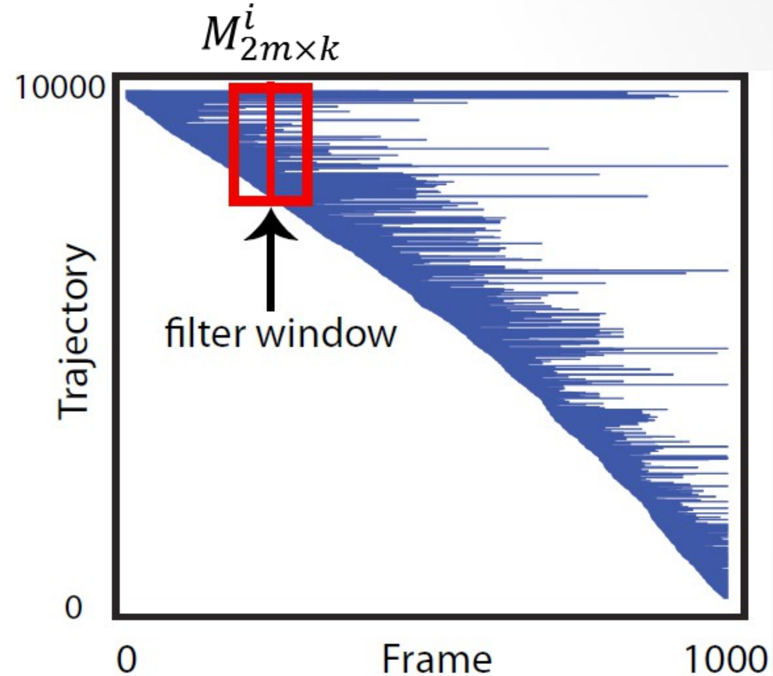

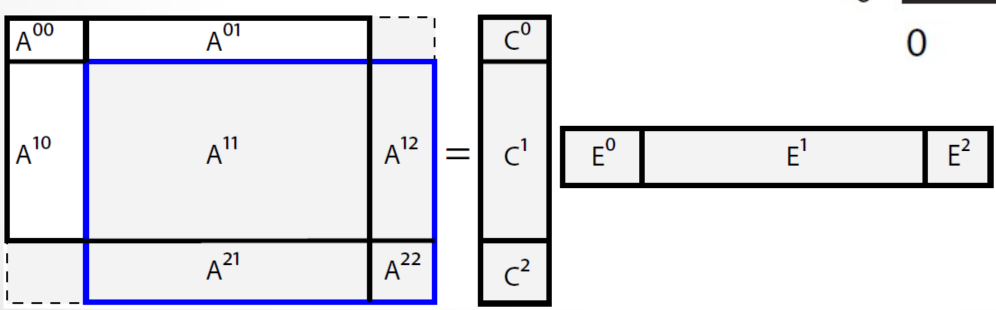

Subspace Stabilization

- Smoothing the Eigen-trajectories

Binary matrix {0, 1}

$$\widehat{M^i} = W * (C(EK))$$

Smoothing Kernel

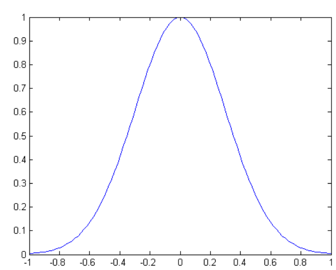

Subspace Stabilization

- Smoothing the Eigen-trajectories

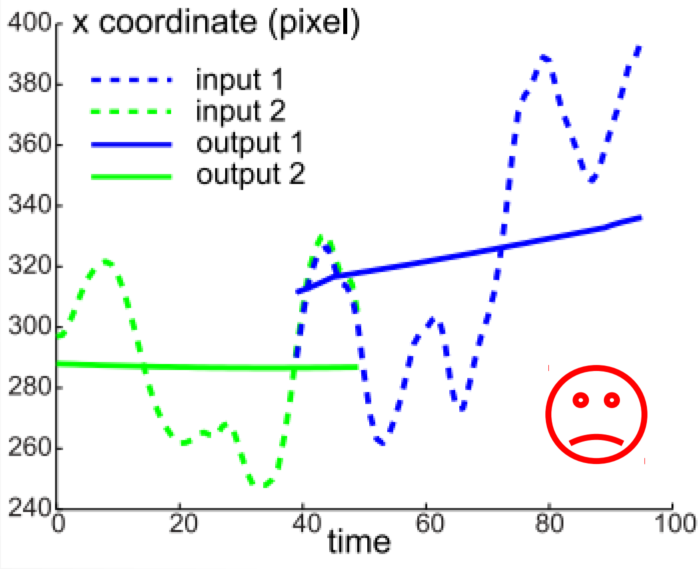

Filtering the trajectories independently

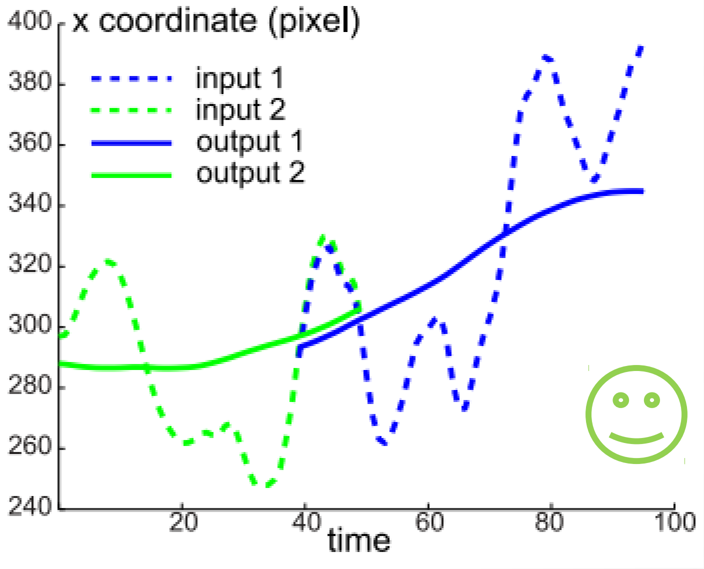

Filtering the Eigen-trajectories

Subspace Stabilization

Subspace Stabilization

- V High quality without computing a 3D reconstruction
- V Efficient
- V Can handle complex Video
- X No motion planning possible
- X Relies on an approximation

- Open question: How to evaluate videos if we cannot compare them on the physically-correct base?

Conclusions

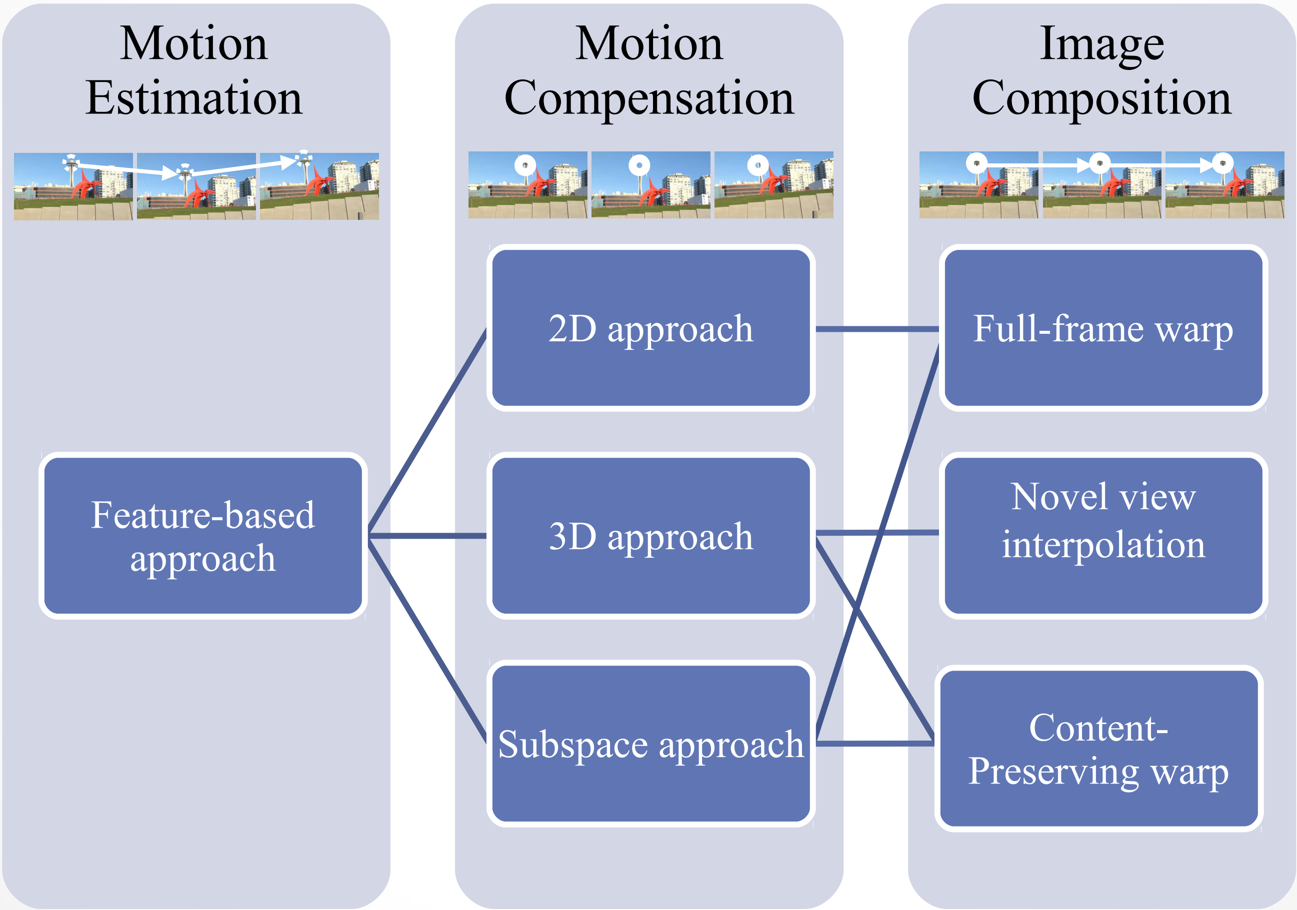

Conclusions

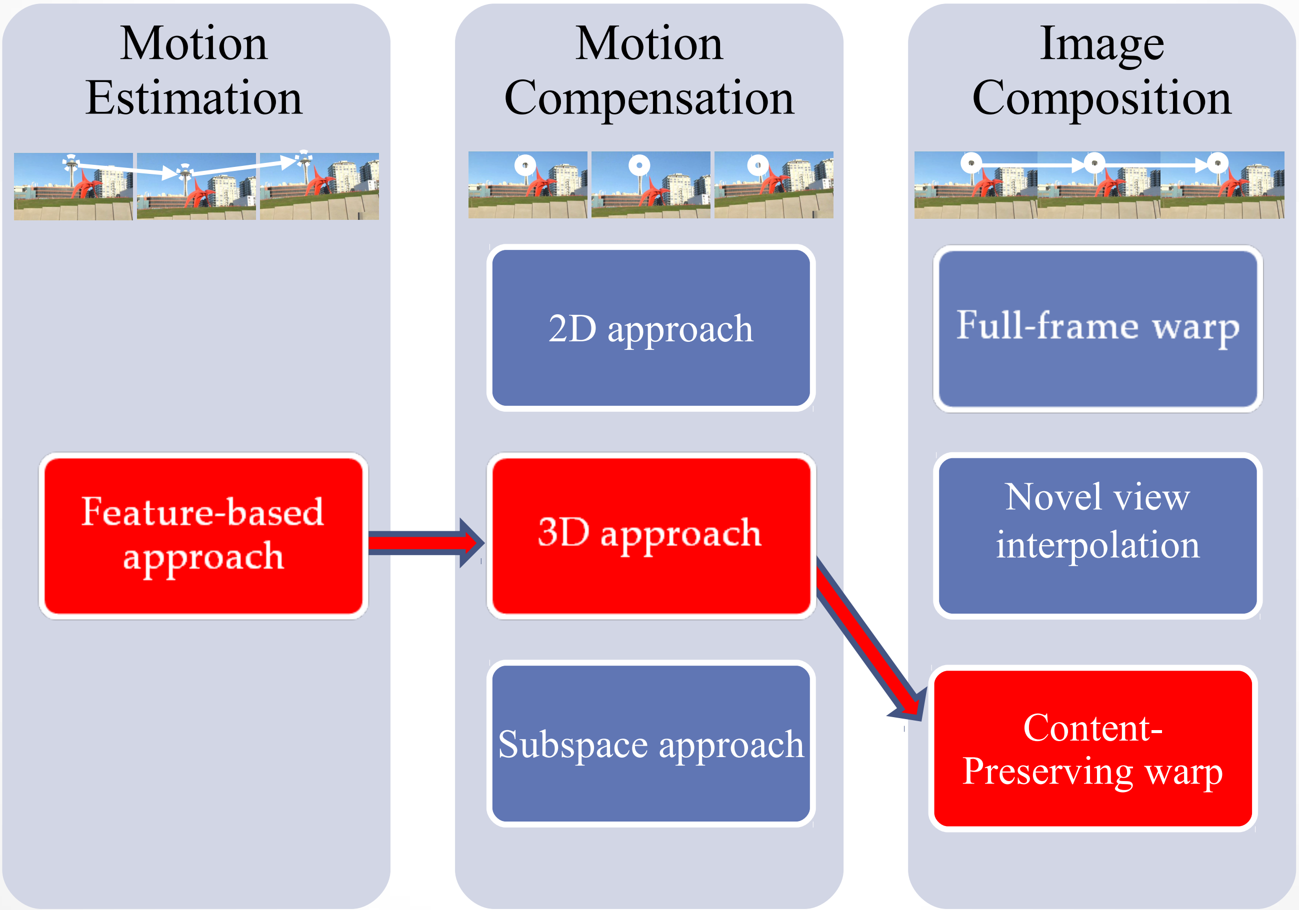

Conclusions

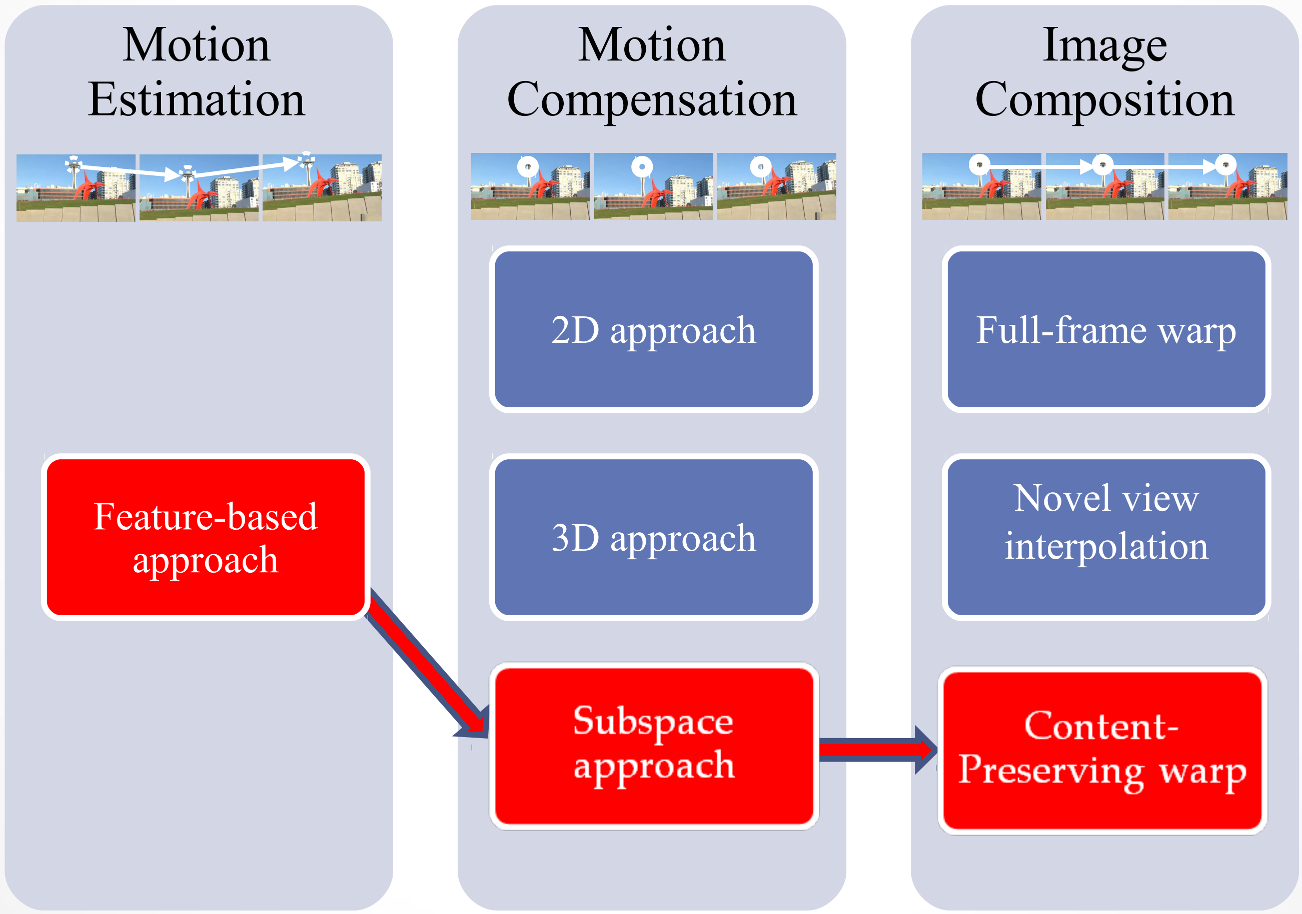

Discussion

Discussion

- Challenges
  - Motion blur
  - Excessive shake
  - Geometry with little texture
  - Large moving objects that occlude/dominate the scene

Discussion

- How far can the output camera diverge from the input one before artefacts occur?

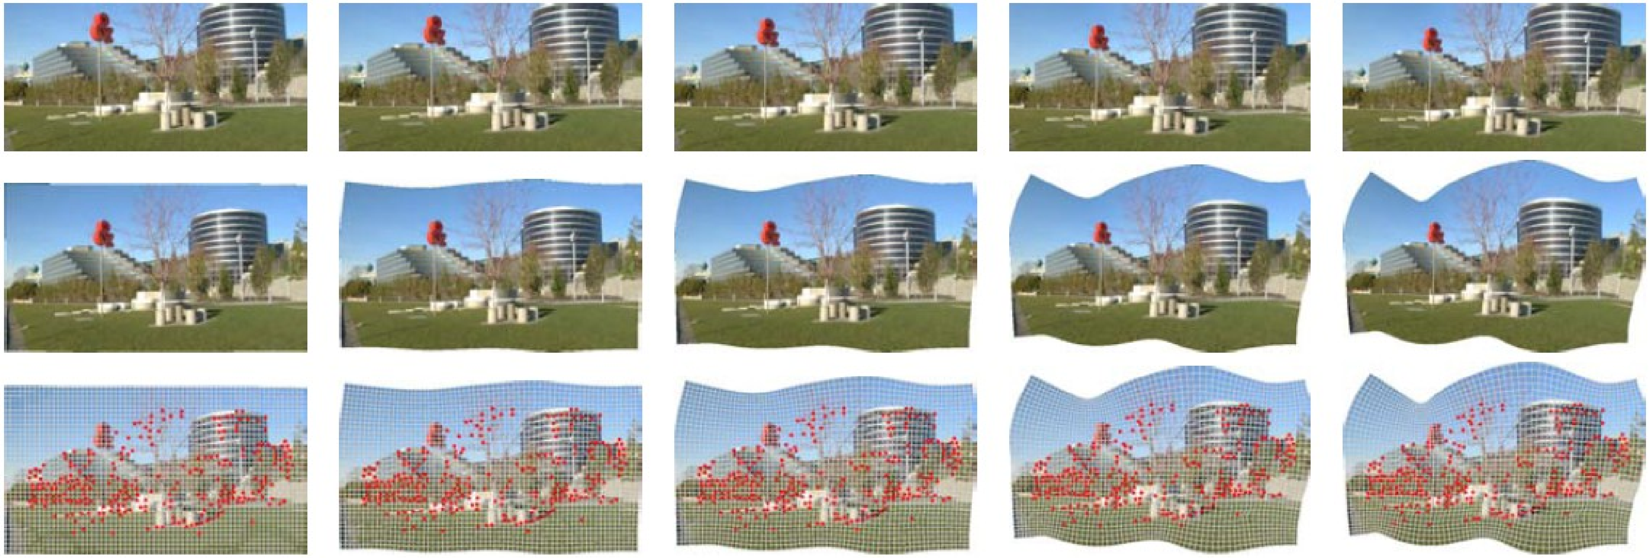

Figure 6: *Each row shows a sequence of warps of a single input video frame created by pulling the camera away from its original location. The top row shows the final cropped result; the middle shows the entire warp result; the bottom shows the corresponding grids and the points that guide the warp. For small camera motions the warps look reasonable, but they become visibly distorted at larger camera displacements.*

Discussion

- What if we use a video-depth camera (with motion sensor) instead of a usual one?

Discussion

- Is it completely correct to warp moving objects in the scene just considering the background displacement? Would they ever appear distorted?

Discussion

- How to evaluate videos if we cannot compare them on the physically-correct base?