

Schedule

- Tuesday, May 10:
 - Motion microscopy, separating shading and paint
- Thursday, May 12:
 - 5-10 min. student project presentations, **projects due.**

Computer vision for photography

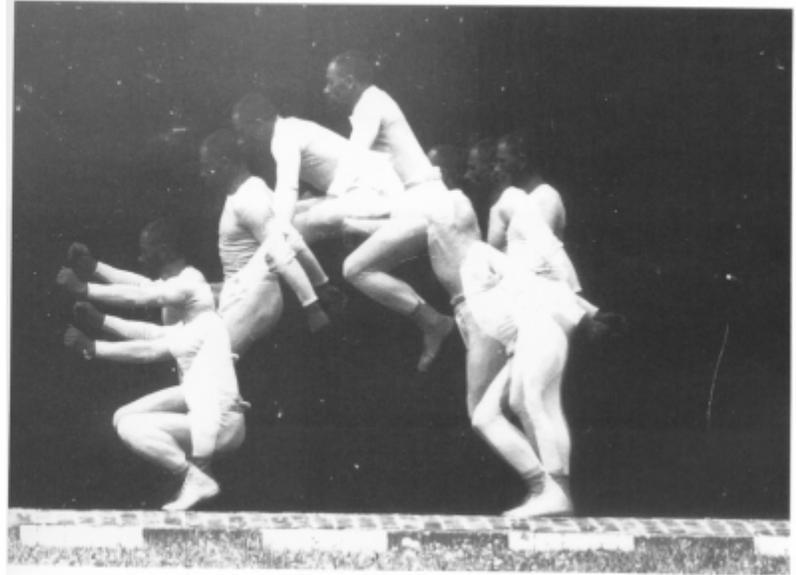
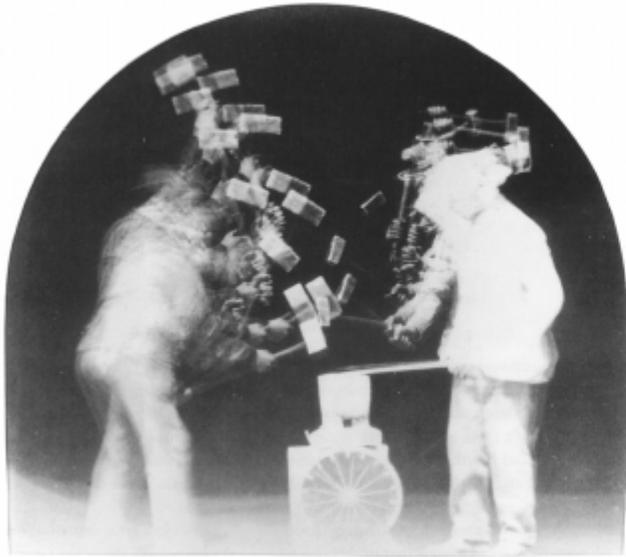
Bill Freeman

Computer Science and Artificial Intelligence

Laboratory,

MIT

Multiple-exposure images by Marey



Edgerton



Computational photography

Update those revealing photographic techniques with digital methods.

- 1) Shapetime photography
- 2) Motion microscopy
- 3) Separating shading and paint

Shapetime photography

Joint work with Hao Zhang, U.C. Berkeley

Video frames



Multiple-exposure

Layer-By-Time



Shape-Time



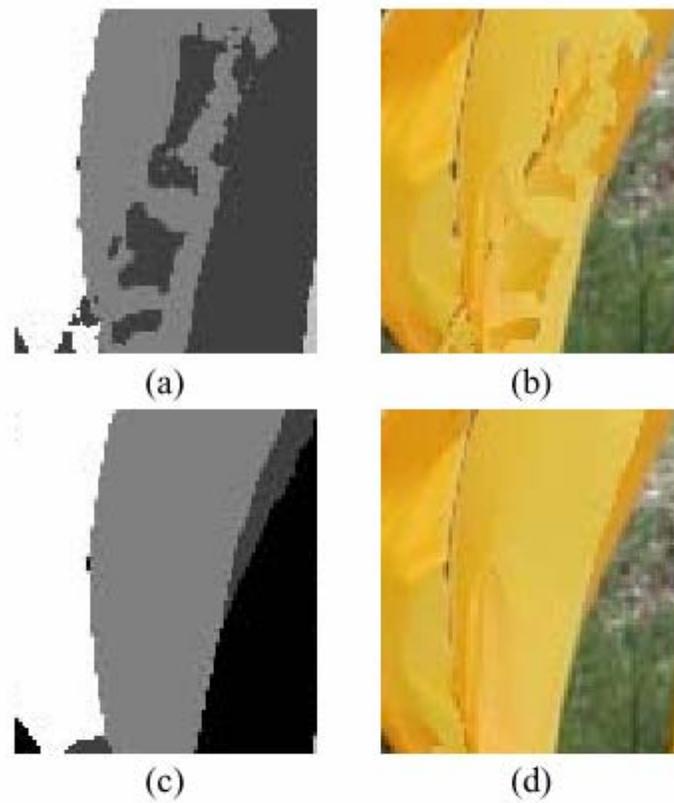


Figure 4: (a) Layer assignments, without MRF processing. (b) Shape-time image based on those assignments. (c) Most probable layer assignments, computed by MRF. (d) Resulting shape-time image.

number of pixels in the image.

$$P(\vec{t}) = \frac{1}{Z} \prod_{(jk)} \phi_{jk}(t_k, t_j) \prod_k \psi_k(t_k), \quad (2)$$

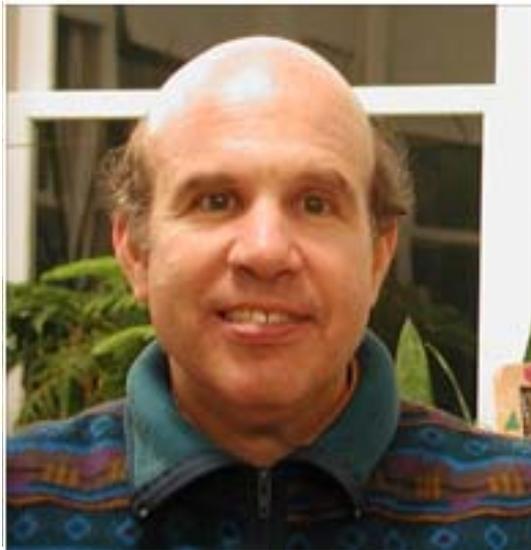
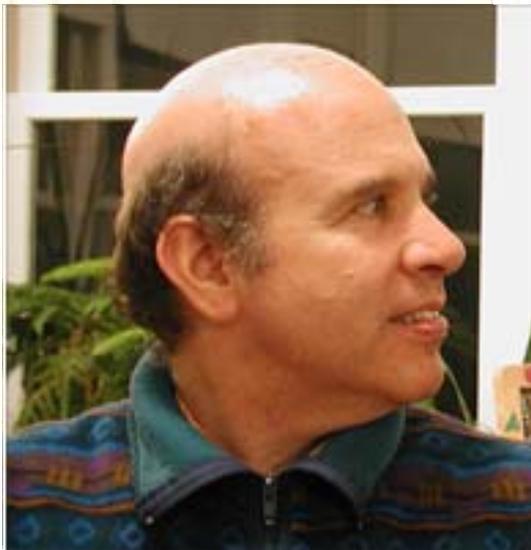


(a)





“how to sew”



Input sequence



Shape-Time
composite



“inside-out”

Insert pictures describing zcam, and
initial results

Motion Magnification



Ce Liu Antonio Torralba William T. Freeman

Fredo Durand Edward H. Adelson

Goal

A microscope for motion

You focus the microscope by specifying which motions to magnify, and by how much;

the motion microscope then re-renders the input sequence with the desired motions magnified.

The naïve solution has artifacts

Original sequence

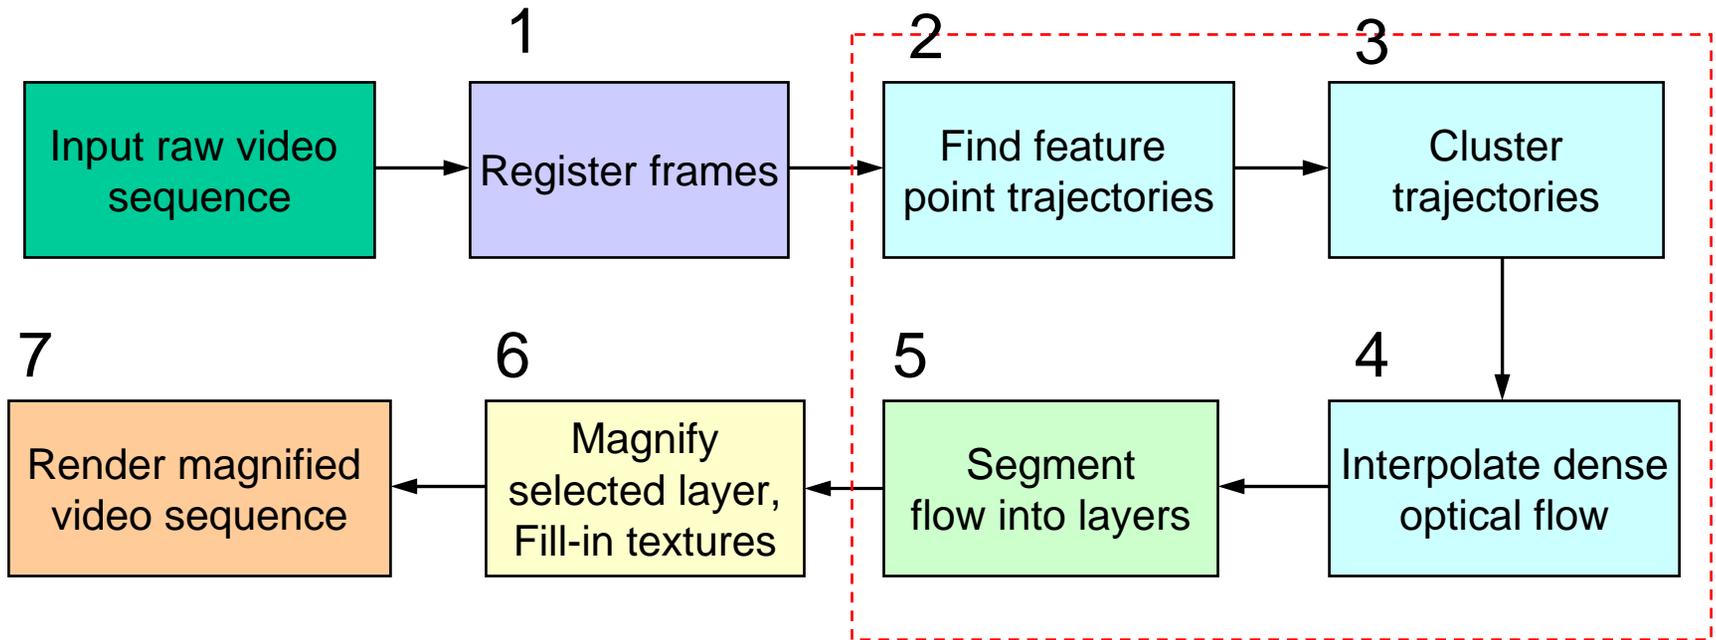


Naïve motion magnification



Amplified dense Lukas-Kanade optical flow.
Note artifacts at occlusion boundaries.

Motion magnification flowchart

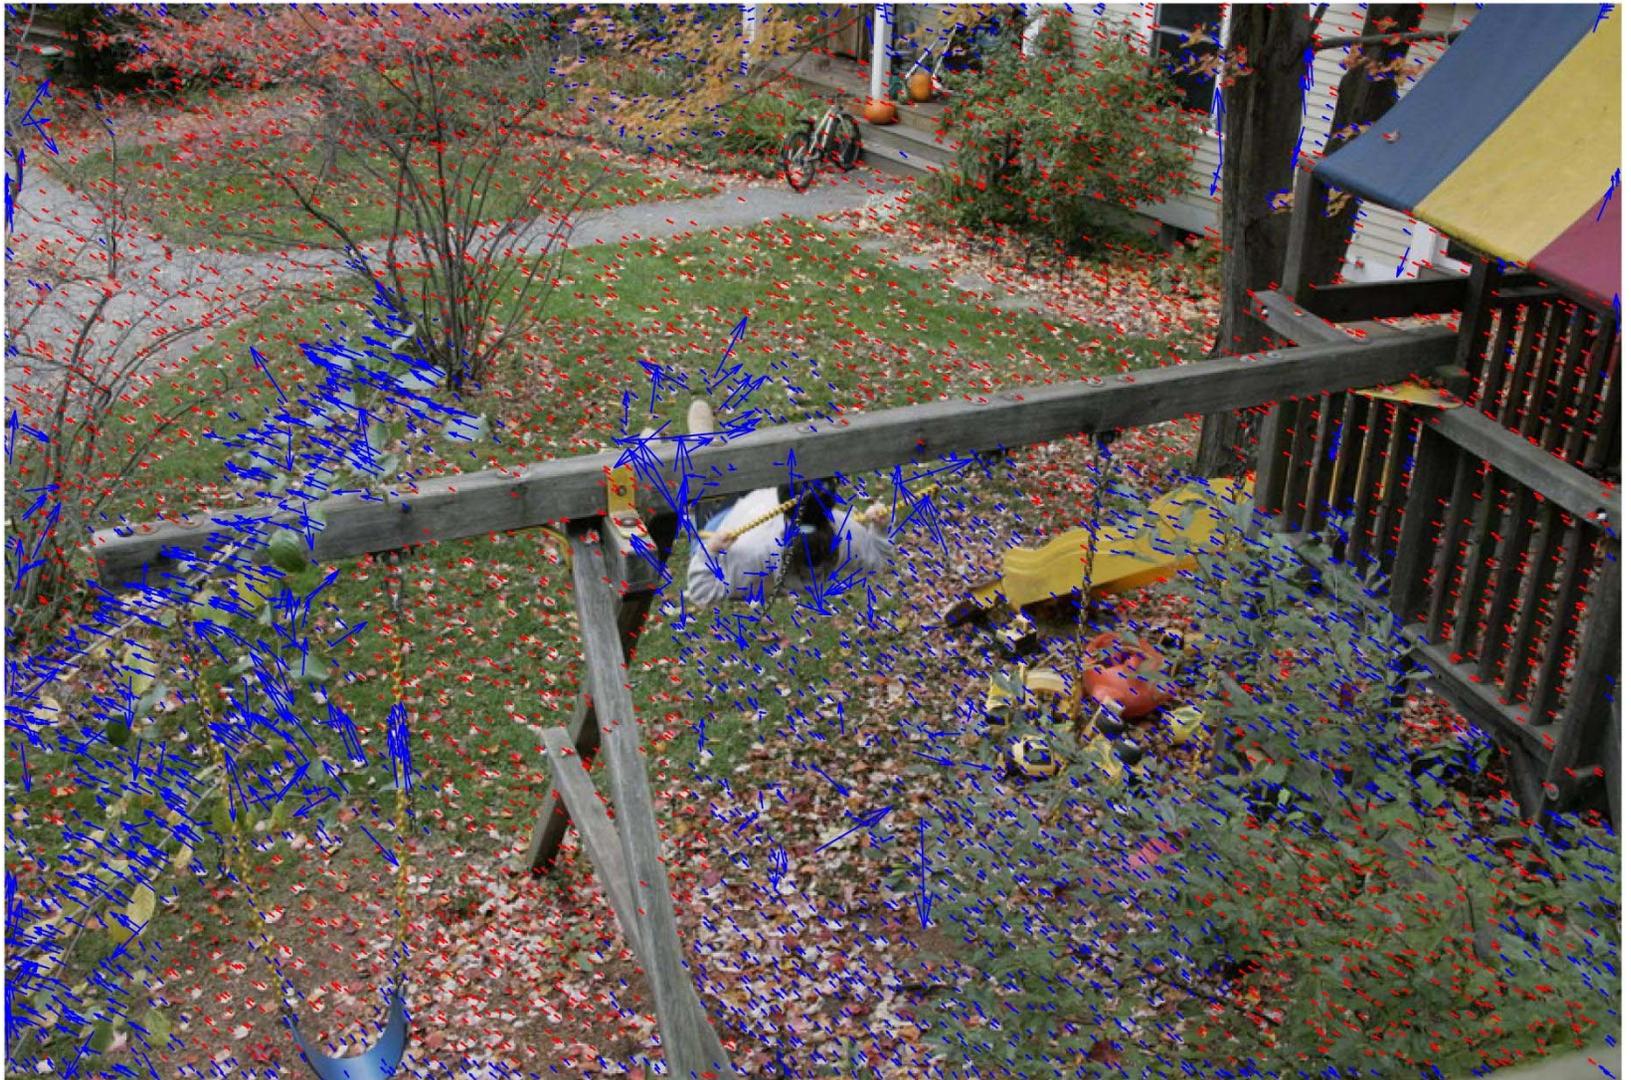


Layer-based motion analysis

1 Video Registration

- To find a reliable set of feature points that are “still” in the sequence
 - Detect and track feature points
 - Estimate the affine motion from the reference frame to each of the rest frames
 - Select feature points that are inliers through all the frames
 - Affine warping based on the inliers

Inliers (red) and outliers (blue)



Registration results



2 Find feature point trajectories

- An EM algorithm to find both trajectory and region of support for each feature point
 - E-step: to use the variance of matching score to compute the weight of the neighboring pixels
 - M-step: to track feature point based on the region of support
- The following feature points are pruned
 - Occluded (matching error)
 - Textureless (motion coherence)
 - Static (zero motion)

Learned regions of support for features

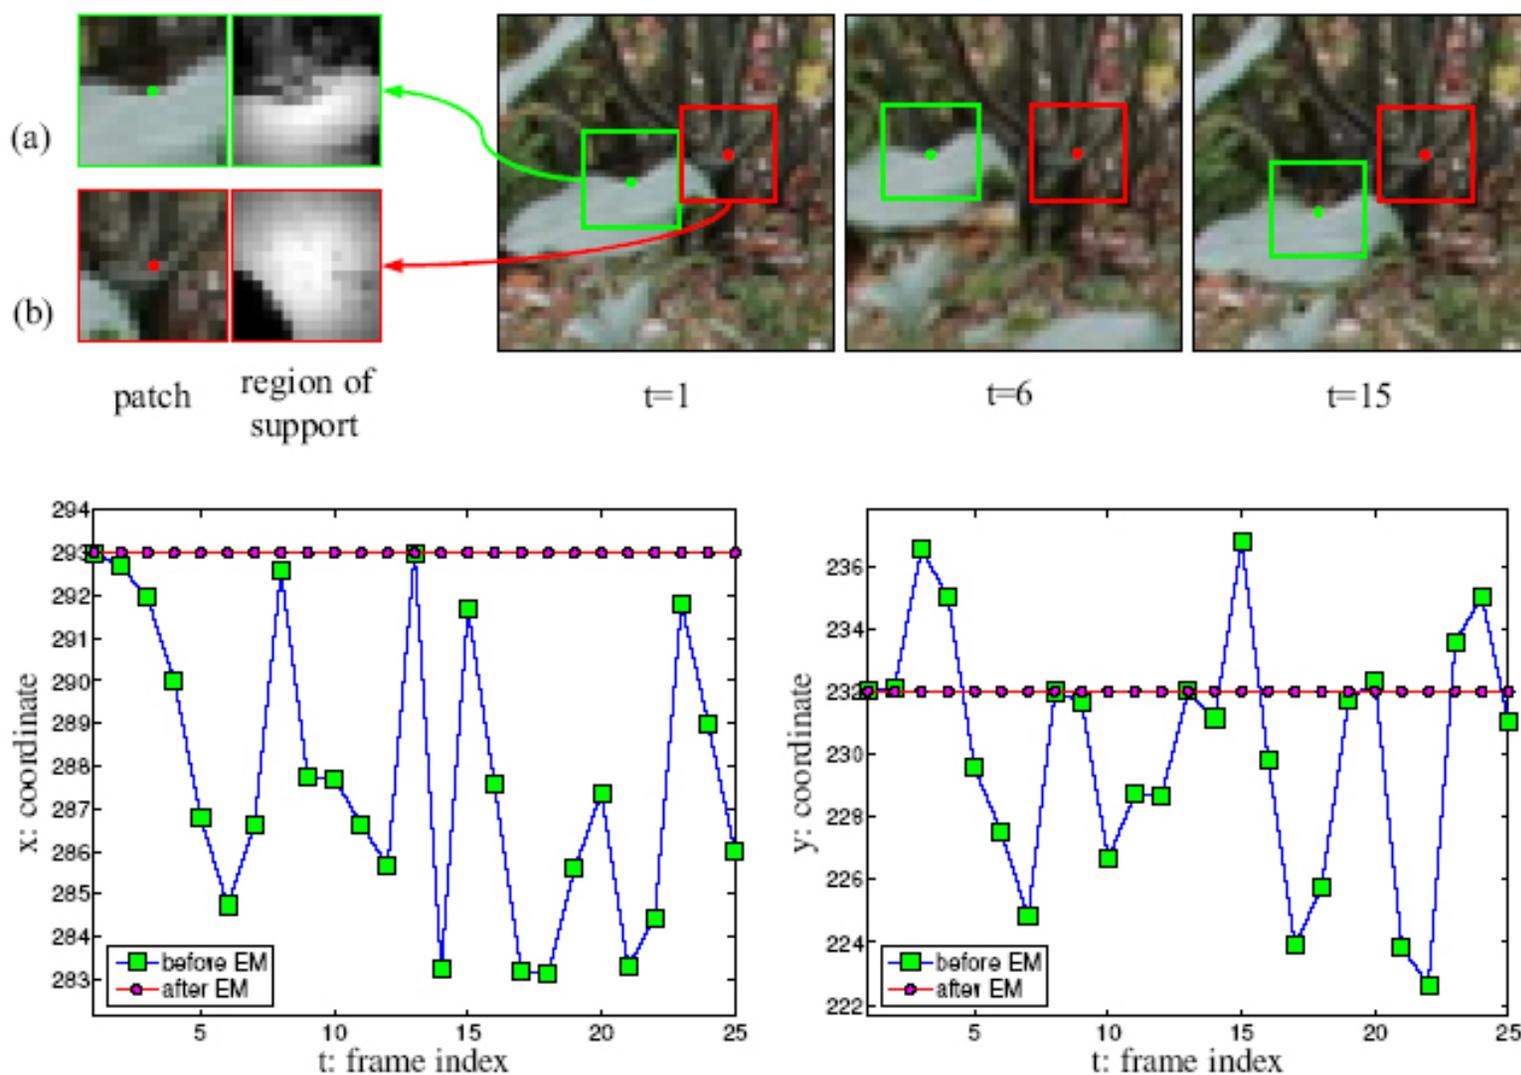


Figure 3: Learned regions of support allow features (a) and (b) to reliably track the leaf and background, respectively, despite partial occlusions. For feature (b) on the stationary background, the plots show the x (left) and y (right) coordinates of the track both with (red) and without (blue) a learned region of support for appearance

Robust feature point tracking

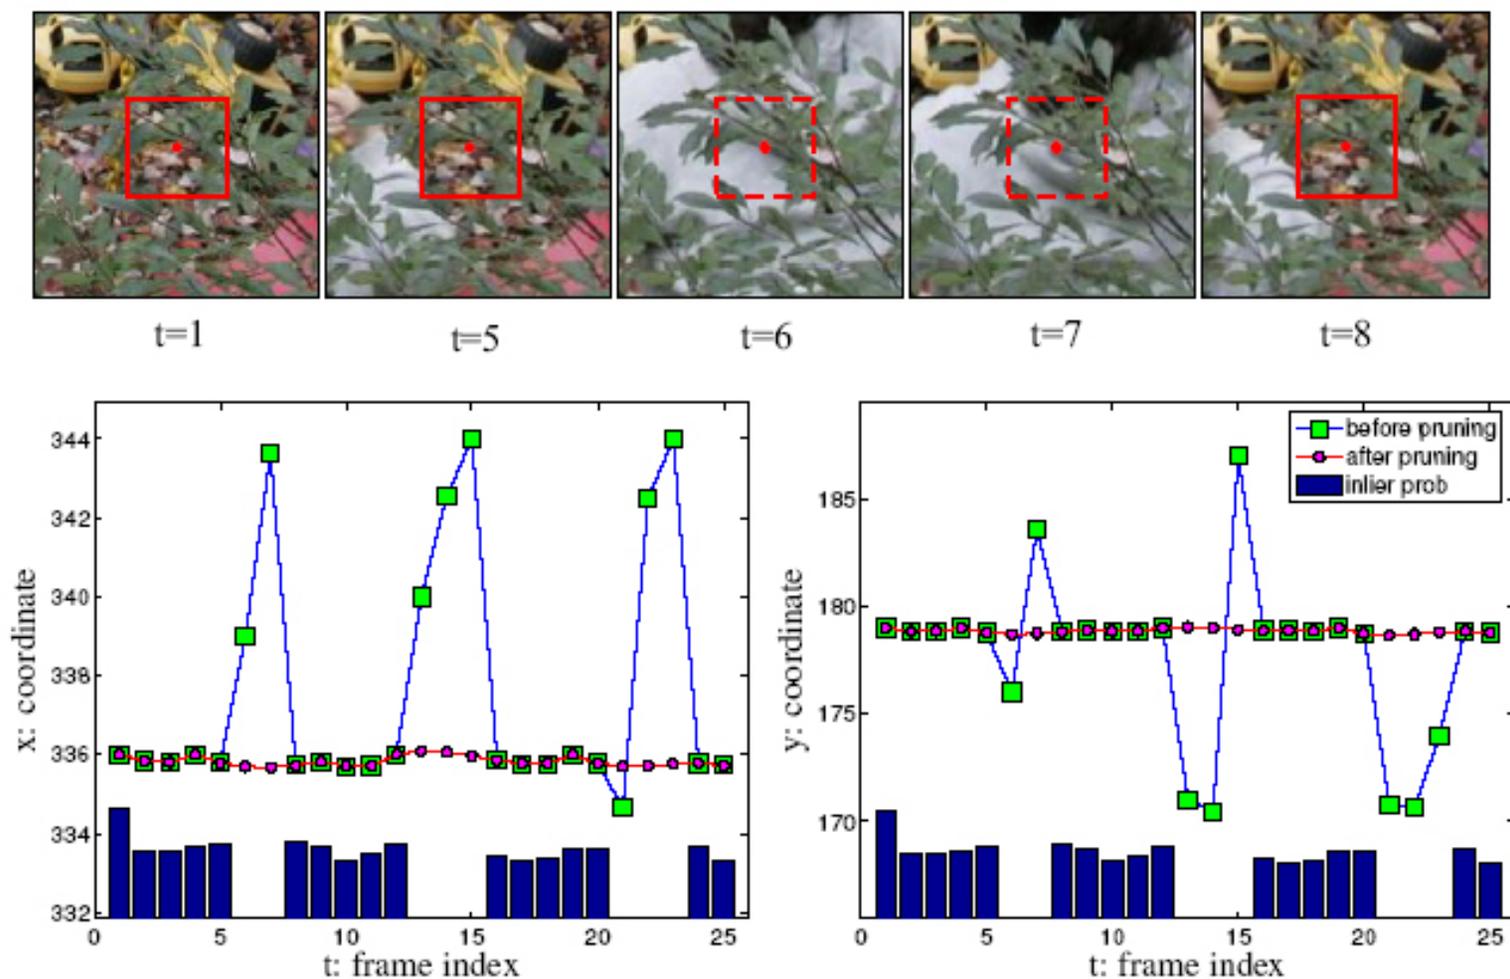
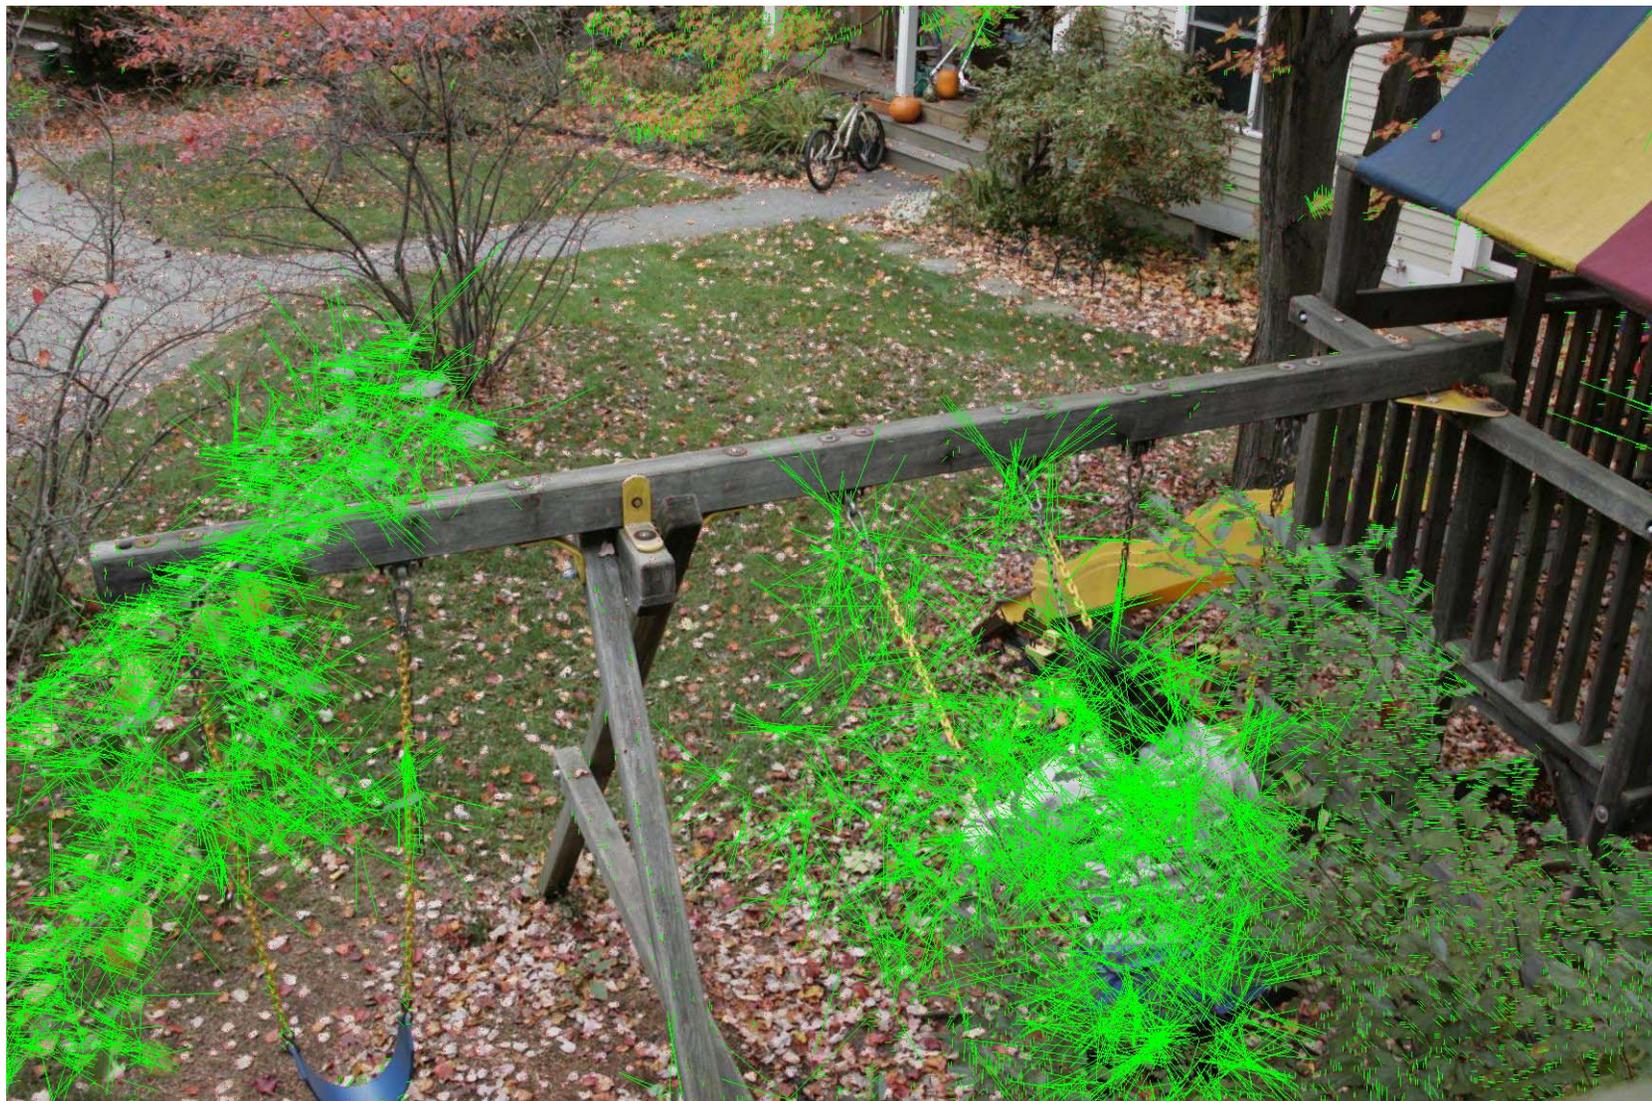


Figure 4: Top images show a feature point on the stationary background layer becoming occluded during frames 6 and 7. Below are the x- and y- coordinates of the tracked feature, showing outlier positions. These can be identified from the inlier probabilities shown as a bar plot (repeated for comparison with each plot) and replaced with smoothed values.

Minimal SSD match to find feature point trajectories



Use EM to find regions of support and prune low-likelihood trajectories

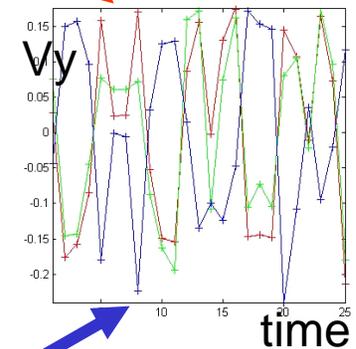
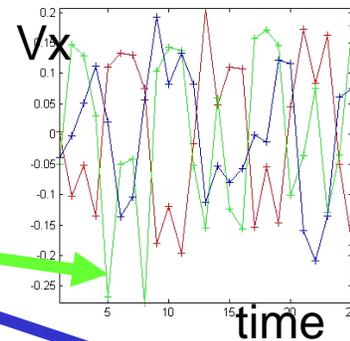
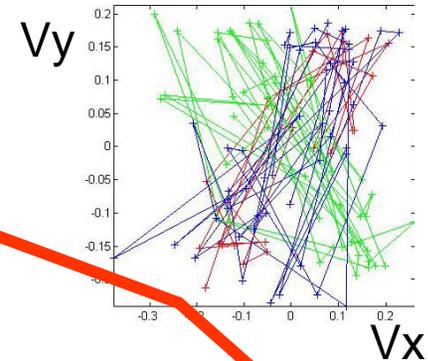
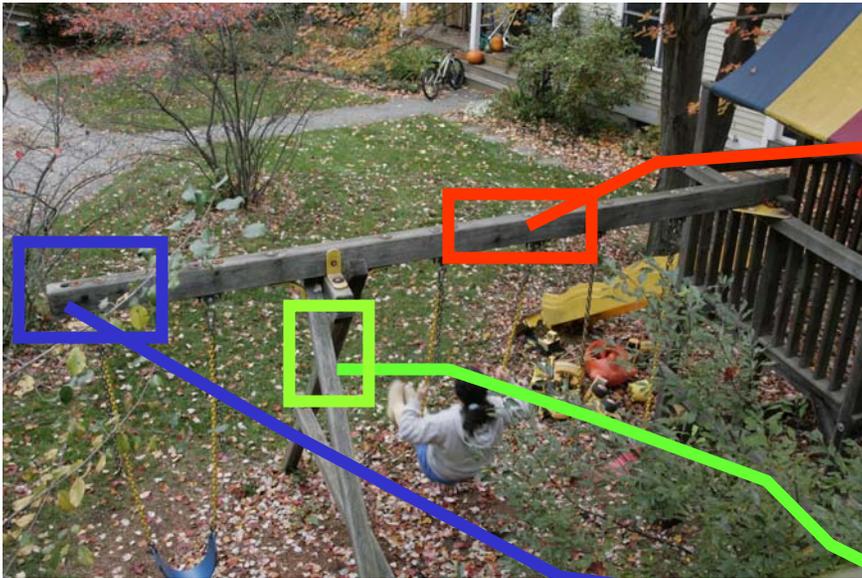


3 Trajectory clustering

We need to cluster trajectories belonging to the same object, despite:

Points have different appearances

Undergo very small motions, of varying amplitudes and directions



3 Compatibility function used to group feature point trajectories

$\rho_{n,m}$: compatibility between the n th and m th point trajectories.

$v_x(n,k)$: the displacement, relative to the reference frame, of the n th feature point in the k th frame.

$$\rho_{n,m} = \left| \frac{\sum_k (v_x(n,k) + jv_y(n,k)) (v_x(m,k) + jv_y(m,k))}{\sqrt{(\sum_k v_x(n,k)^2 + v_y(n,k)^2) (\sum_k v_x(m,k)^2 + v_y(m,k)^2)}} \right|$$

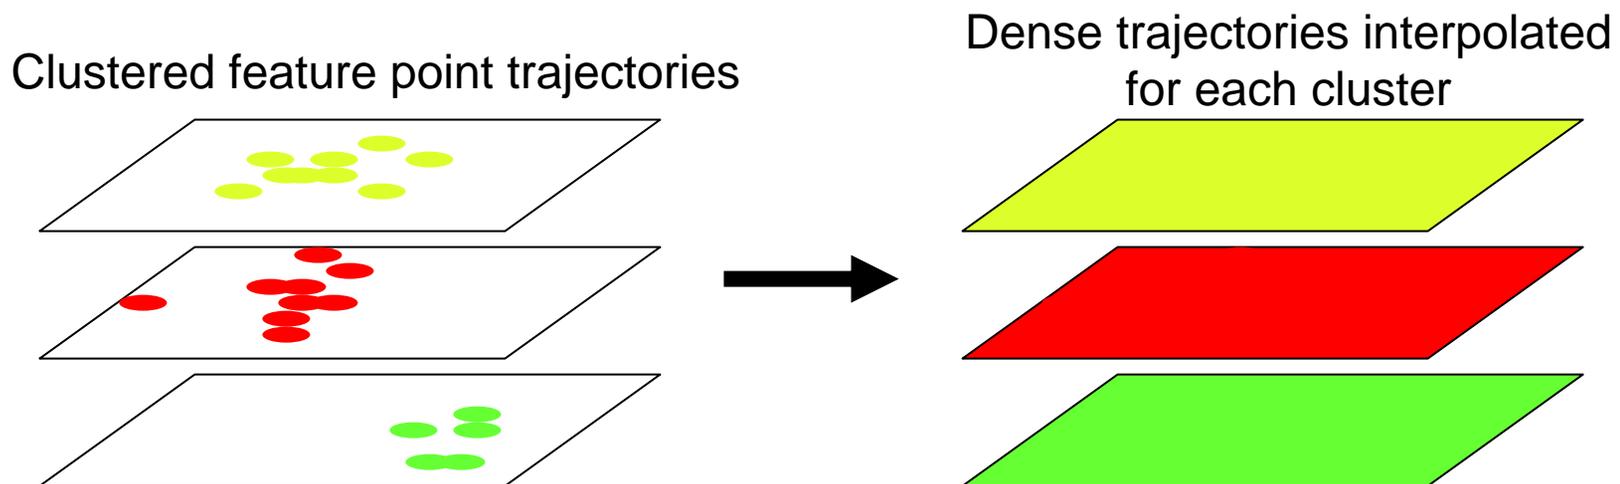
Using the $\rho_{n,m}$ compatibilities, cluster the point trajectories using normalized cuts.

Clustering results



4 Dense optical flow field interpolation

- For each layer (cluster) a dense optical flow field (per pixel) is interpolated
- Use local weighted linear regression to interpolate between feature point trajectories.



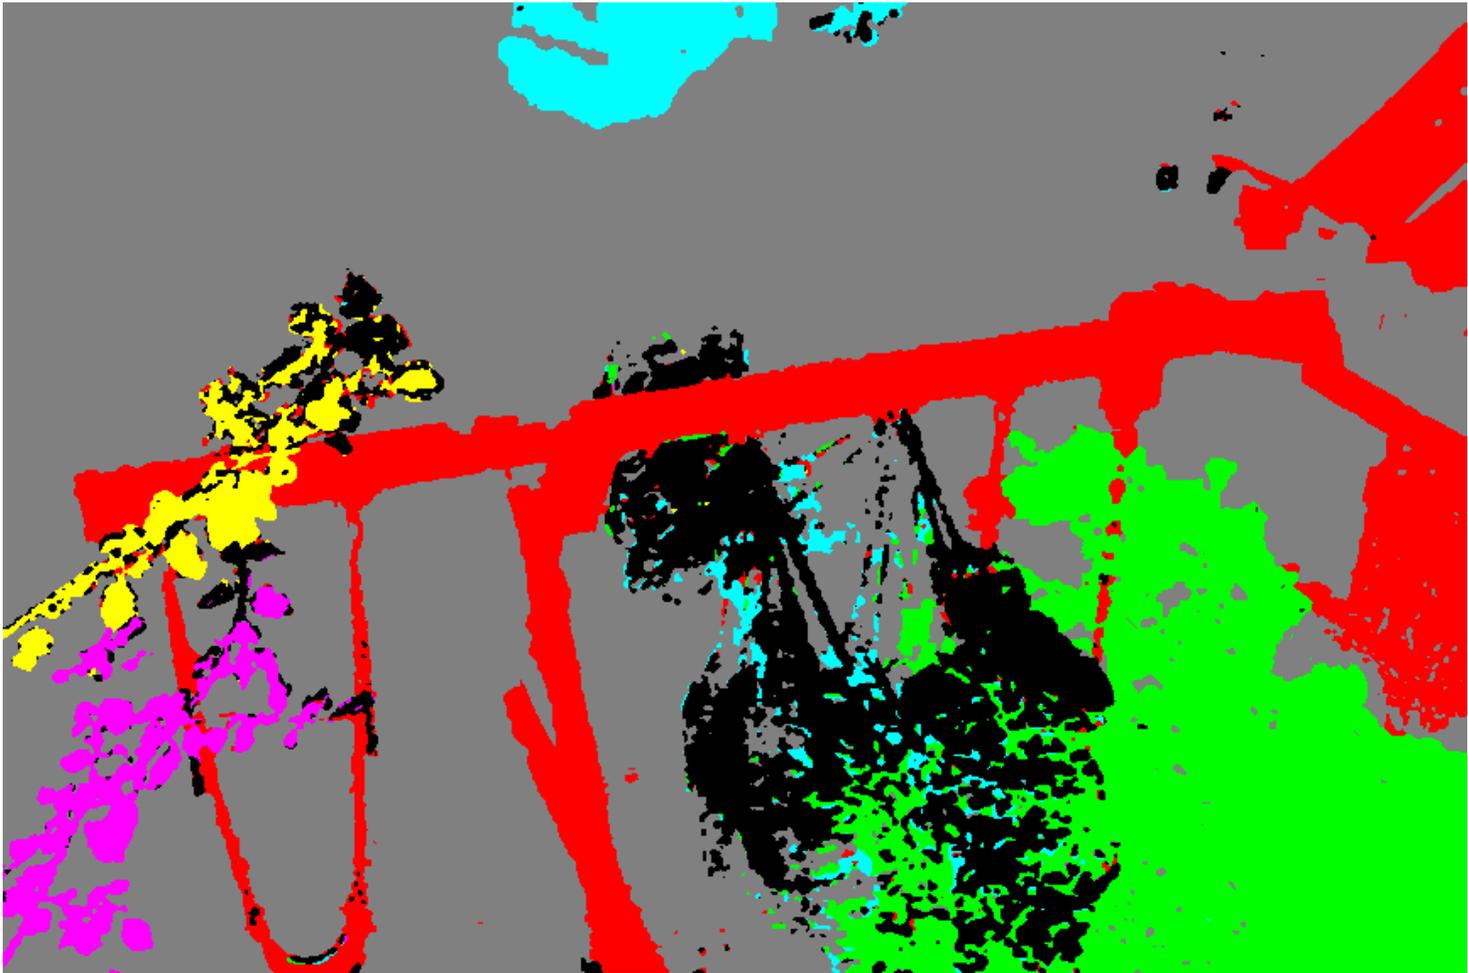
5 Segment flow into layers

- Assign each pixel to a motion cluster layer, using four cues:
 - Motion likelihood
 - Color likelihood
 - Spatial connectivity
 - Temporal coherence
- Energy minimization using graph cuts

$$\begin{aligned} L = \operatorname{argmin} & - \sum_{(x,y)} \log P_M(I_t(x,y) | L(x,y)) \\ & - \rho \sum_{(x,y)} \log P_C(I_t(x,y) | L(x,y)) \\ & + \gamma \sum_{(x,y)} \sum_{(i,j) \in N_{(x,y)}} V(I_t(x,y), I_t(x+i, y+j), L(x,y), L(x+i, y+j)) \end{aligned}$$

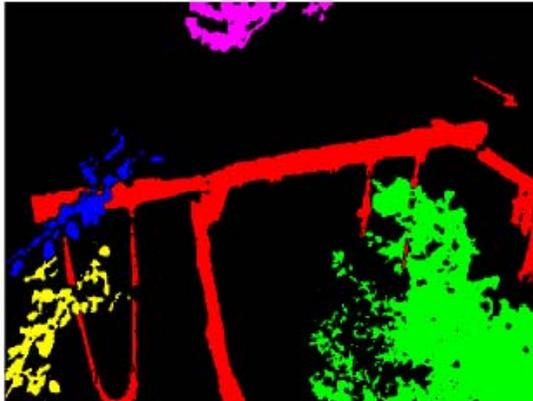
Motion segmentation results

Note we have 2 special layers: the background layer (gray), and the outlier layer (black).



6, 7 Magnification, texture fill-in and rendering

- Amplify the motion of the selected layers by warping the reference image pixels accordingly.
- Render unselected layers without magnification.
- Fill-in holes revealed in background layer using Efron-Leung texture synthesis
- Directly pass-through pixel values of the outlier layer.



(d) Layers of related motion and appearance



(e) Motion magnified, showing holes



(f) Final motion-magnified output

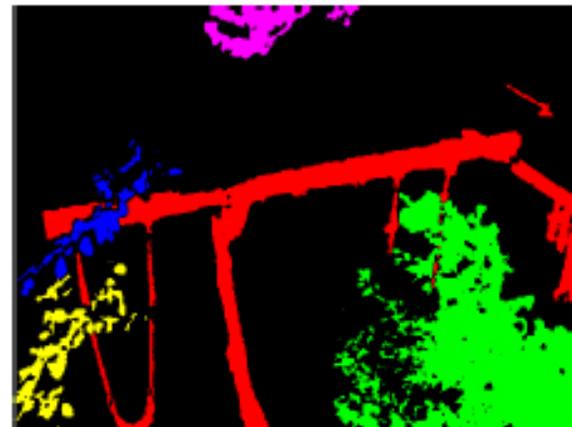
Summary of motion magnification steps



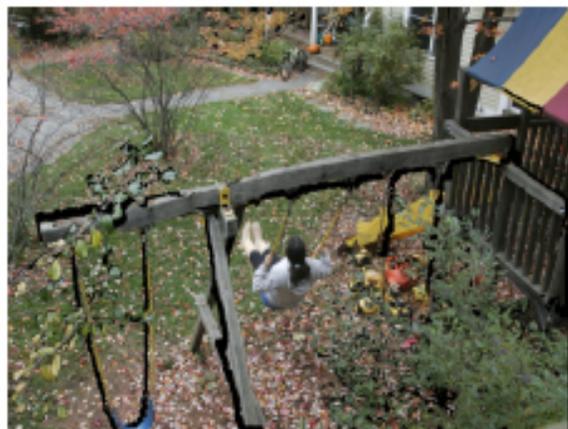
(a) Registered input frame



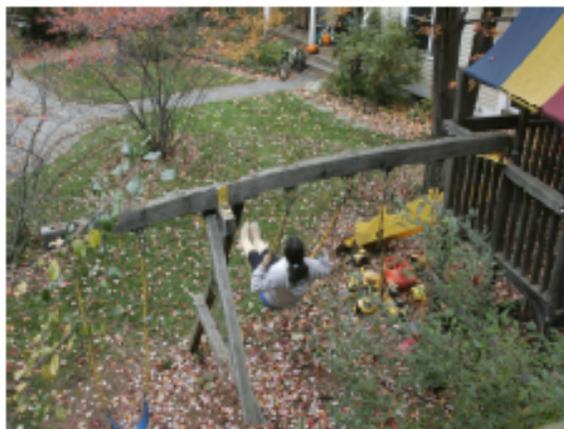
(b) Clustered trajectories of tracked features



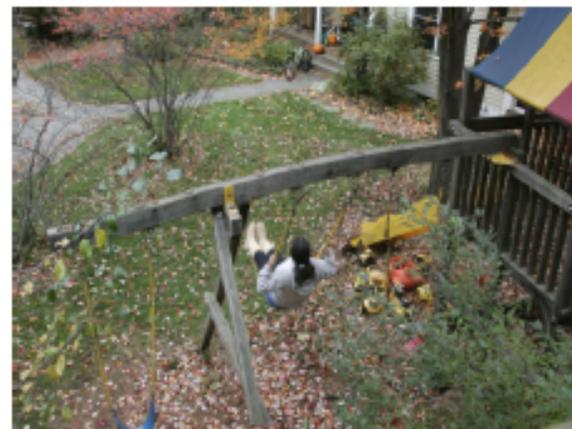
(c) Layers of related motion and appearance



(d) Motion magnified, showing holes



(e) After texture in-painting to fill holes



(f) After user's modifications to segmentation map in (c)

Results

- Demo

Layered representation

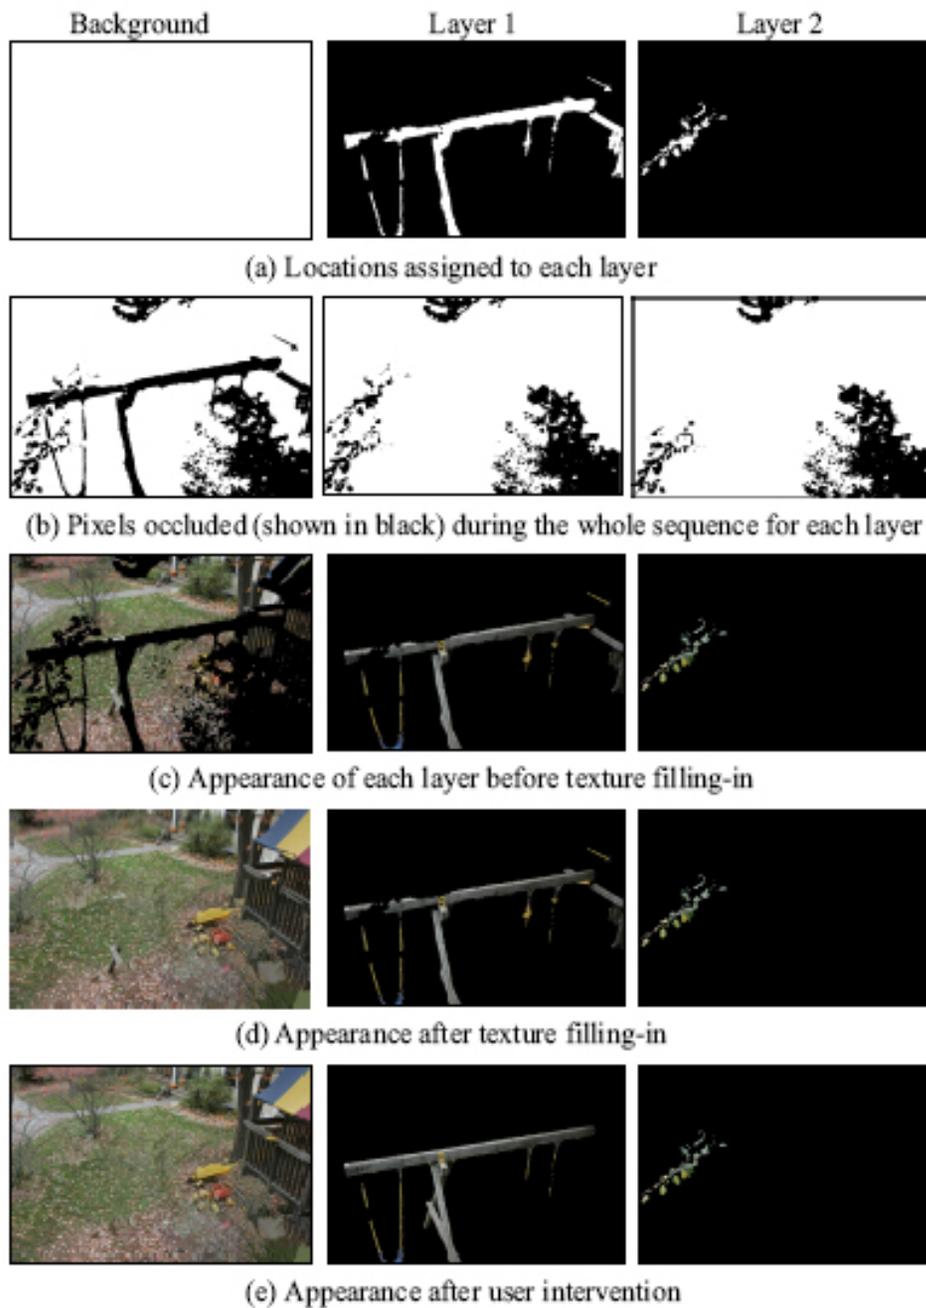


Figure 5: Layered representation for the swing sequence. Only the background and two layers (out of six) are shown.

Motion Magnification

Paper 0420

Is the signal really in the video? (yes)

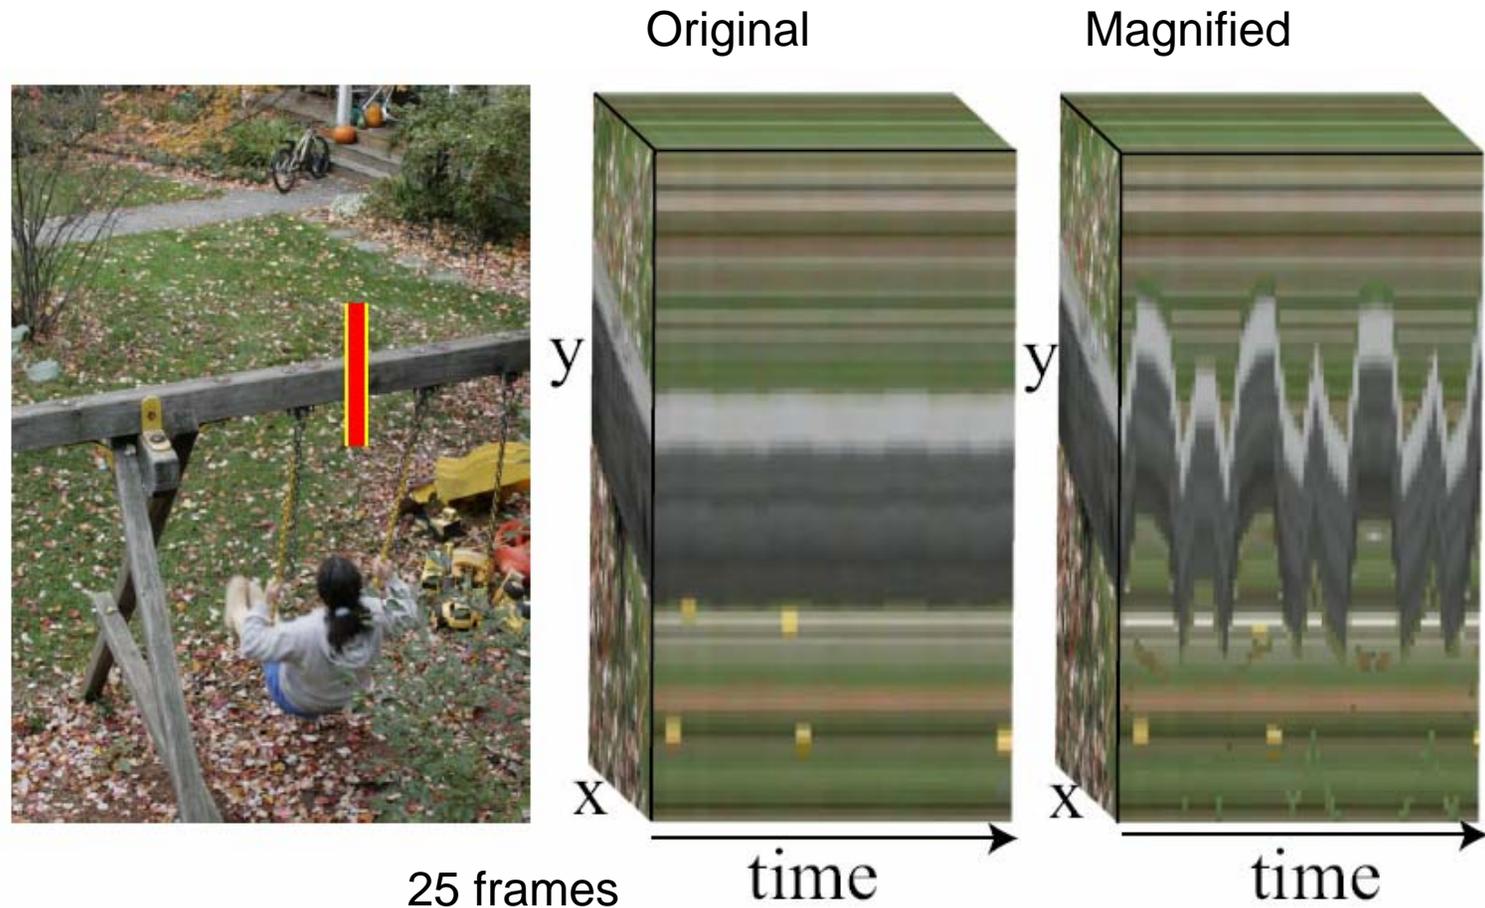


Figure 6: Section of the x, y, t volume from the original sequence (left) and the sequence after motion magnification. The detail shows a vertical section of the beam. The volume illustrates the am-

Swingset details

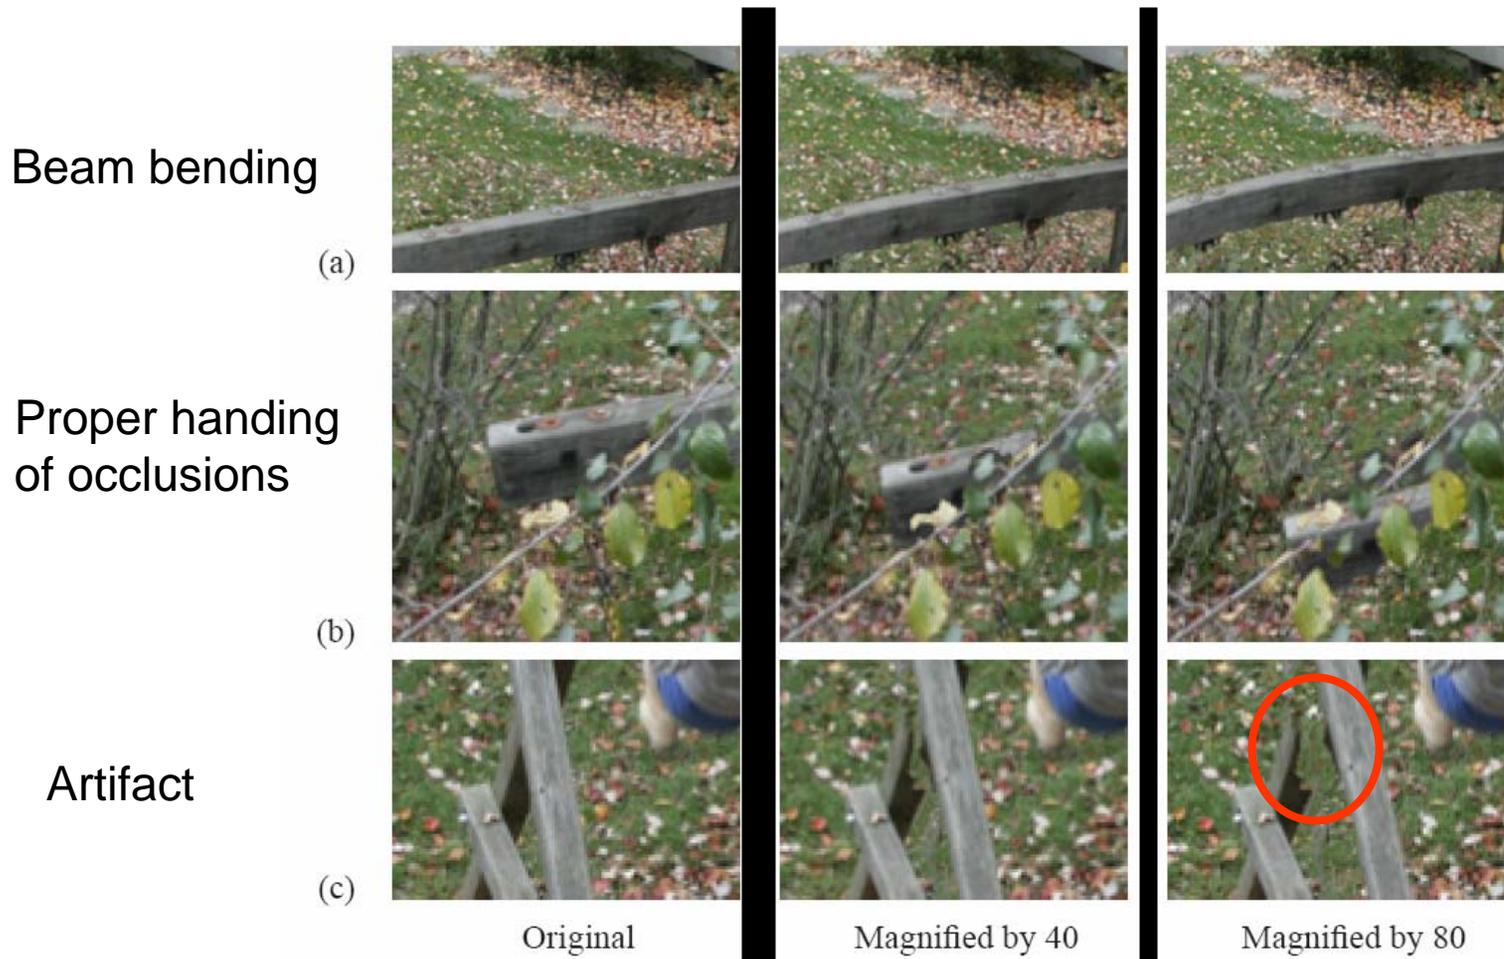


Figure 7: Details from frames of original and motion magnified swingset sequence, showing (a) beam curvature, (b) proper handling of occlusions, and (c) an artifact from imperfect segmentation.

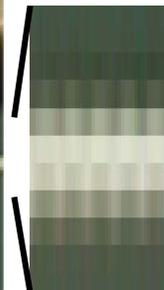
Bookshelf

44 frames

480x640 pixels



Original



time

Magnified



time

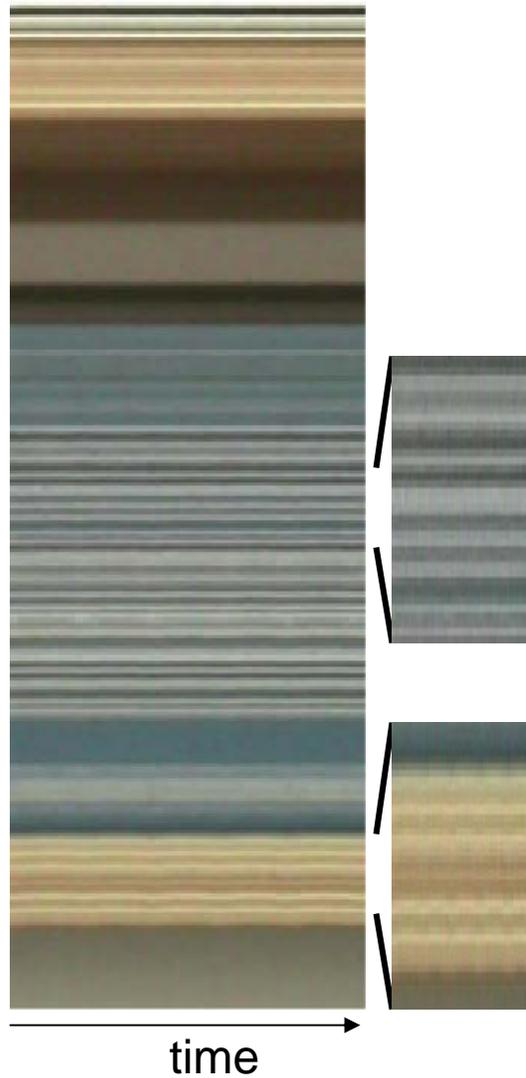
Bookshelf deforms less, further from force



44 frames

480x640 pixels

Original



Magnified



Outtakes from imperfect segmentations

Breathing Mike

Original sequence...



Breathing Mike

Feature points



2 clusters



4 clusters



8 clusters



Breathing Mike

Sequence after magnification...



Standing Mike

Sequence after magnification...

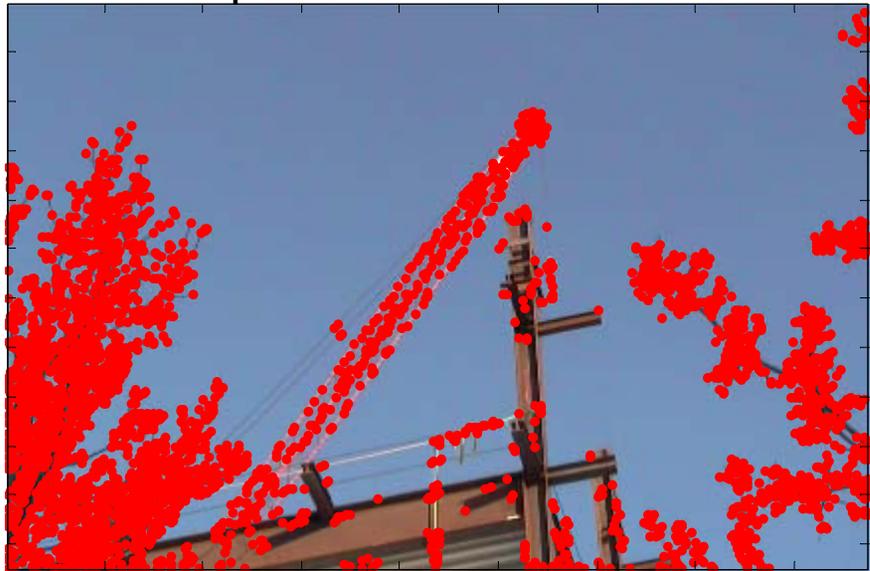


Crane

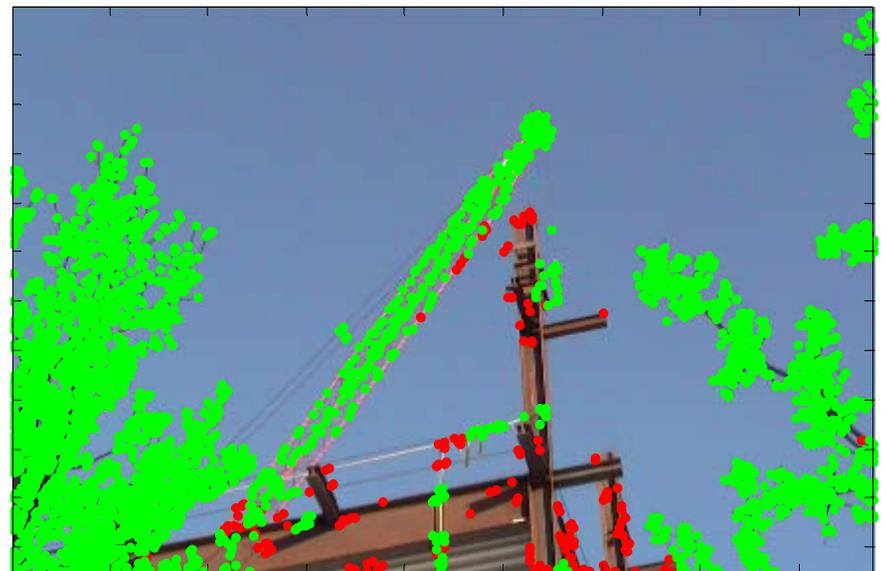


Crane

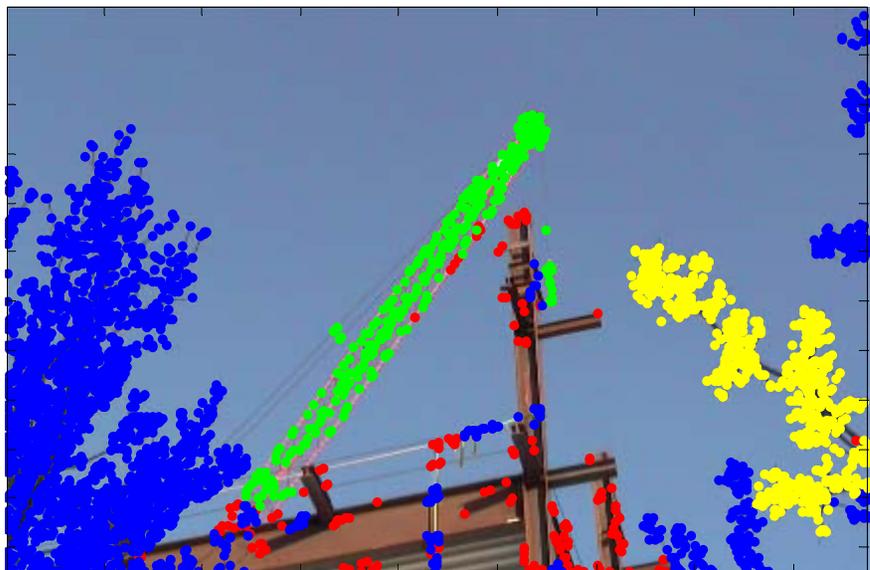
Feature points



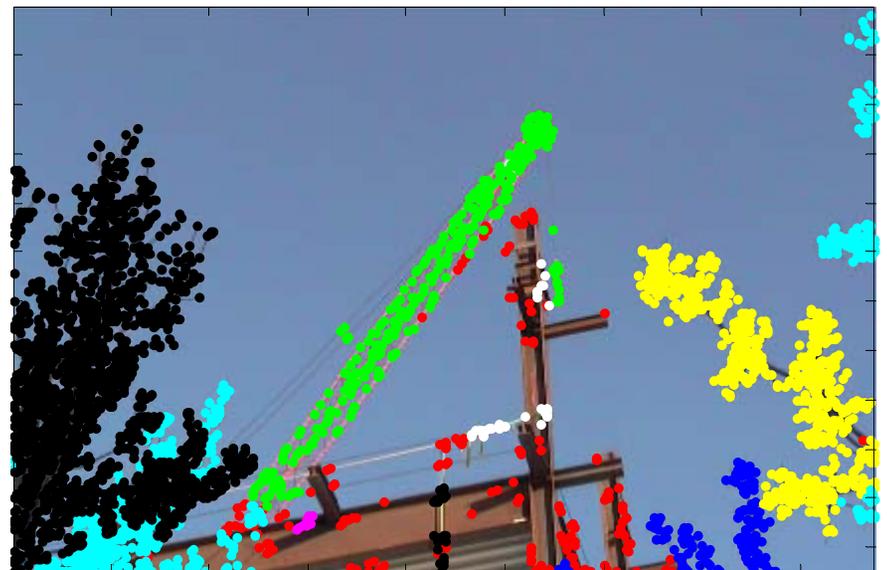
2 clusters



4 clusters



8 clusters



Crane



Things can go horribly wrong sometimes...

What next

- Continue improving the motion segmentation.
 - Motion magnification is “segmentation artifact amplification”—a good test bed.
- Real applications
 - Videos of inner ear
 - Connections with mechanical engineering dept.
- Generalization: amplifying small differences in motion.
 - What’s up with Tiger Woods’ golf swing, anyway?