Particle Video: Long-Range Video Motion Estimation using Point Trajectories

Peter Sand
Seth Teller
MIT CSAIL

Long-Range Motion Estimation

Long-Range Motion Estimation

Long-Range Motion Estimation

Long-Range Motion Estimation

Long-Range Motion Estimation

Long-Range Motion Estimation

Long-Range Motion Estimation

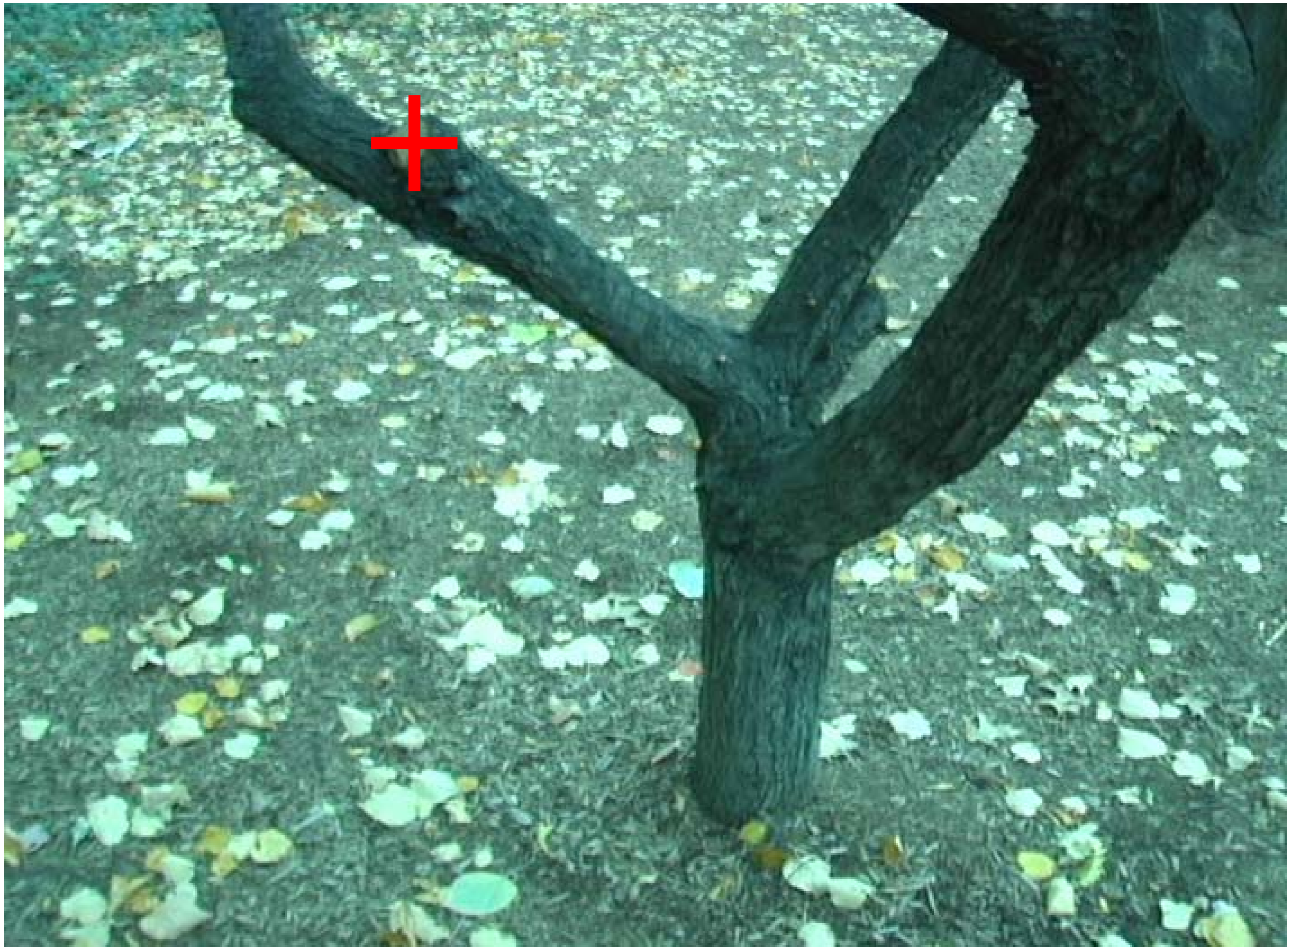

Long-Range Motion Estimation

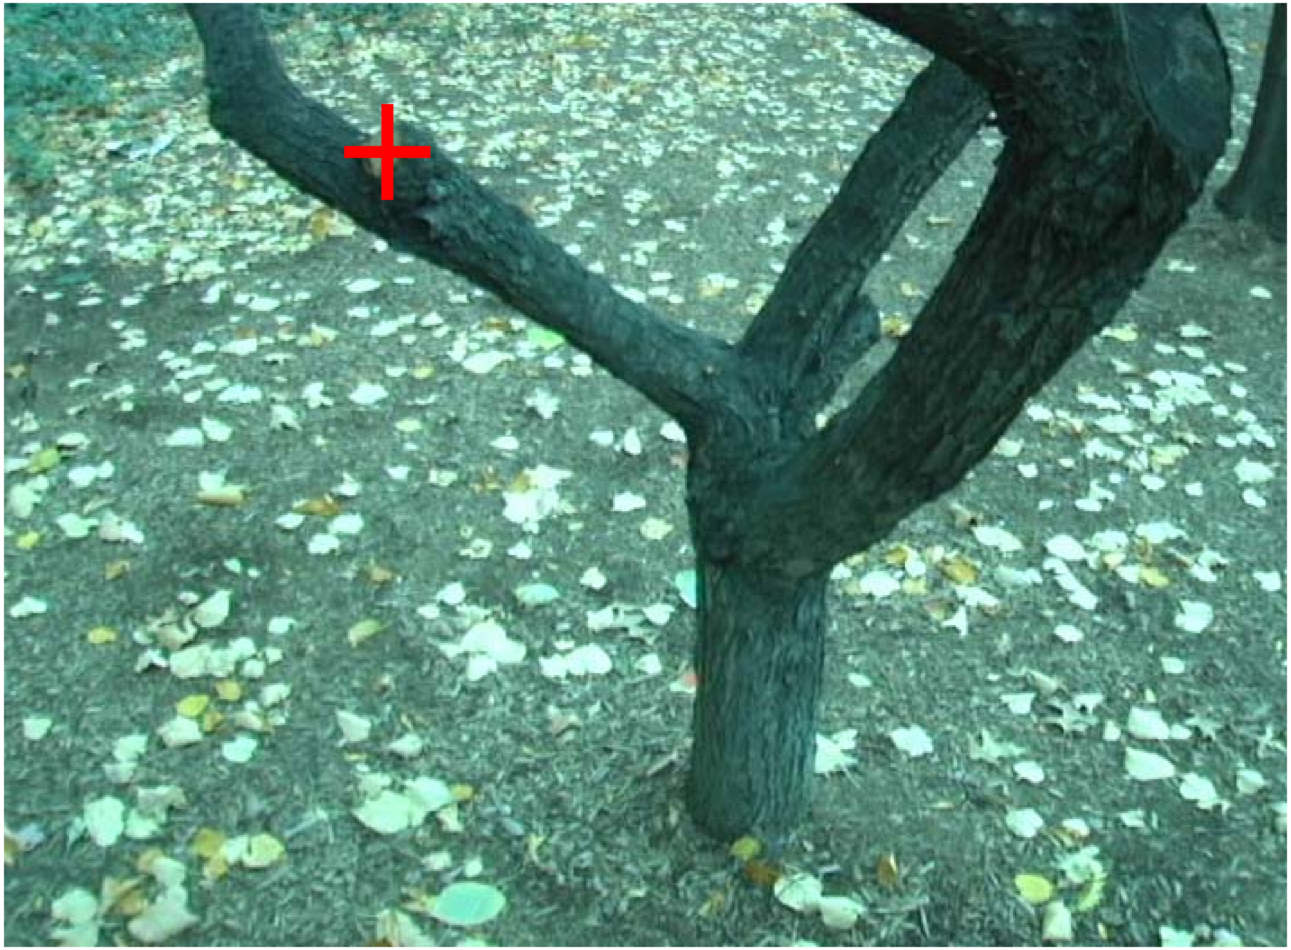

Long-Range Motion Estimation

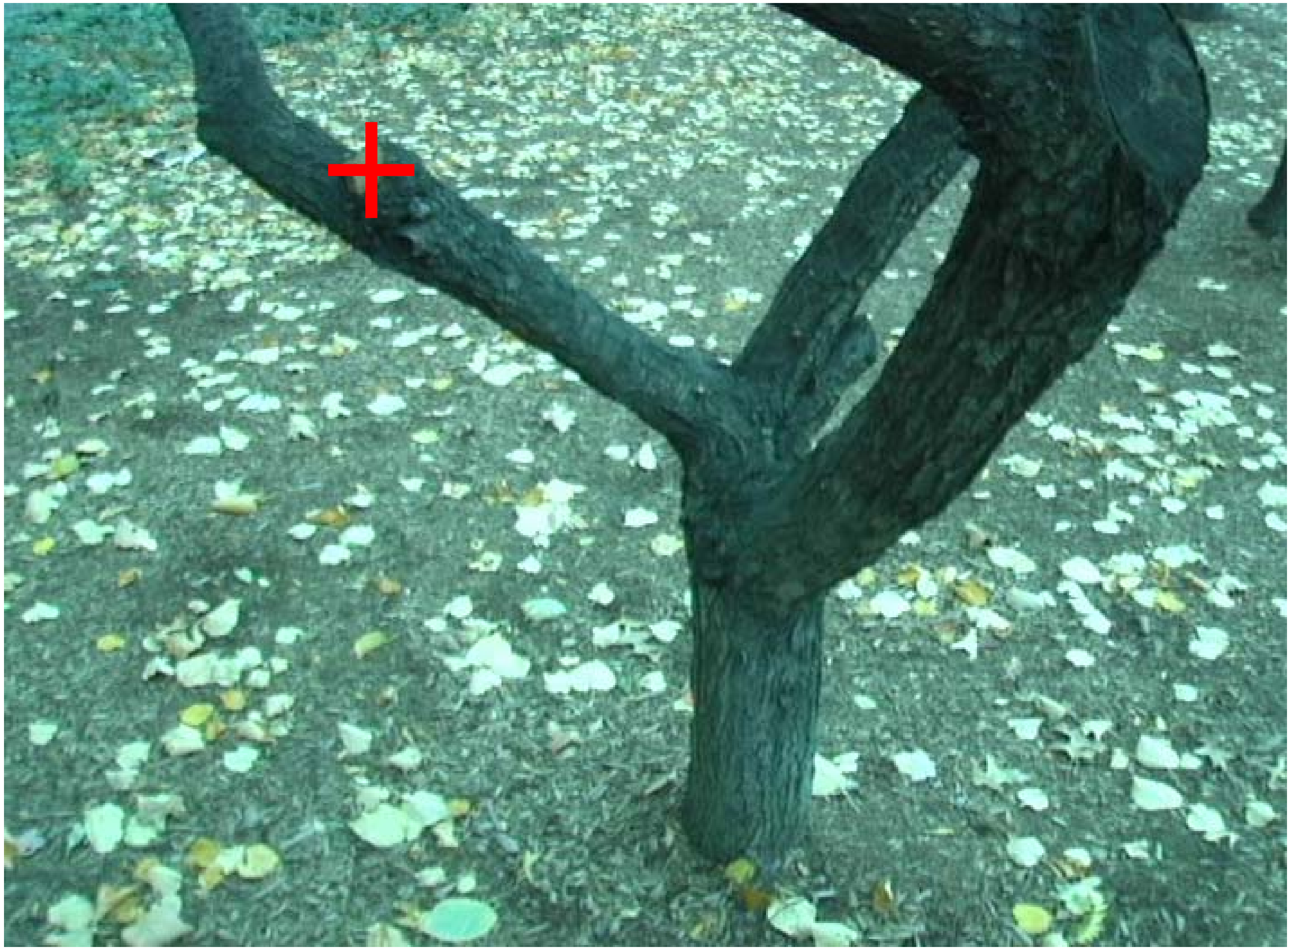

Long-Range Motion Estimation

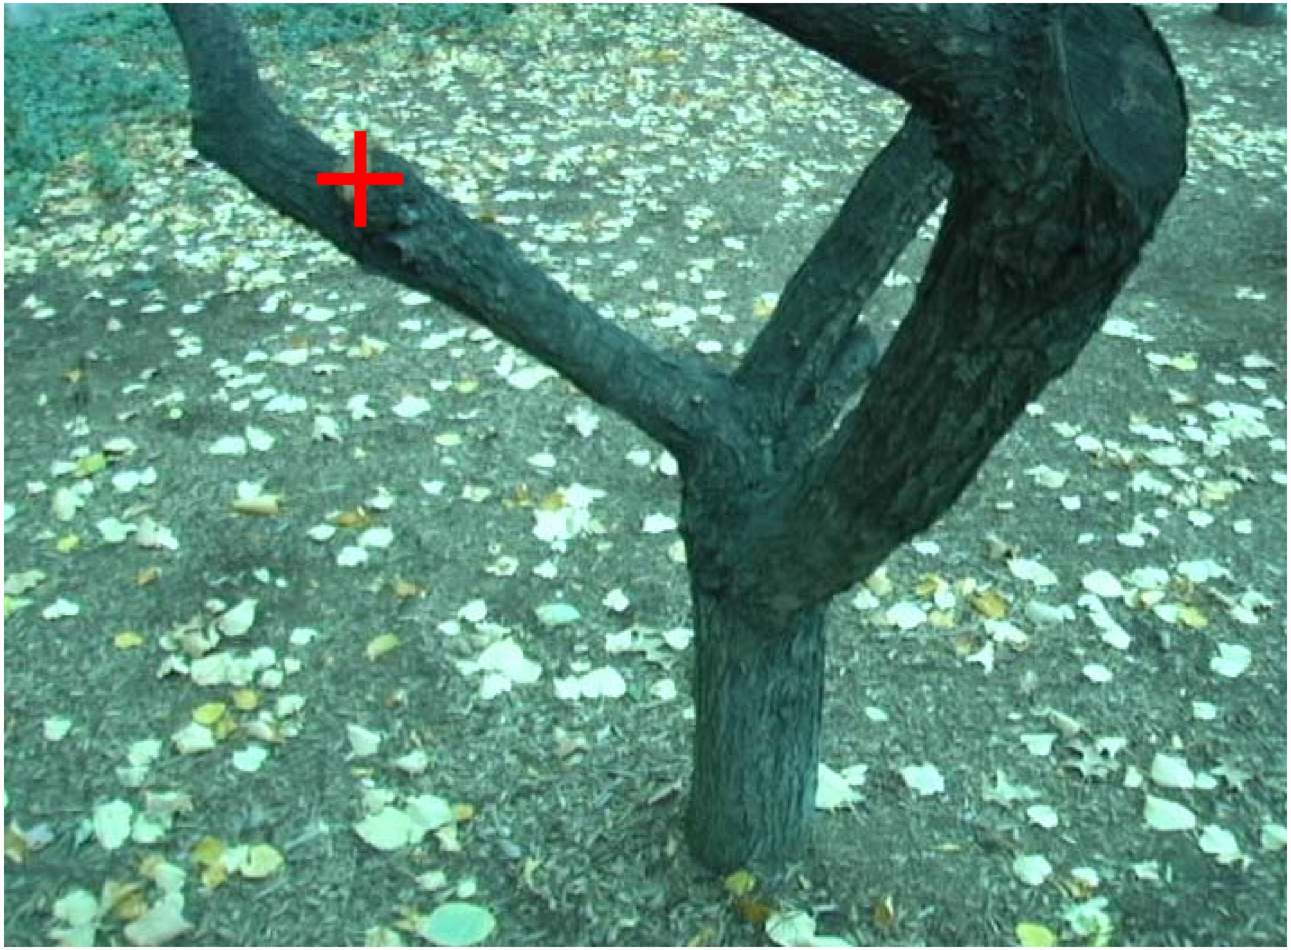

Long-Range Motion Estimation

Long-Range Motion Estimation

Long-Range Motion Estimation

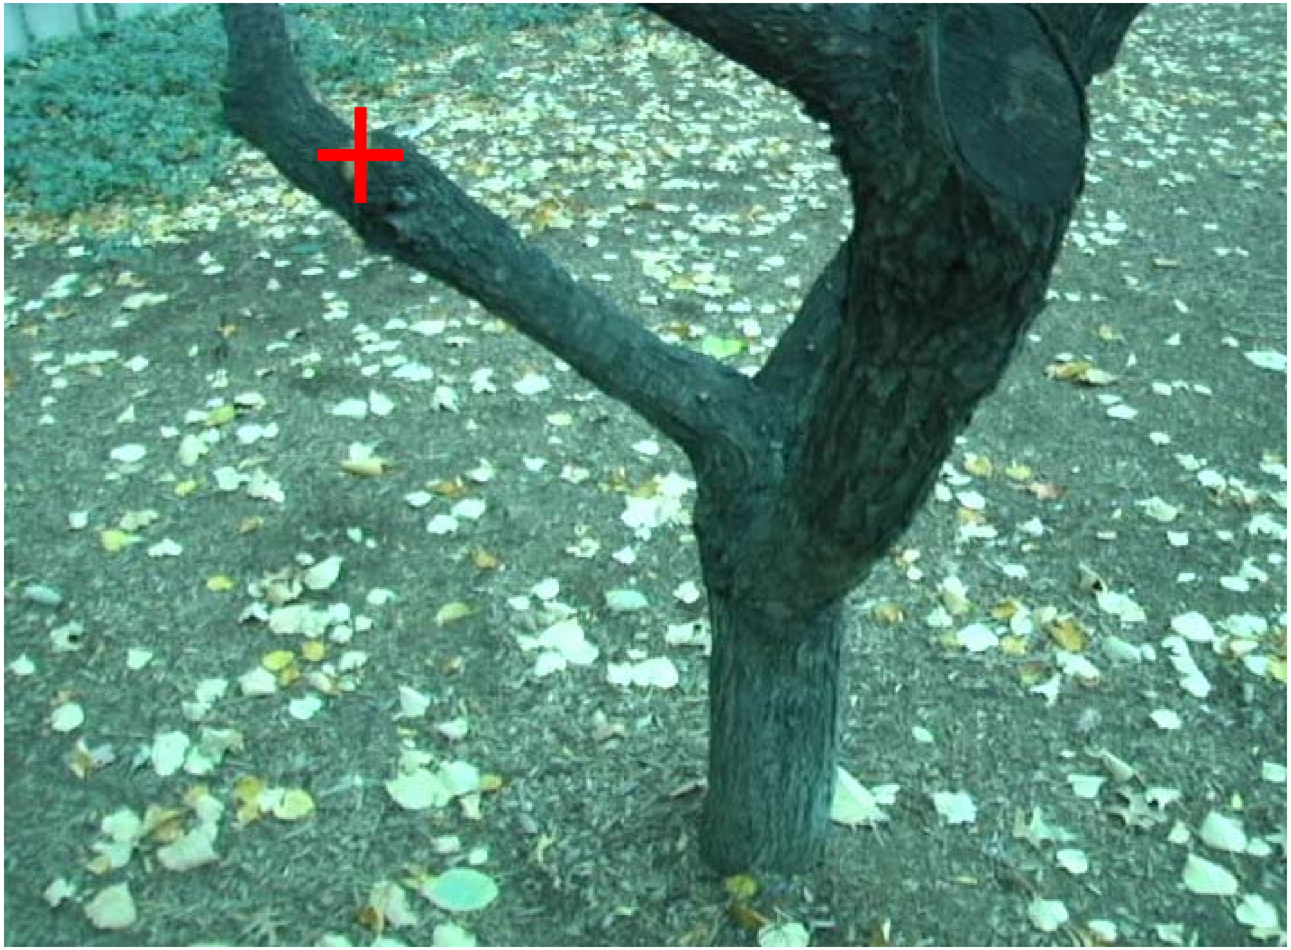

Long-Range Motion Estimation

Long-Range Motion Estimation

Long-Range Motion Estimation

Applications

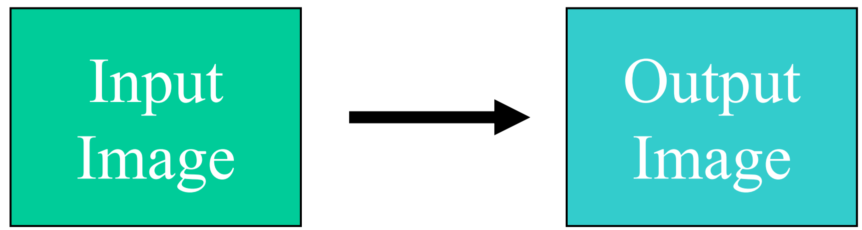

Photoshop, etc.

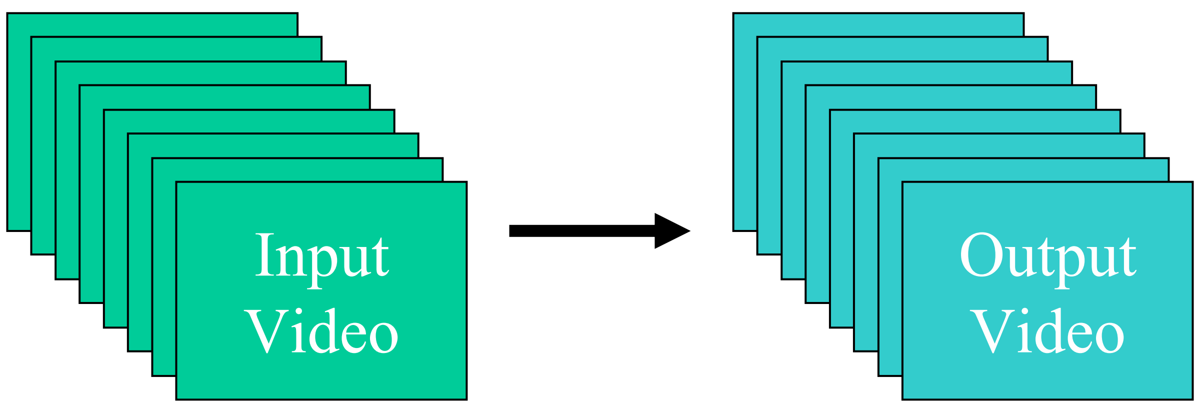

Particle-based editor

Applications

- Super-resolution
- Noise removal
- High dynamic range video
- Image filtering
- Video segmentation
- Matting / rotoscoping
- Object removal

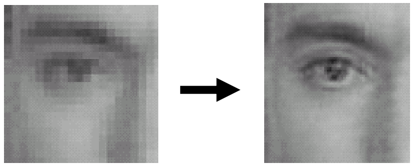

[Capel and Zisserman 2001]

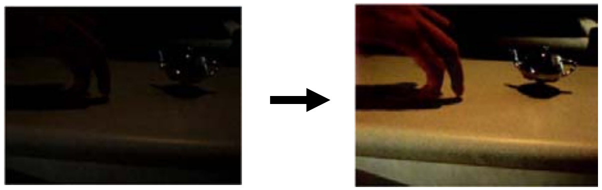

[Bennett and McMillan 2005]

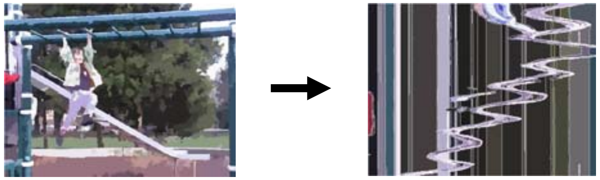

[Jue Wang *et al*. 2004]

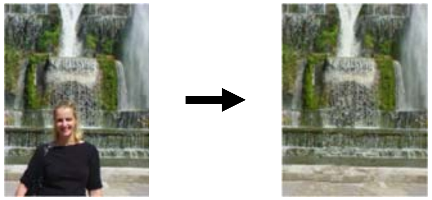

[Criminisi *et al*. 2003]

Applications

Long-range motion estimation is a step toward a larger goal: video decomposition.

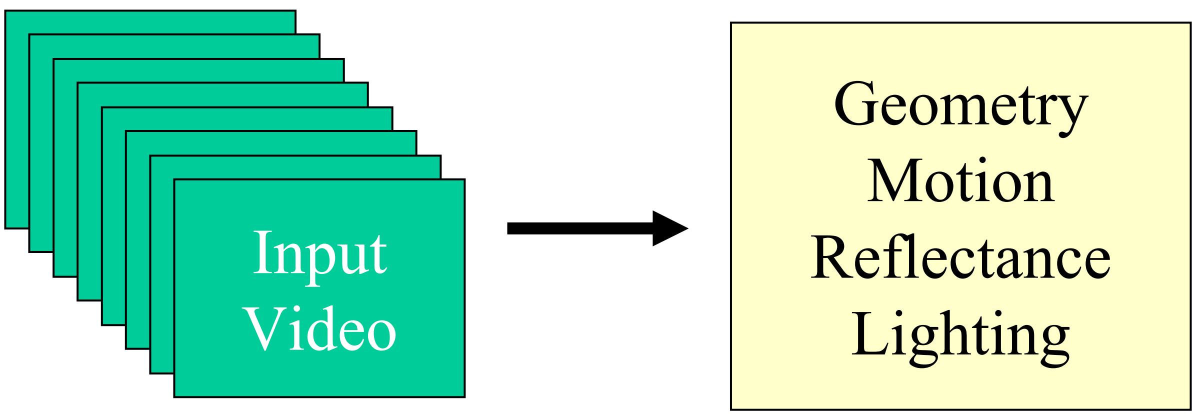

Design Goals

Design Goals

Design Goals

Related Work

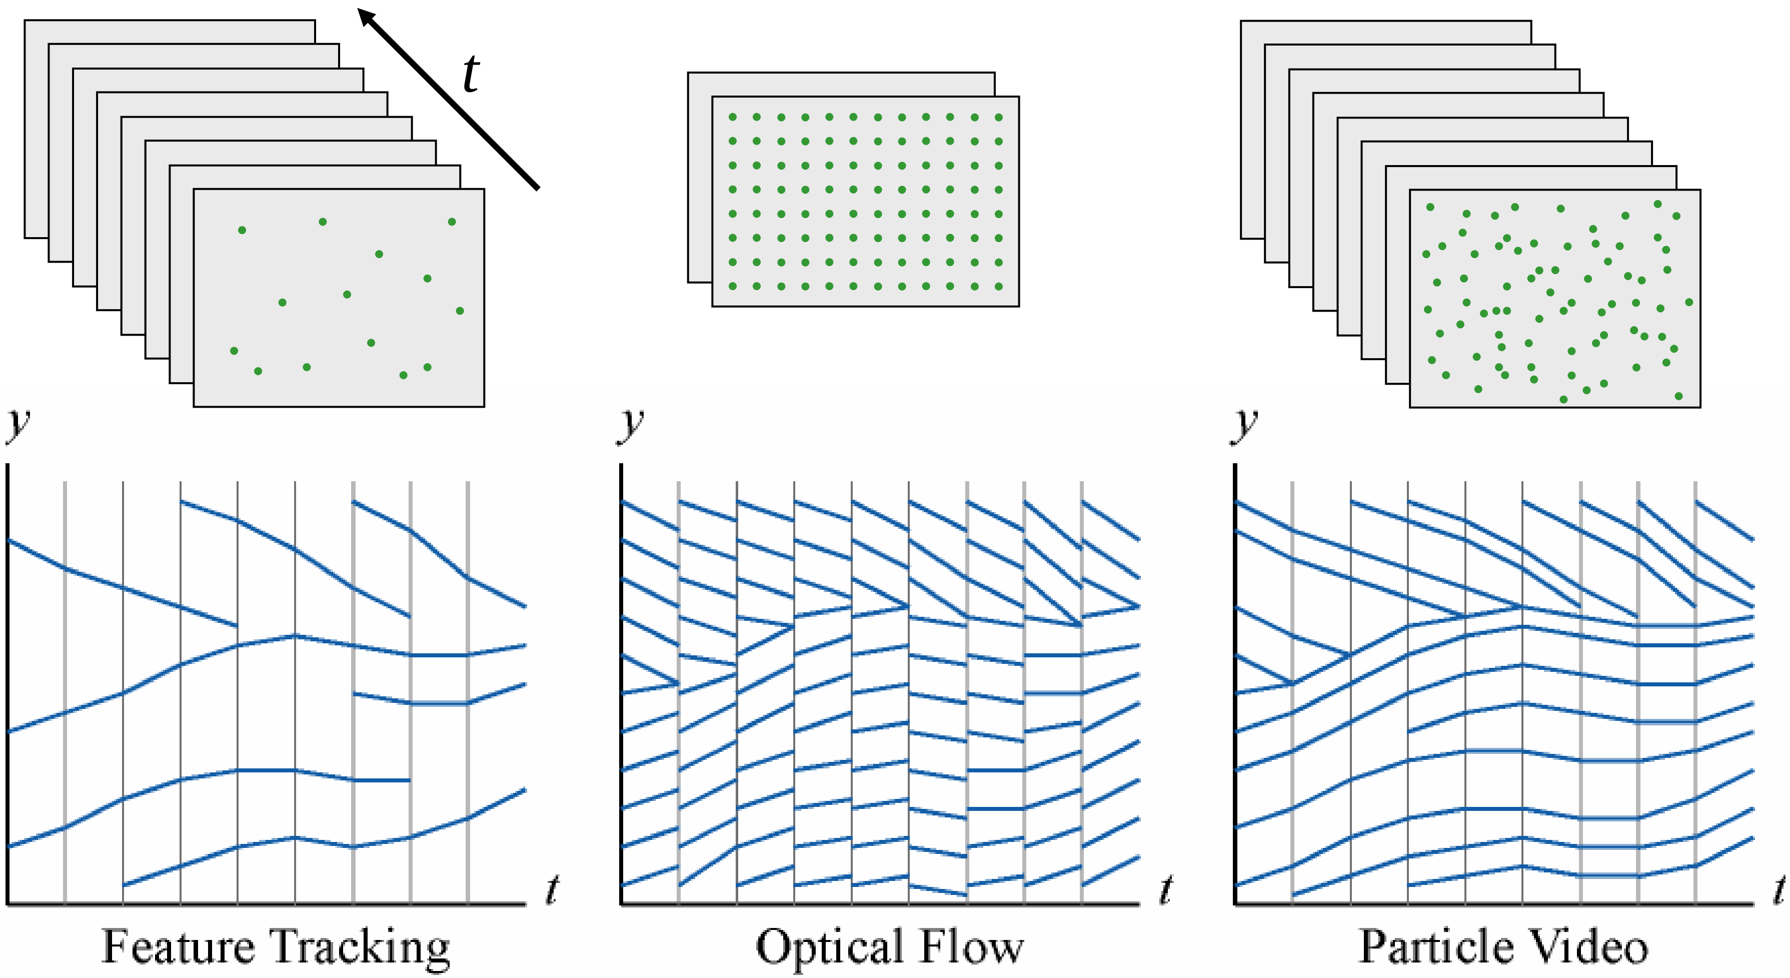

Related Work: Optical Flow

Temporal smoothness assumption

[Black and Anandan 1991, Chin *et al*. 1994, Elad and Feuer 1998, Shi and Malik 1998]

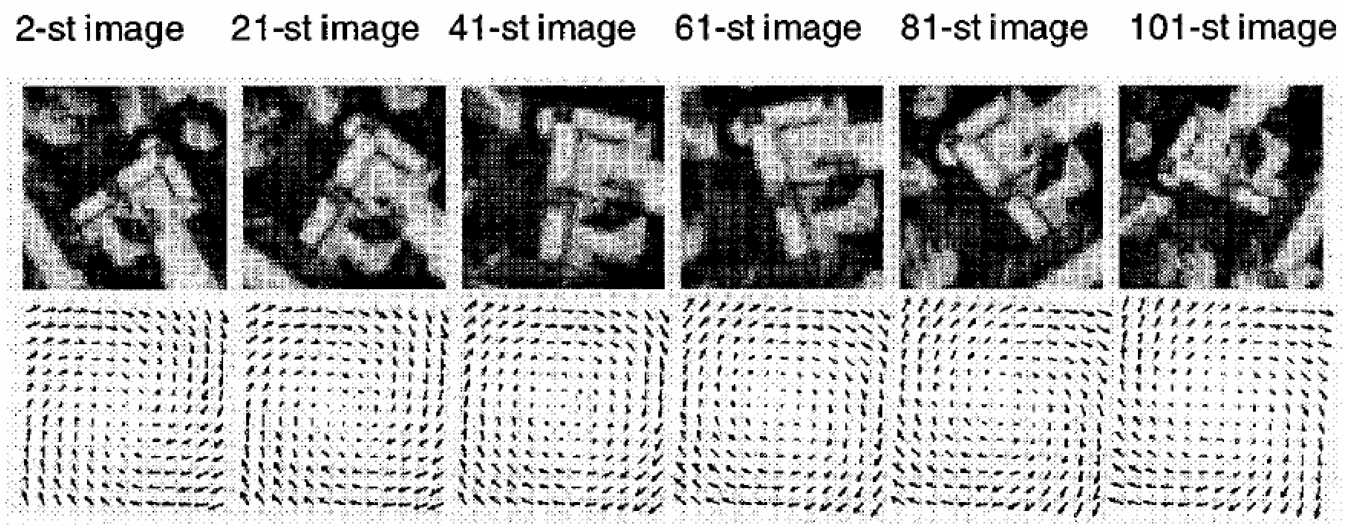

[Elad and Feuer 1998]

Related Work: Optical Flow

Rank-based optical flow

[Irani 1999, Brand 2001]

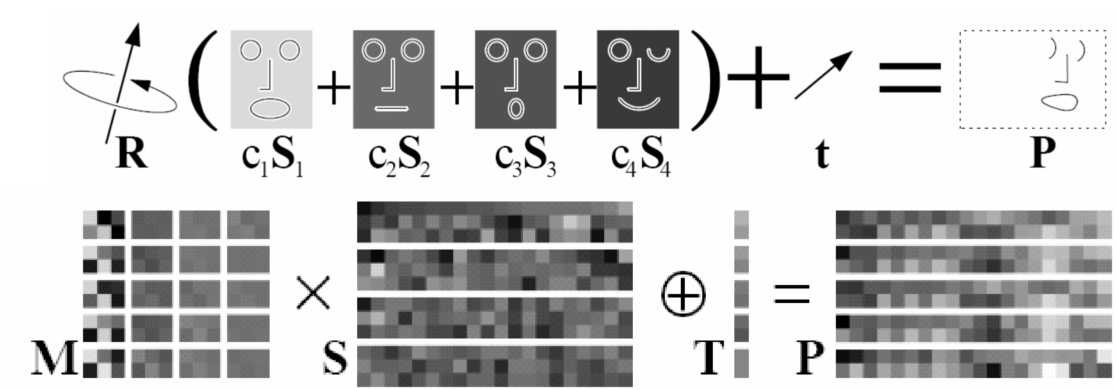

[Brand 2001]

Related Work: Optical Flow

Occlusion labeling:

- Pixel dissimilarity
  [Silva and Santos-Victor 2001, Xiao *et al*. 2006]
- Forward / backward mismatch
  [Alvarez *et al*. 2002]

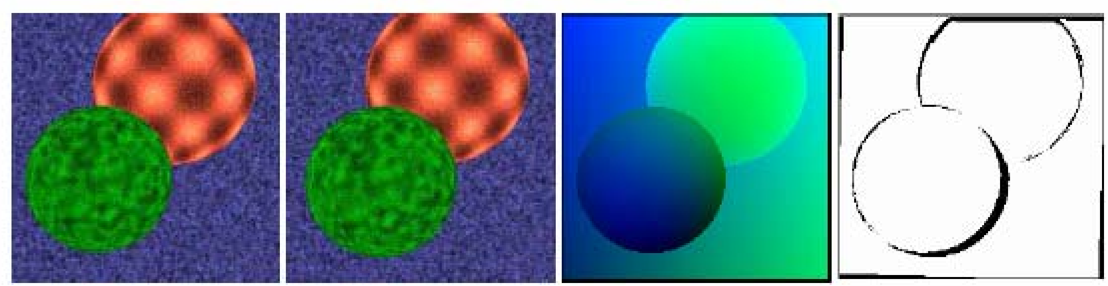

Fig. 2. *One of the synthetic scenes used in the experiments. Left two images: input; middle right: ground truth optical flow; right: ground truth occlusions*

[Stretcha *et al*. 2004]

Particle Approach

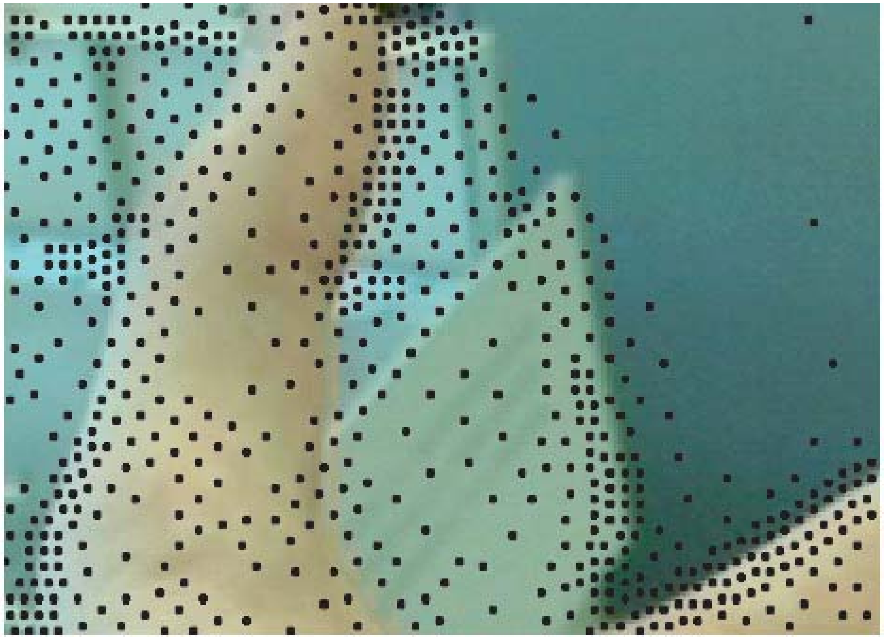

- Particles are small
- Adaptive density
- Does not assume temporal motion smoothness

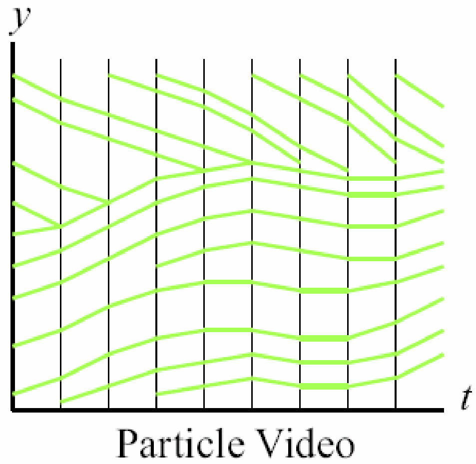

Particle Approach

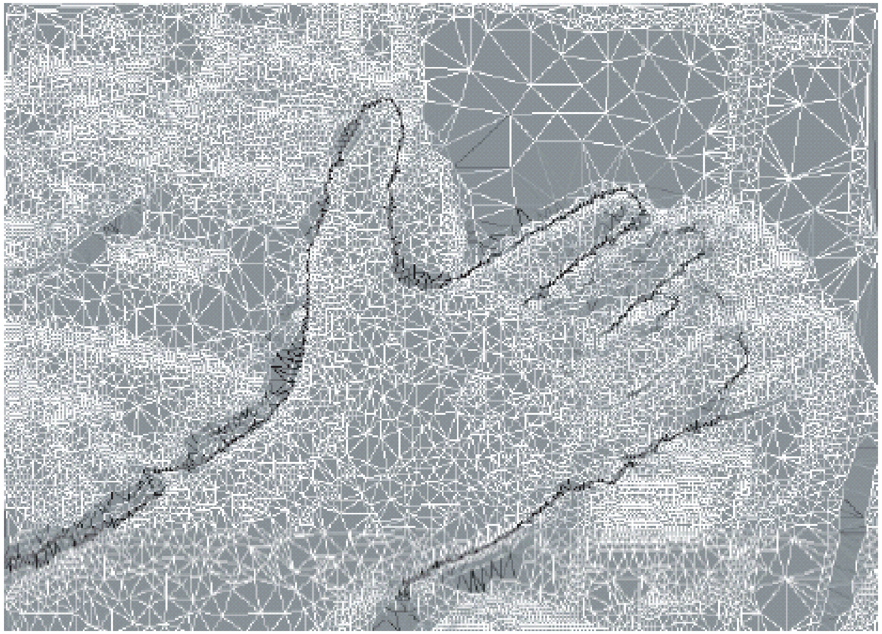

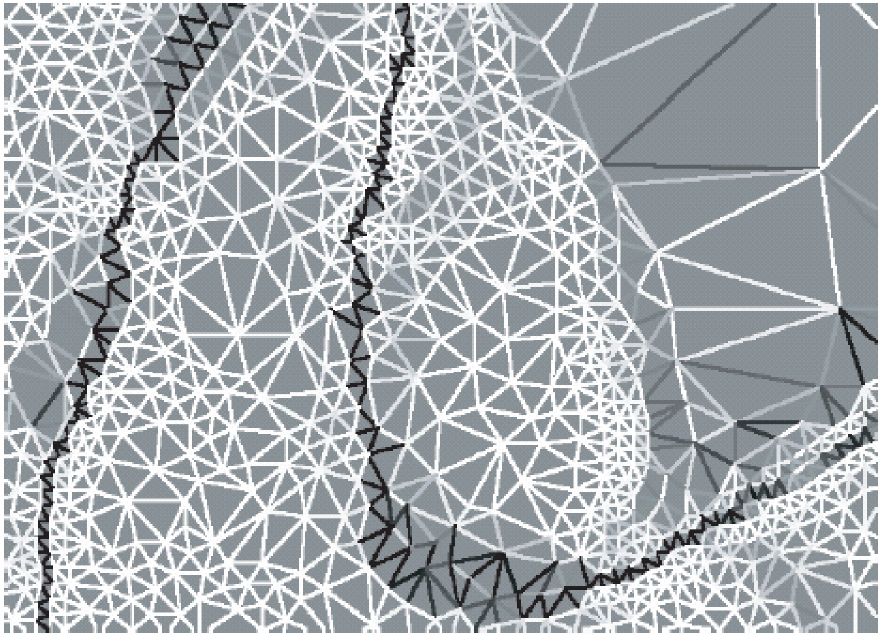

- Triangulation implicitly represents particle grouping
- Non-parametric
- Not layer-based
- No segmentation
- Not planar or rigid components

Optical Flow as Input

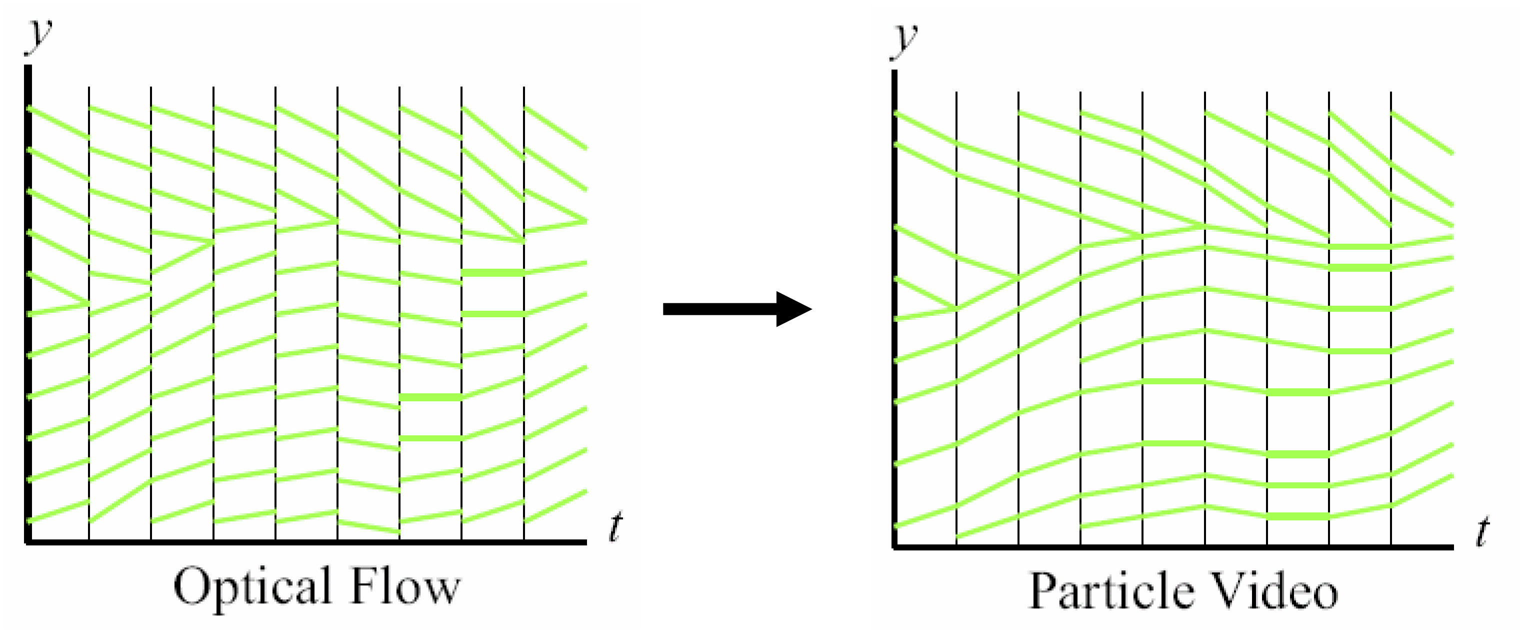

An Optical Flow Algorithm

At each resolution level:

- Variational flow update
  - Similar to [Brox *et al*. 2004]
- Label occluded regions
- Bilateral flow filter
  - Similar to [Xiao *et al*. 2006]

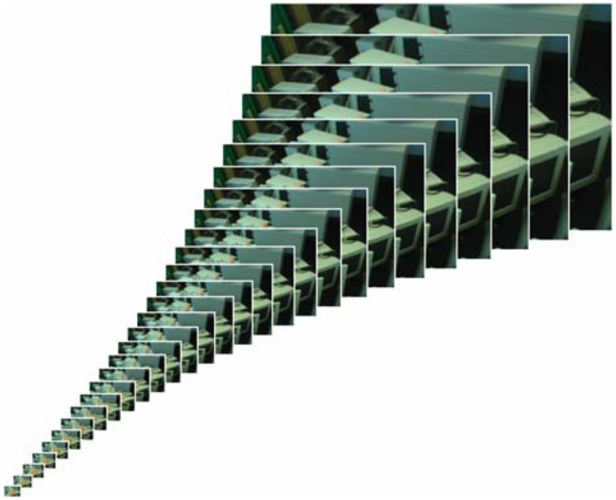

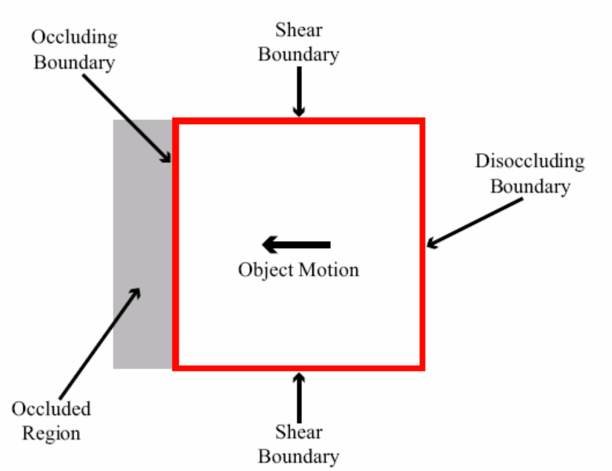

(see paper for more details)

Optical Flow Results

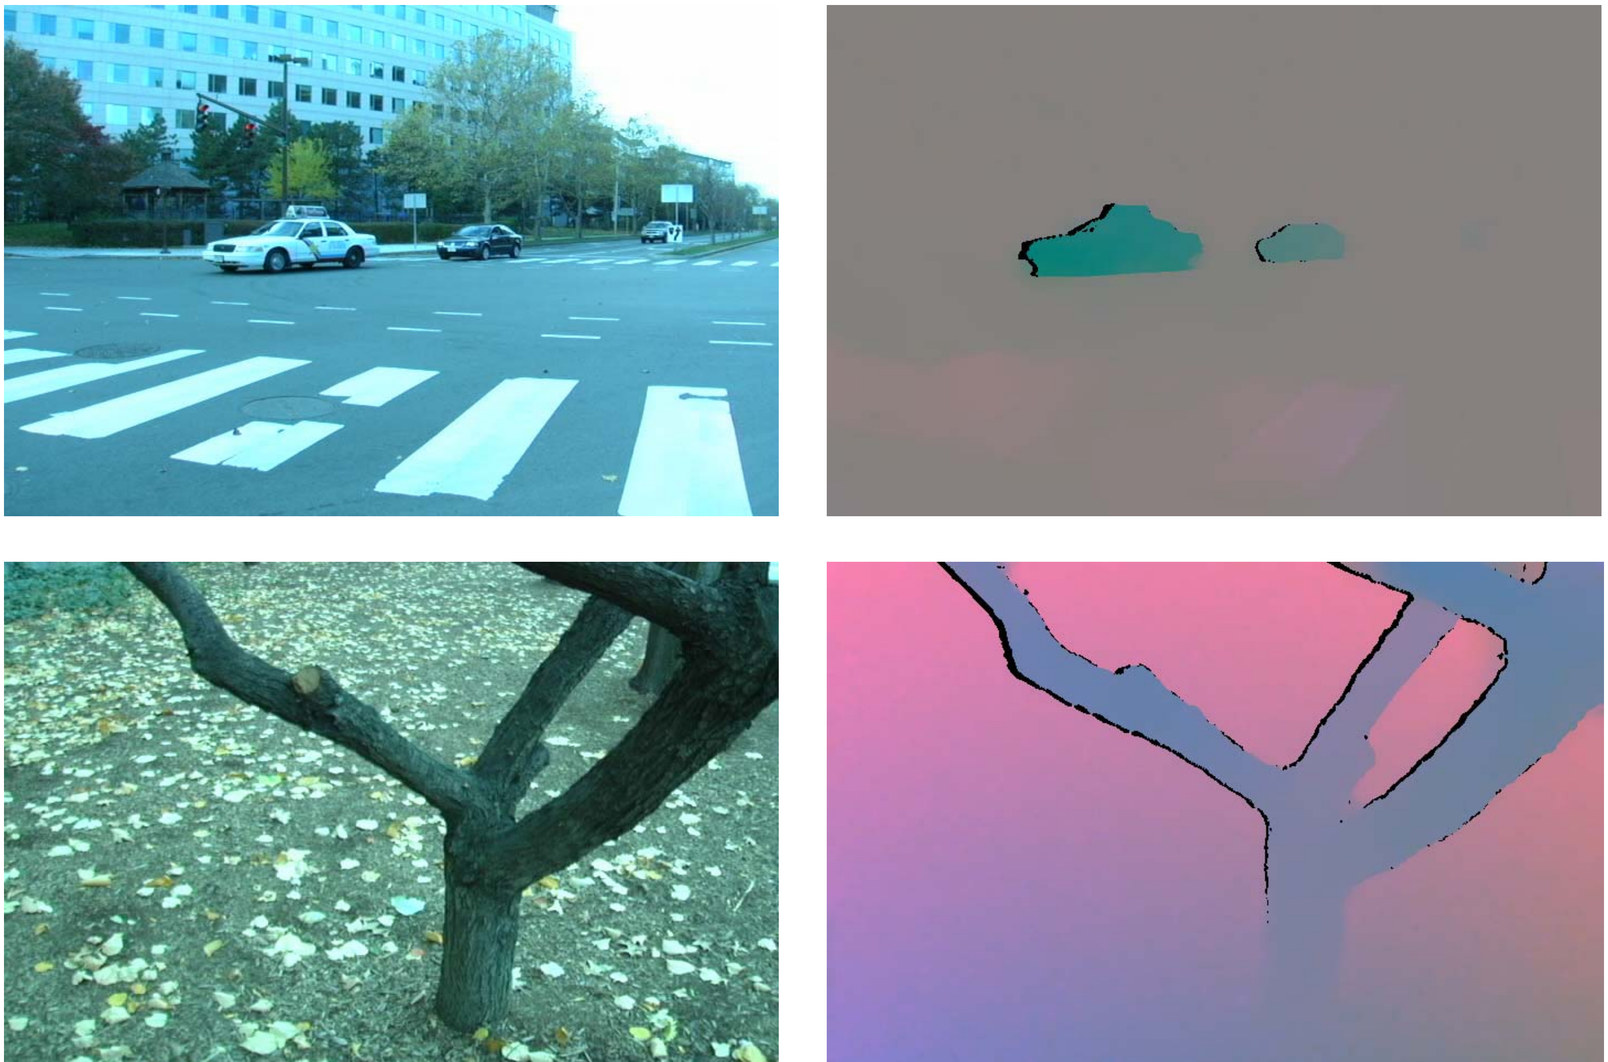

Particle Video Algorithm

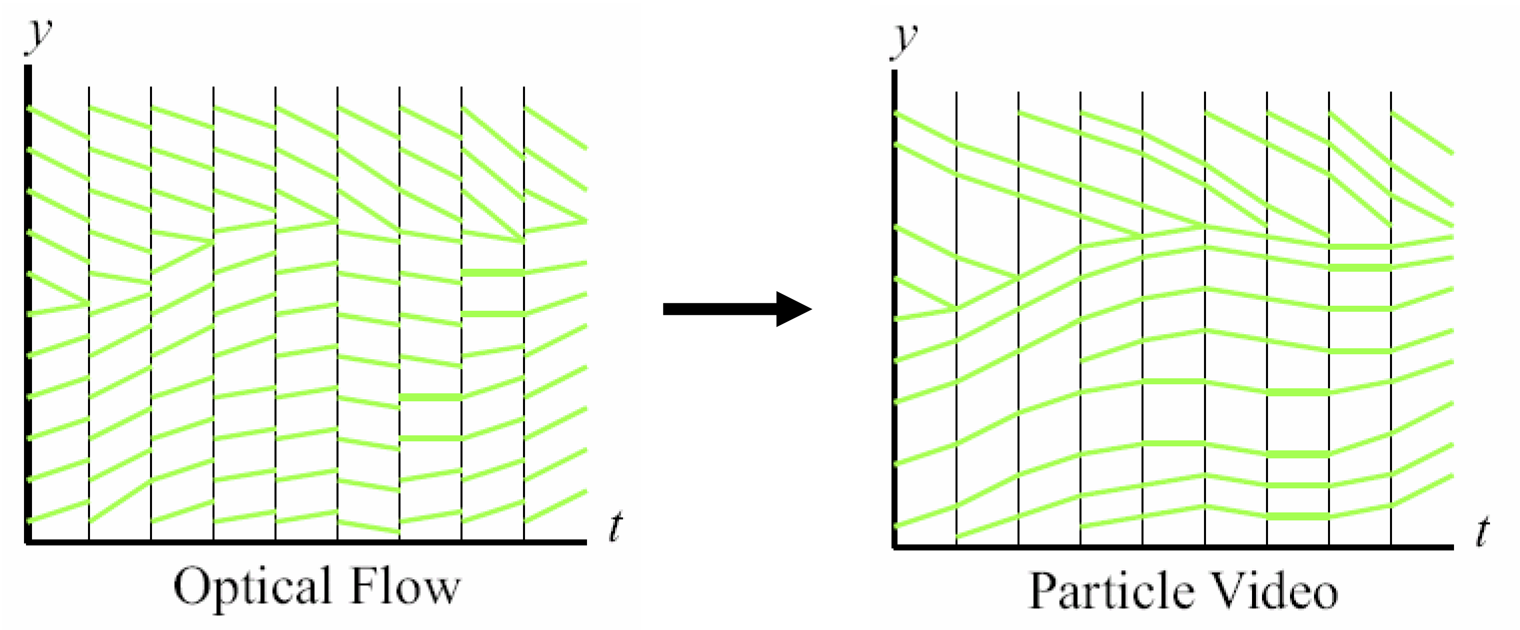

Particle Video Steps

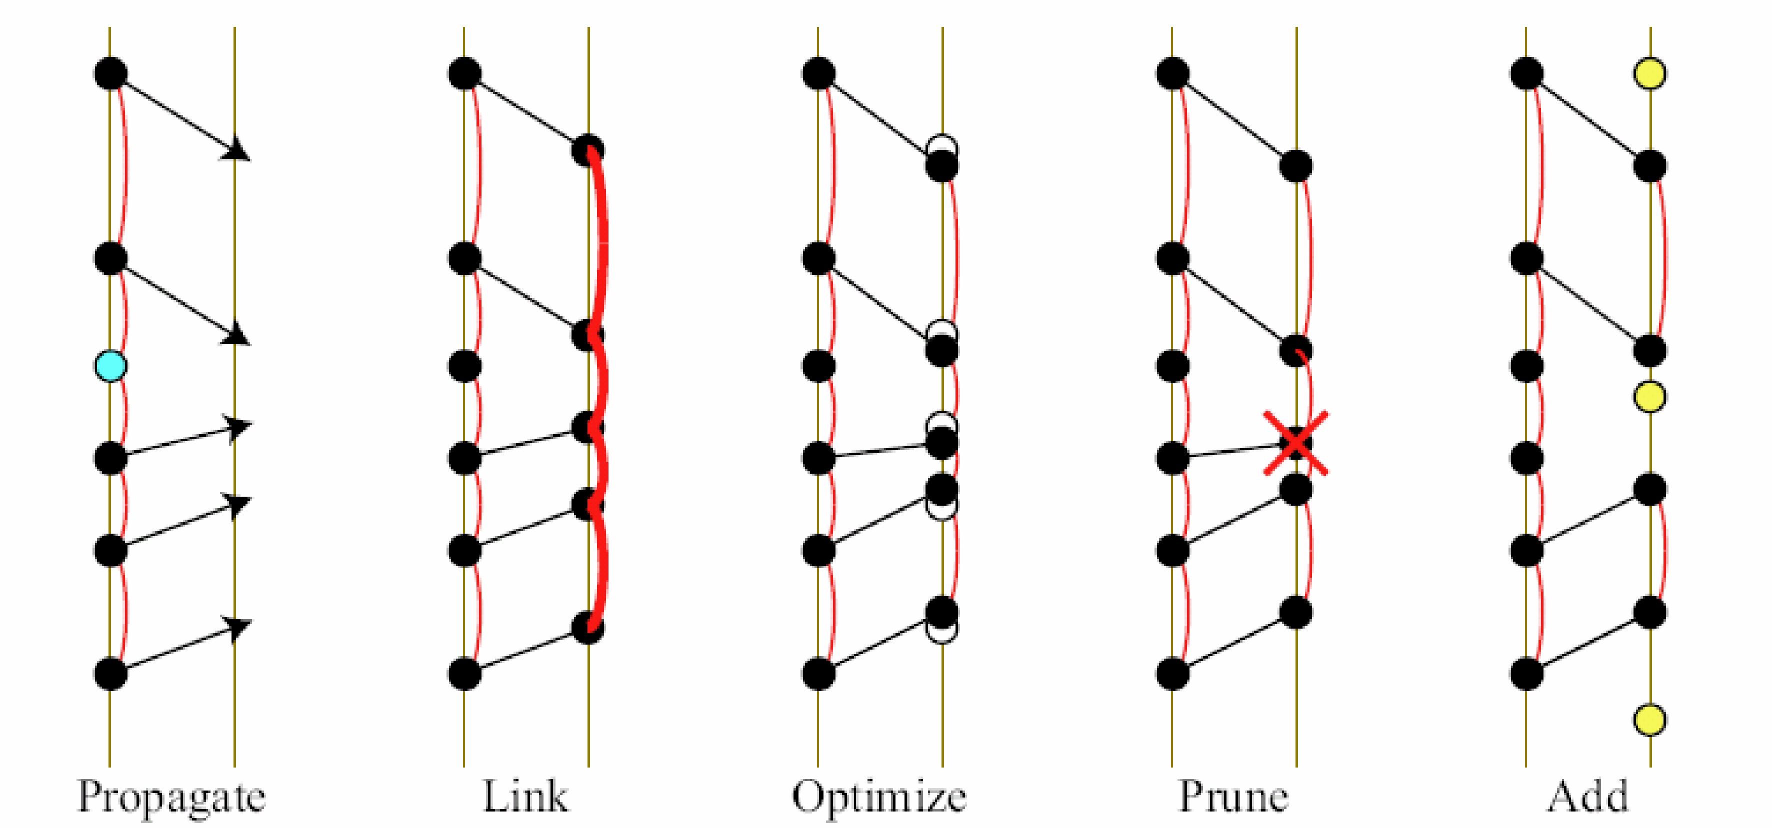

Particle Propagation

Forward propagation:

$$
\begin{aligned}
x_i(t) &= x_i(t-1) + u(x_i(t-1), y_i(t-1), t-1), \\
y_i(t) &= y_i(t-1) + v(x_i(t-1), y_i(t-1), t-1).
\end{aligned}
$$

Particles in occluded regions are not propagated.

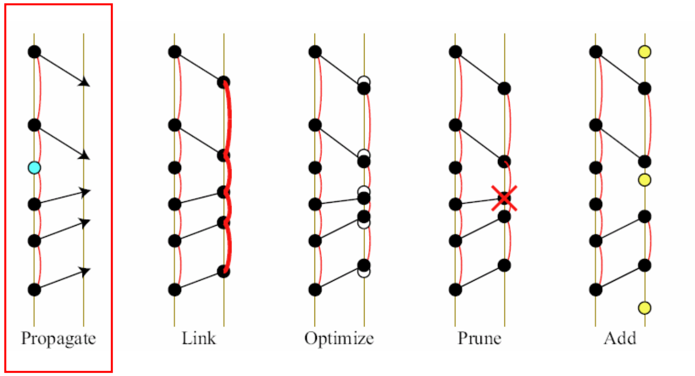

Particle Linking

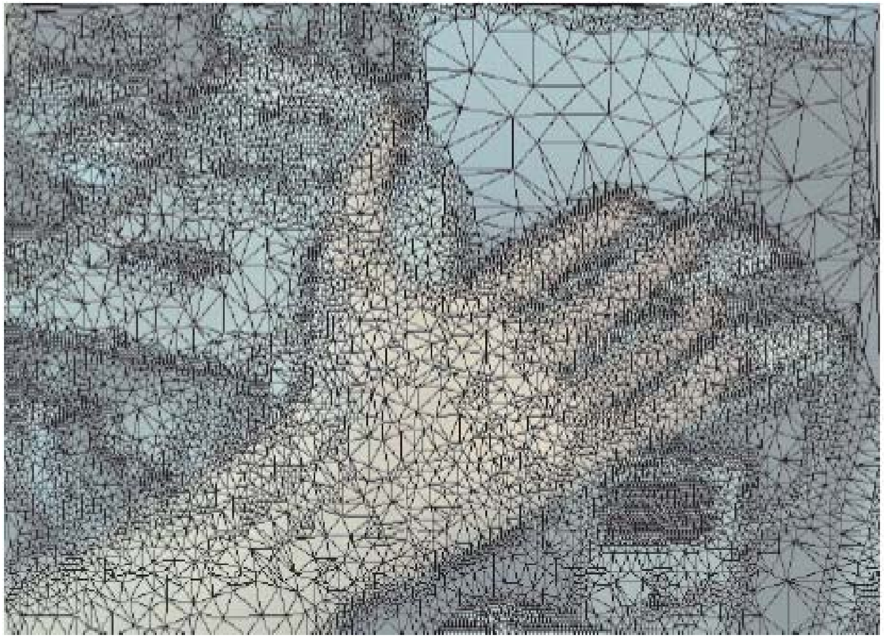

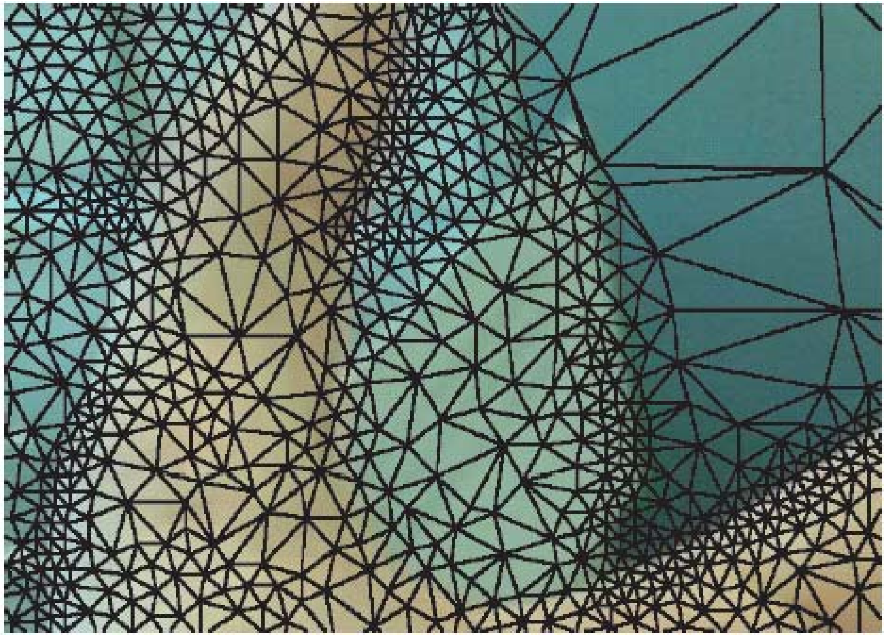

- Delaunay Triangulation [Lischinksi 1994]
- Create link if Delaunay edge exists in current frame or adjacent frame

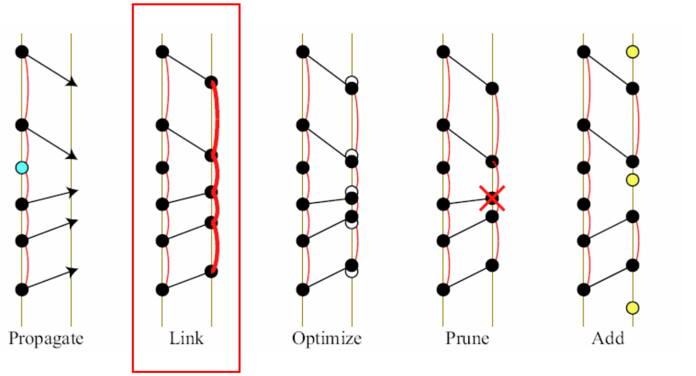

Particle Linking

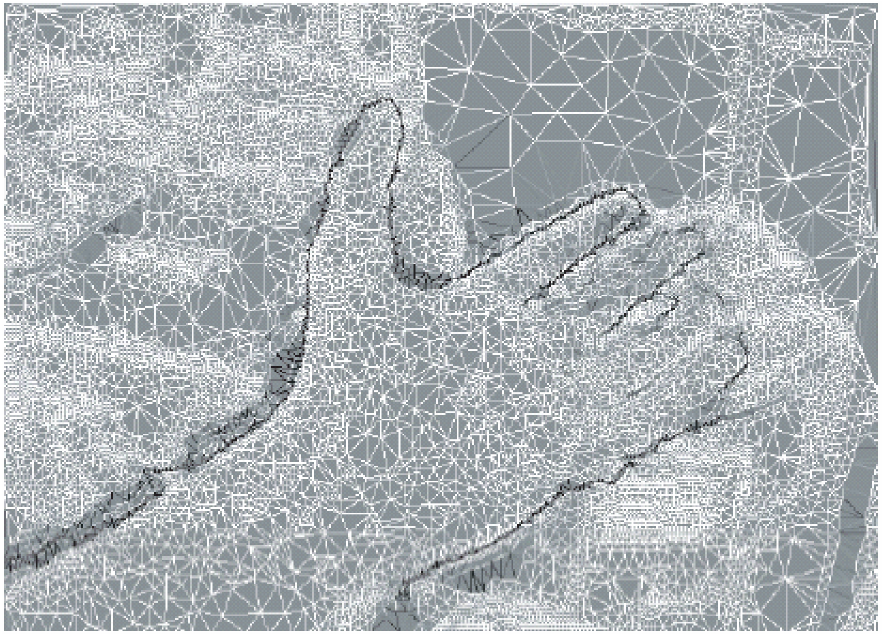

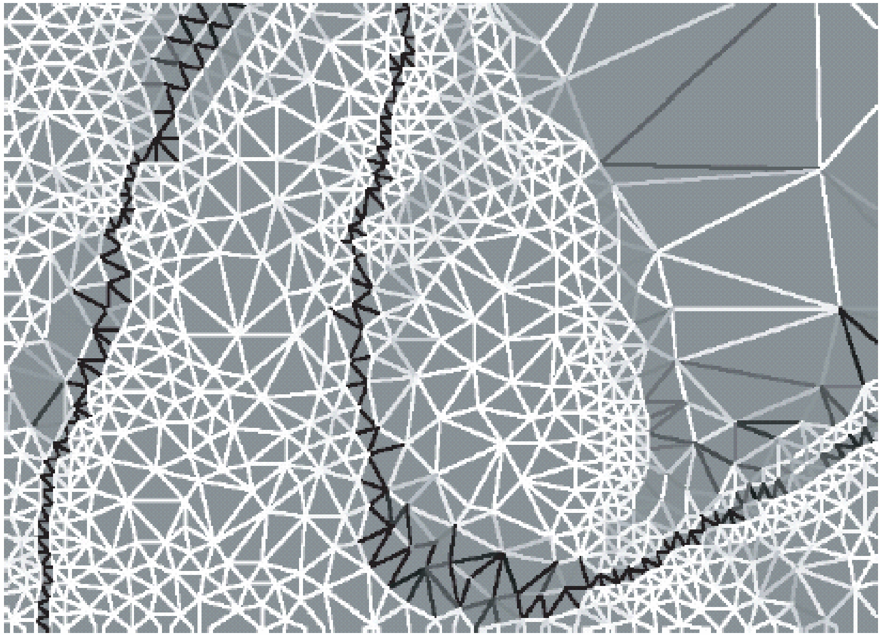

Link Weighting based on Flow Gradient

(lighter = stronger)

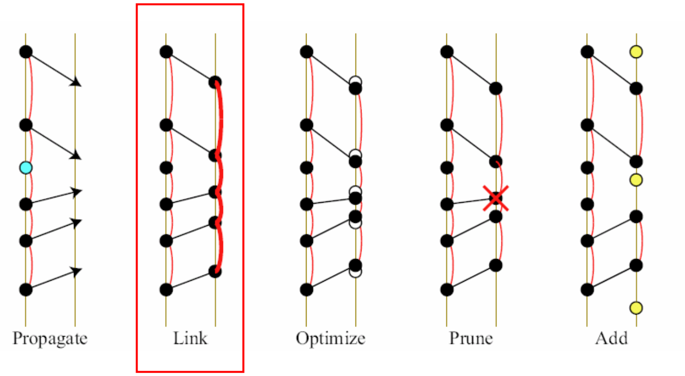

Particle Linking

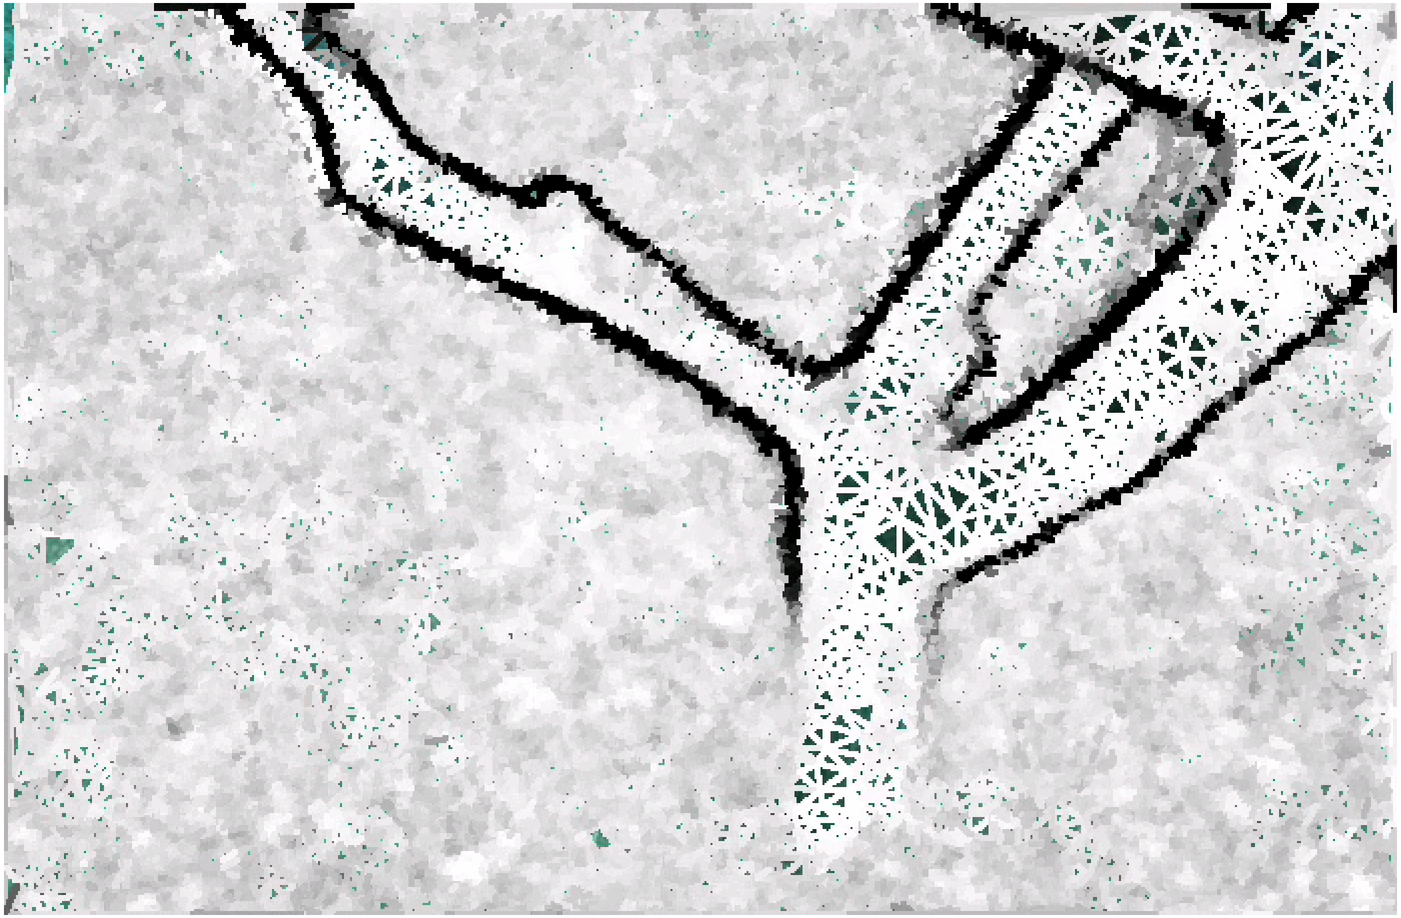

Particle Optimization

Optimization objective function:

$$E(i,t) = \sum_{k \in K_i(t)} E_{Data}^{[k]}(i,t) + \alpha_d \sum_{j \in L_i(t)} E_{Distort}(i,j,t)$$

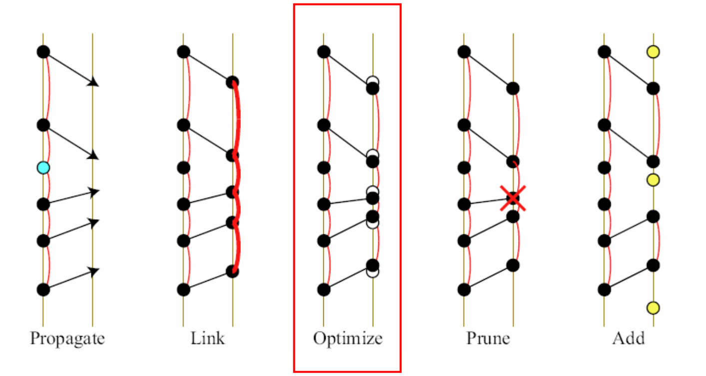

Particle Optimization: Data

Data term:

$$E_{Data}^{[k]}(i,t) = \Psi([I^{[k]}(x_i(t), y_i(t), t) - \hat{c}_i^{[k]}(t)]^2)$$


Brightness


Green - Red


Green - Blue

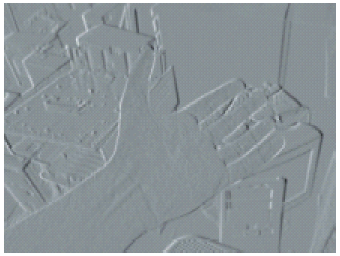
x Gradient

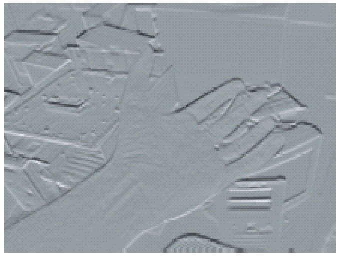
y Gradient

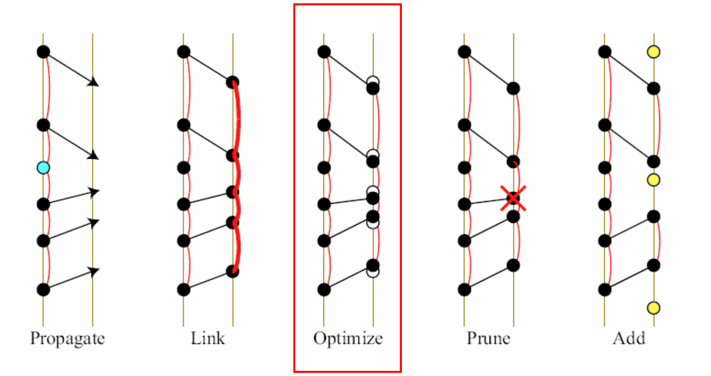

Particle Optimization: Data

Data term:

$$E^{[k]}_{Data}(i,t) = \Psi([I^{[k]}(x_i(t),y_i(t),t) - \hat{c}^{[k]}_i(t)]^2)$$

Observed Channel Value

Filtered Channel Value

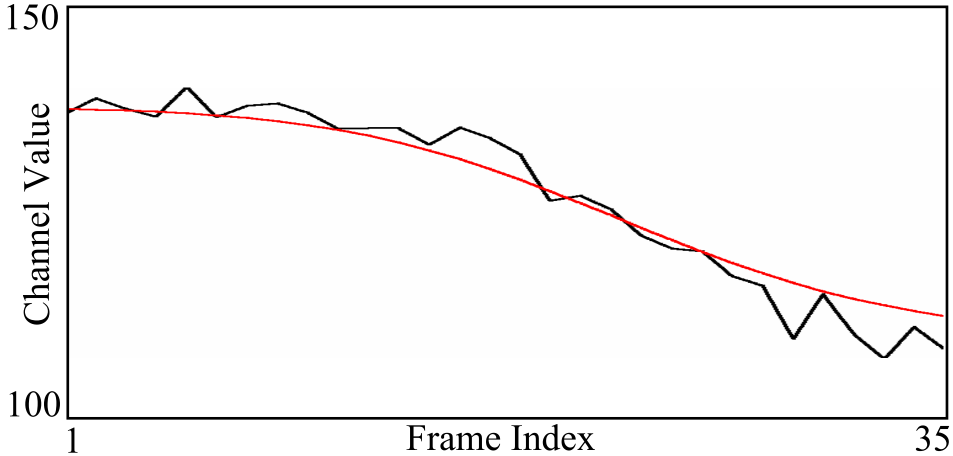

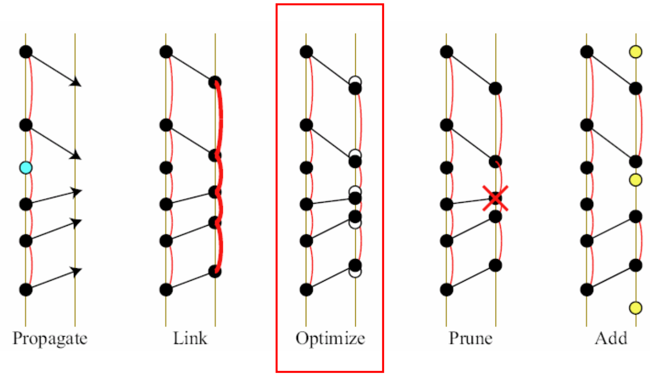

Particle Optimization: Distortion

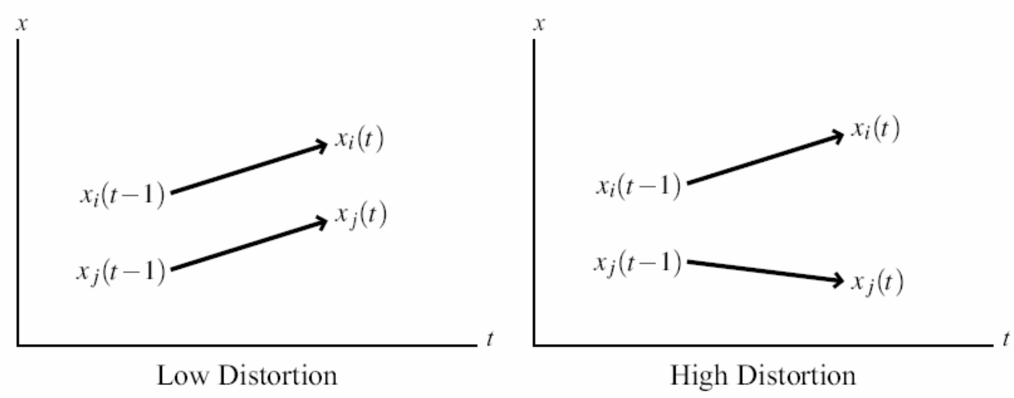

Distortion term:

$$E_{Distort}(i,j,t) = l_{ij}(t)\Psi([u_i(t)-u_j(t)]^2+[v_i(t)-v_j(t)]^2)$$

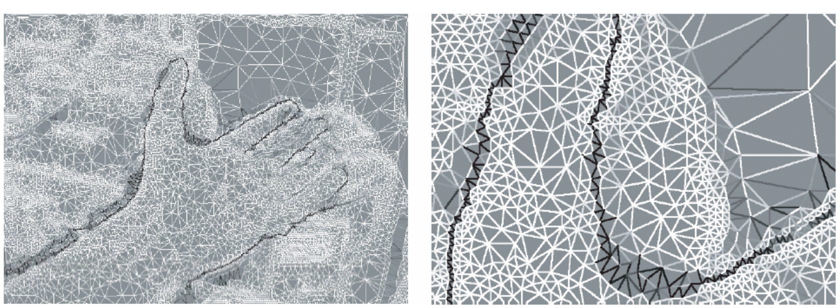

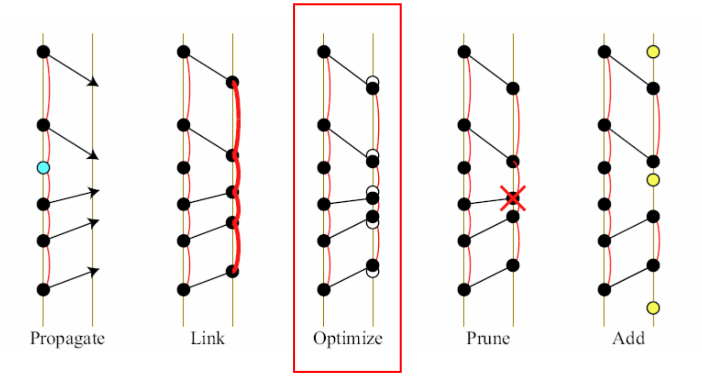

Particle Optimization

Loop until convergence:

- Solve system for $dx_i(t)$, $dy_i(t)$ using SOR:

$$\left\{ \frac{\delta E}{\delta dx_i(t)} = 0, \frac{\delta E}{\delta dy_i(t)} = 0 \mid i \in P, t \in F \right\}$$

- $x_i(t) \leftarrow x_i(t) + dx_i(t)$
- $y_i(t) \leftarrow y_i(t) + dy_i(t)$
- Update link weights, etc.

(see paper for more details)

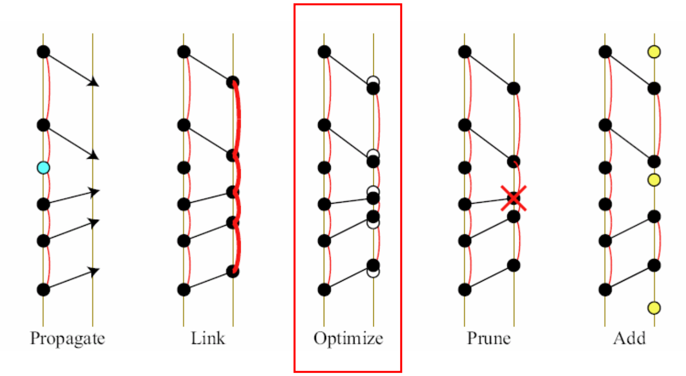

Particle Pruning

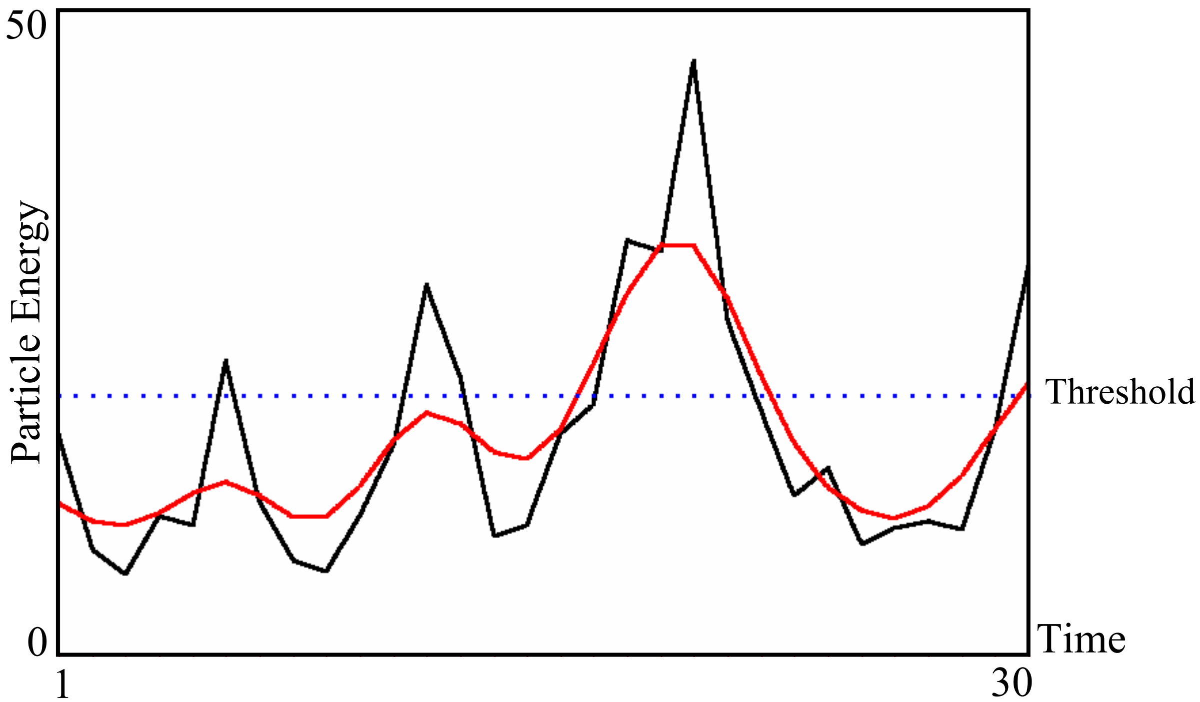

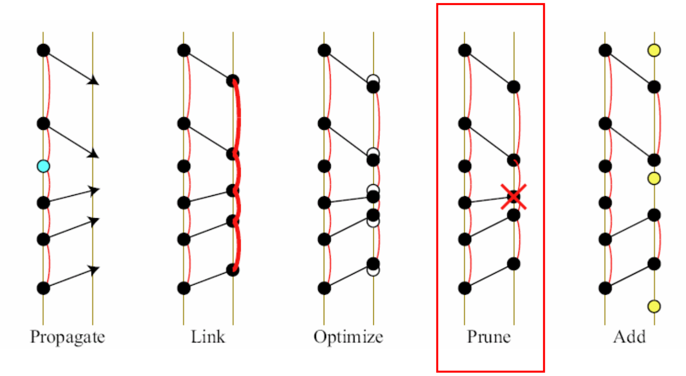

Particle Pruning

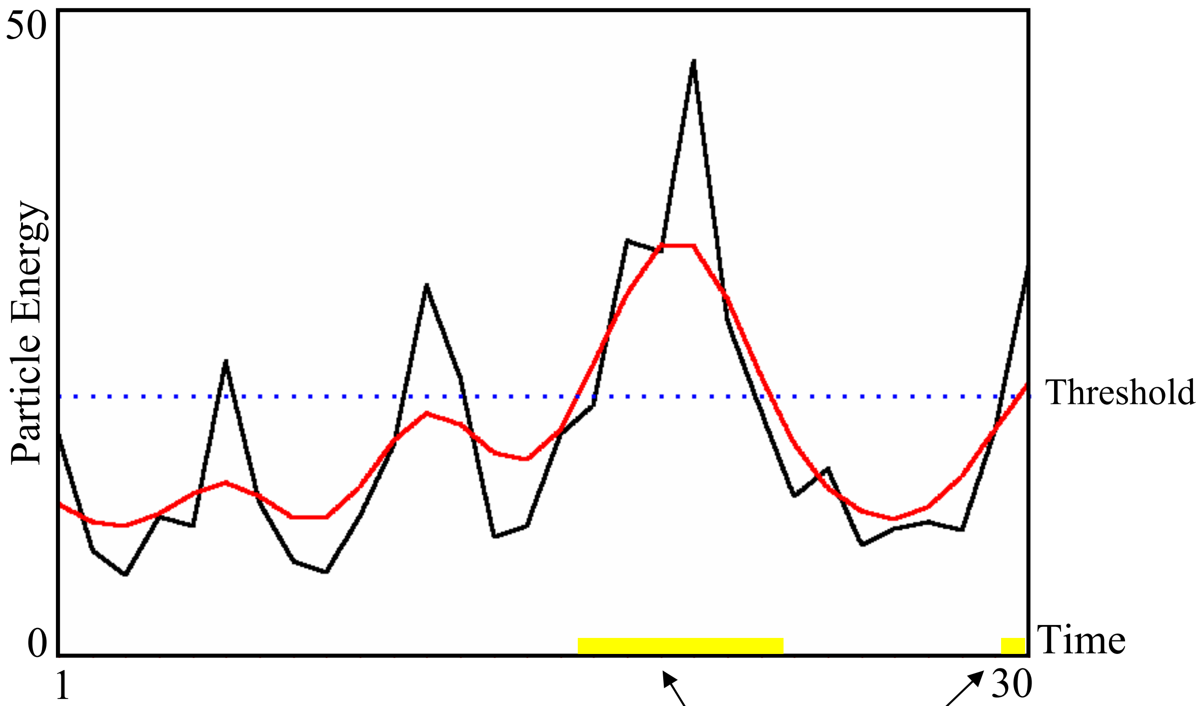

Above Threshold:
Deactivate

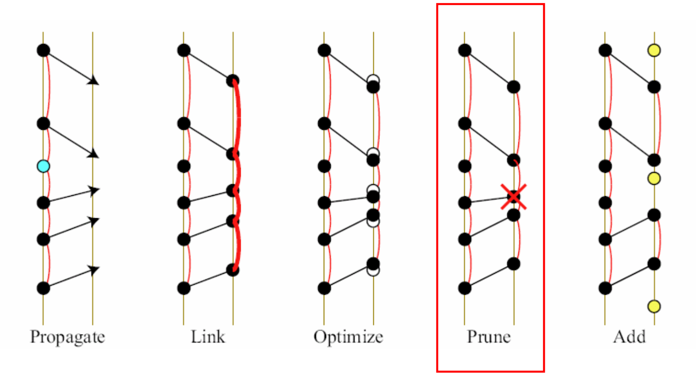

Particle Addition

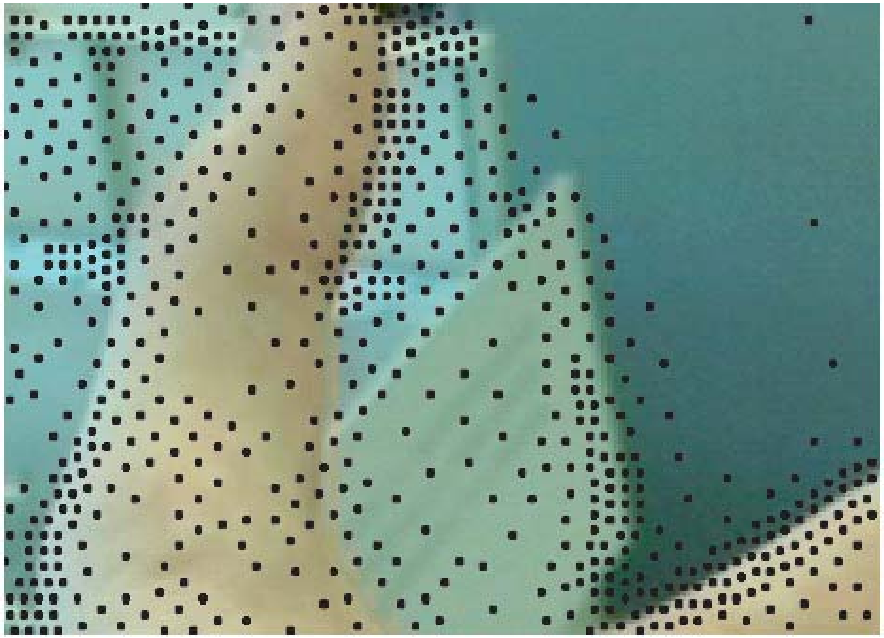

Particle Placement

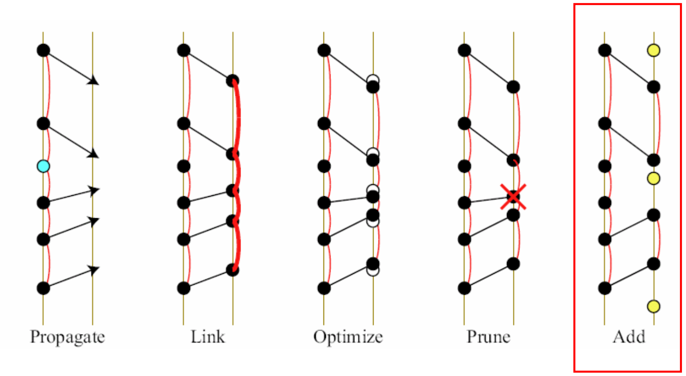

Particle Addition

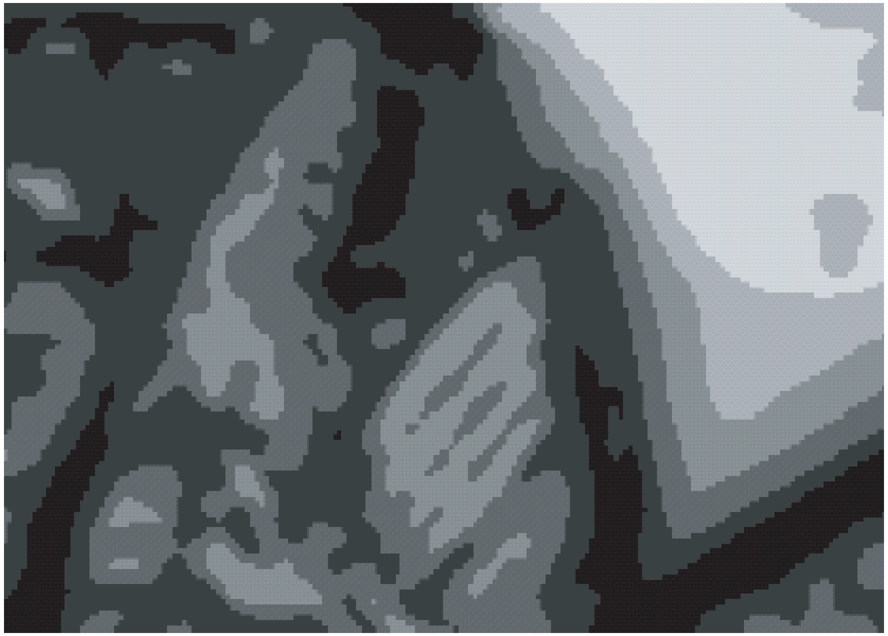

Scale Map

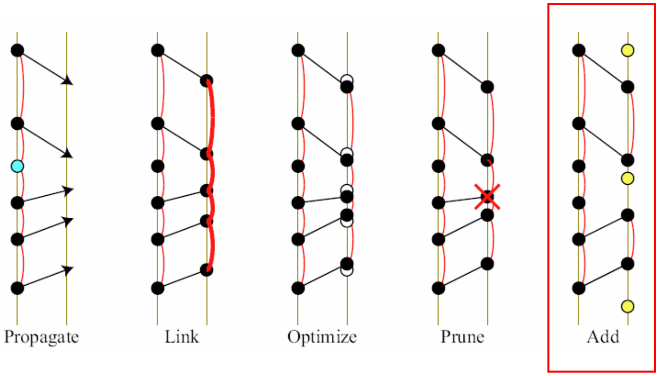

Particle Video Algorithm

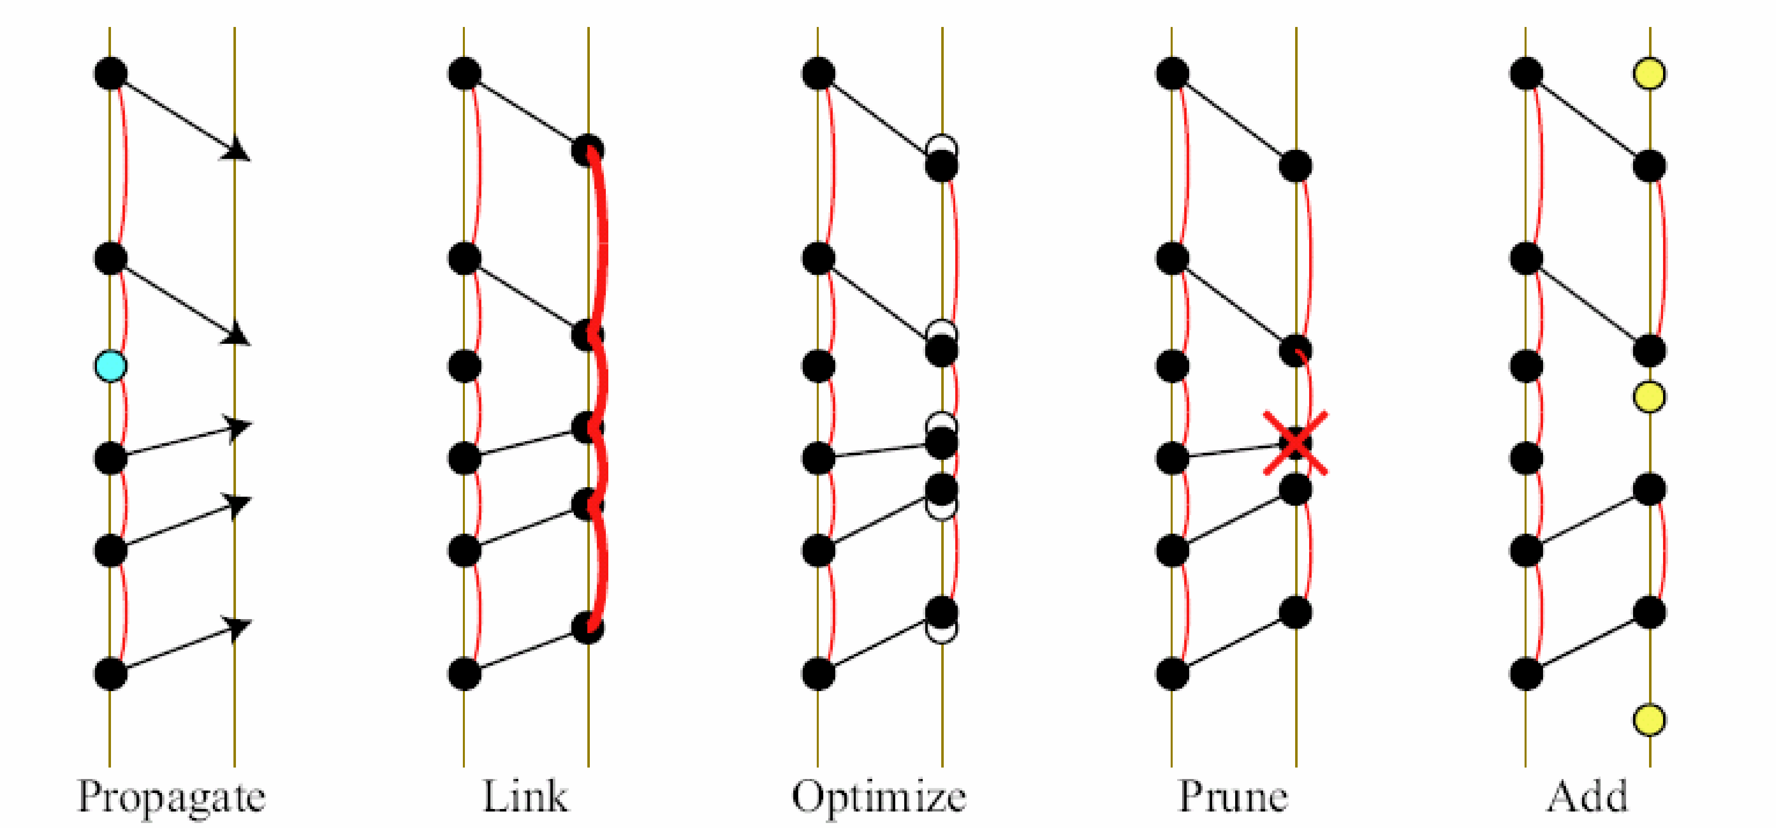

Evaluation Videos

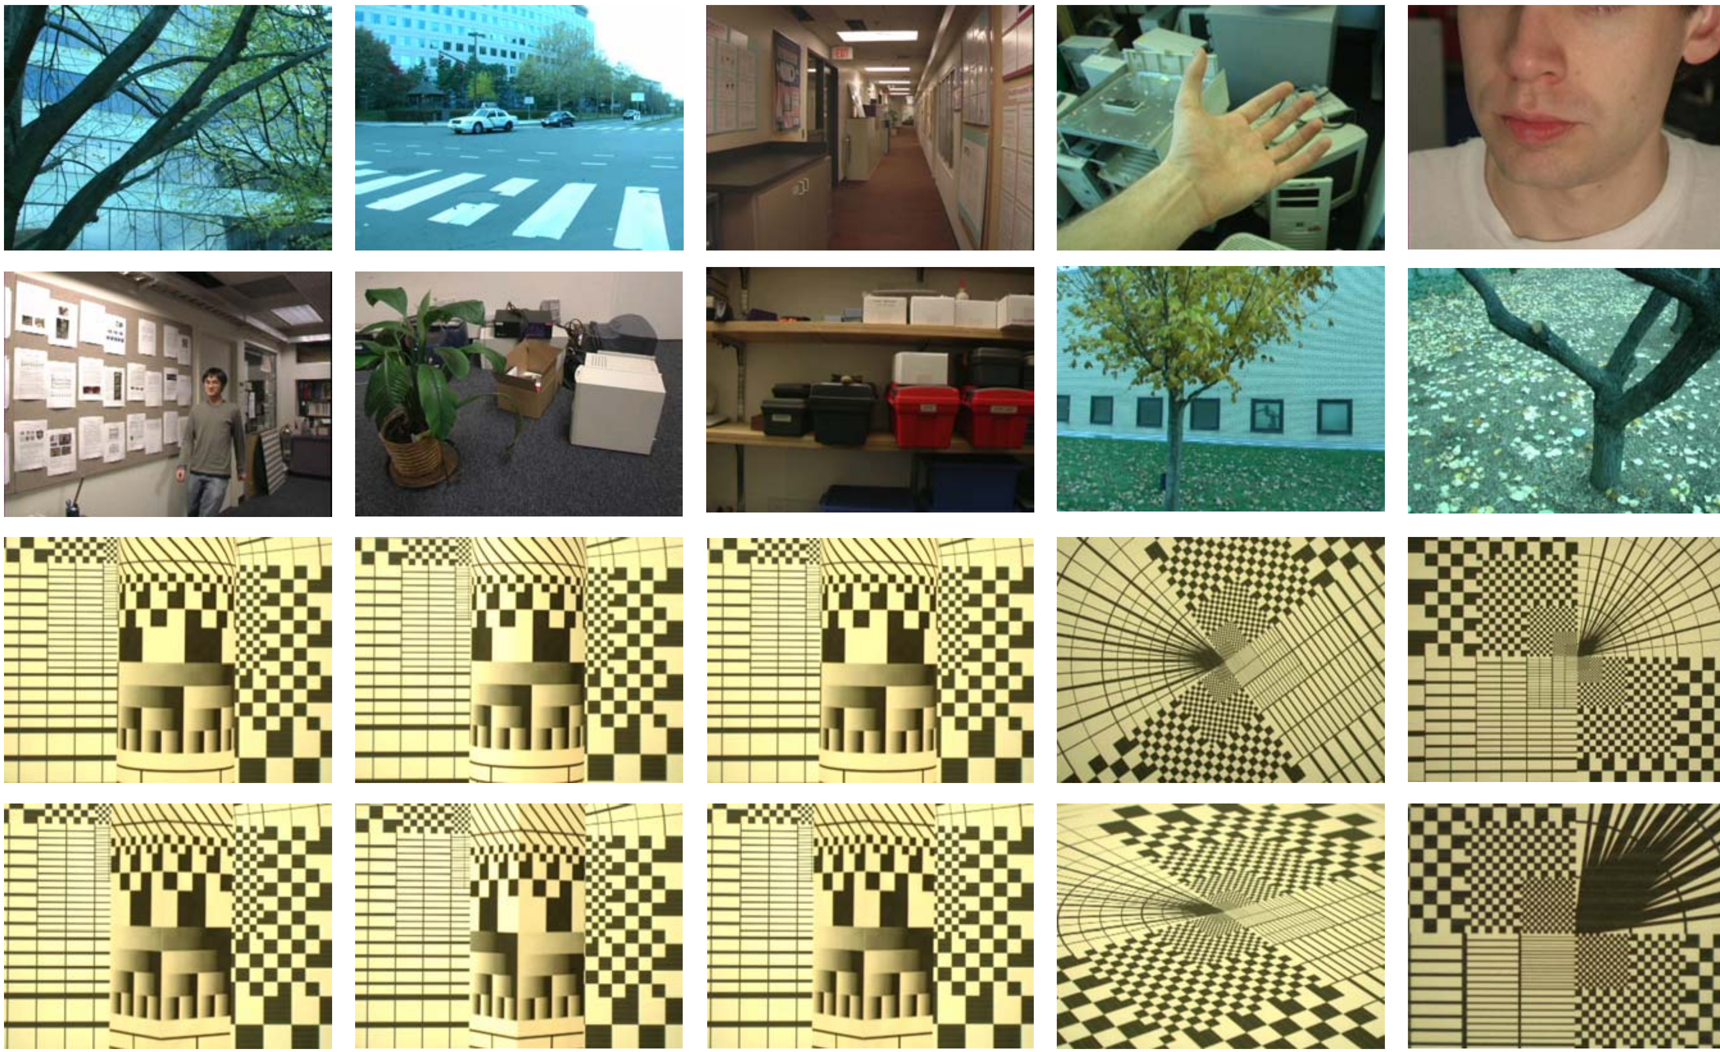

Results / Evaluation

Construct videos that return to the start frame:

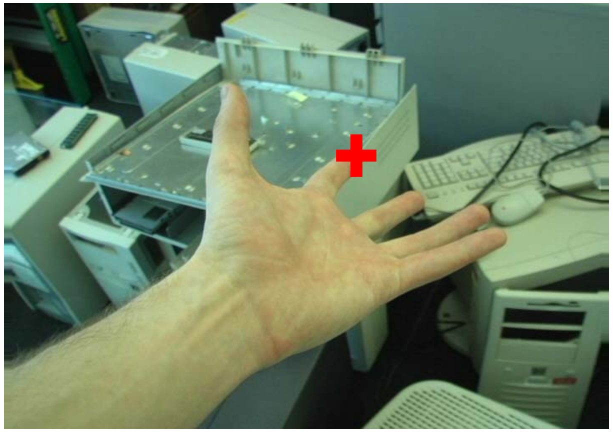

1, 2, 3, … …, N-1, N, N-1, … …3, 2, 1

Results / Evaluation

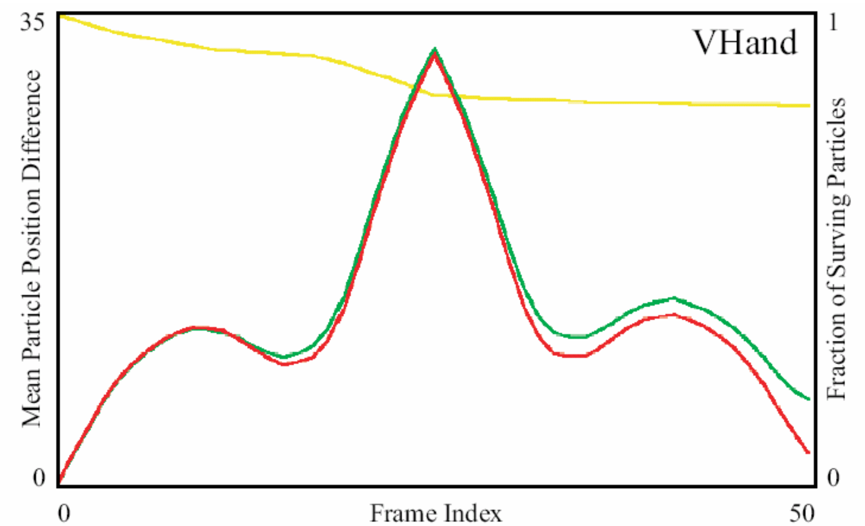

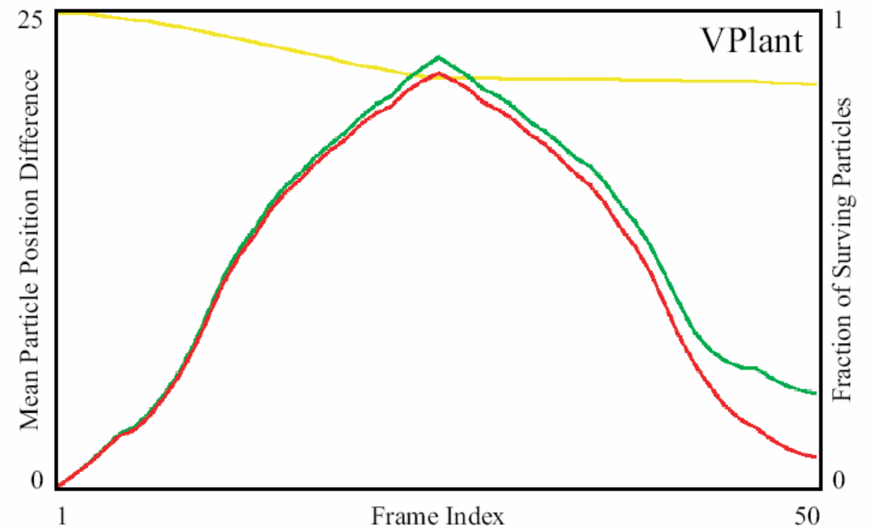

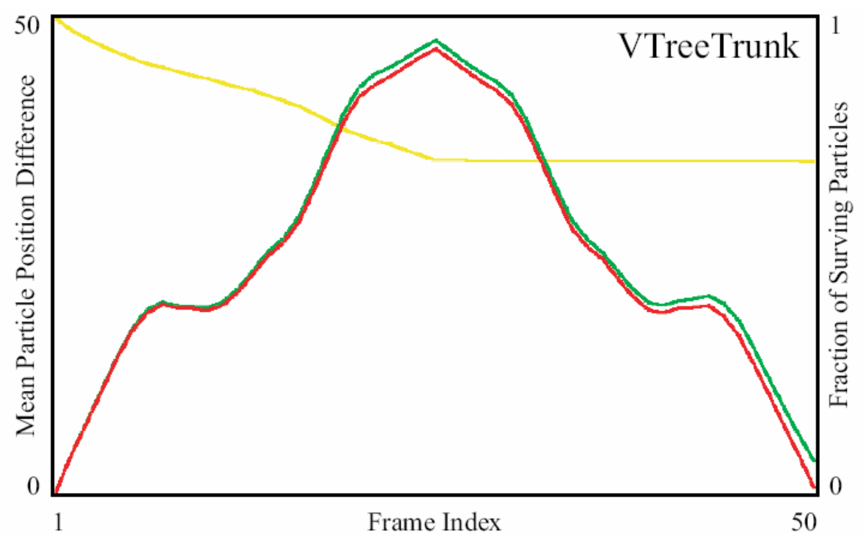

Particle distance: red
Concatenated flow distance: green
Fraction surviving: yellow

Results / Evaluation

Results / Evaluation

Results / Evaluation

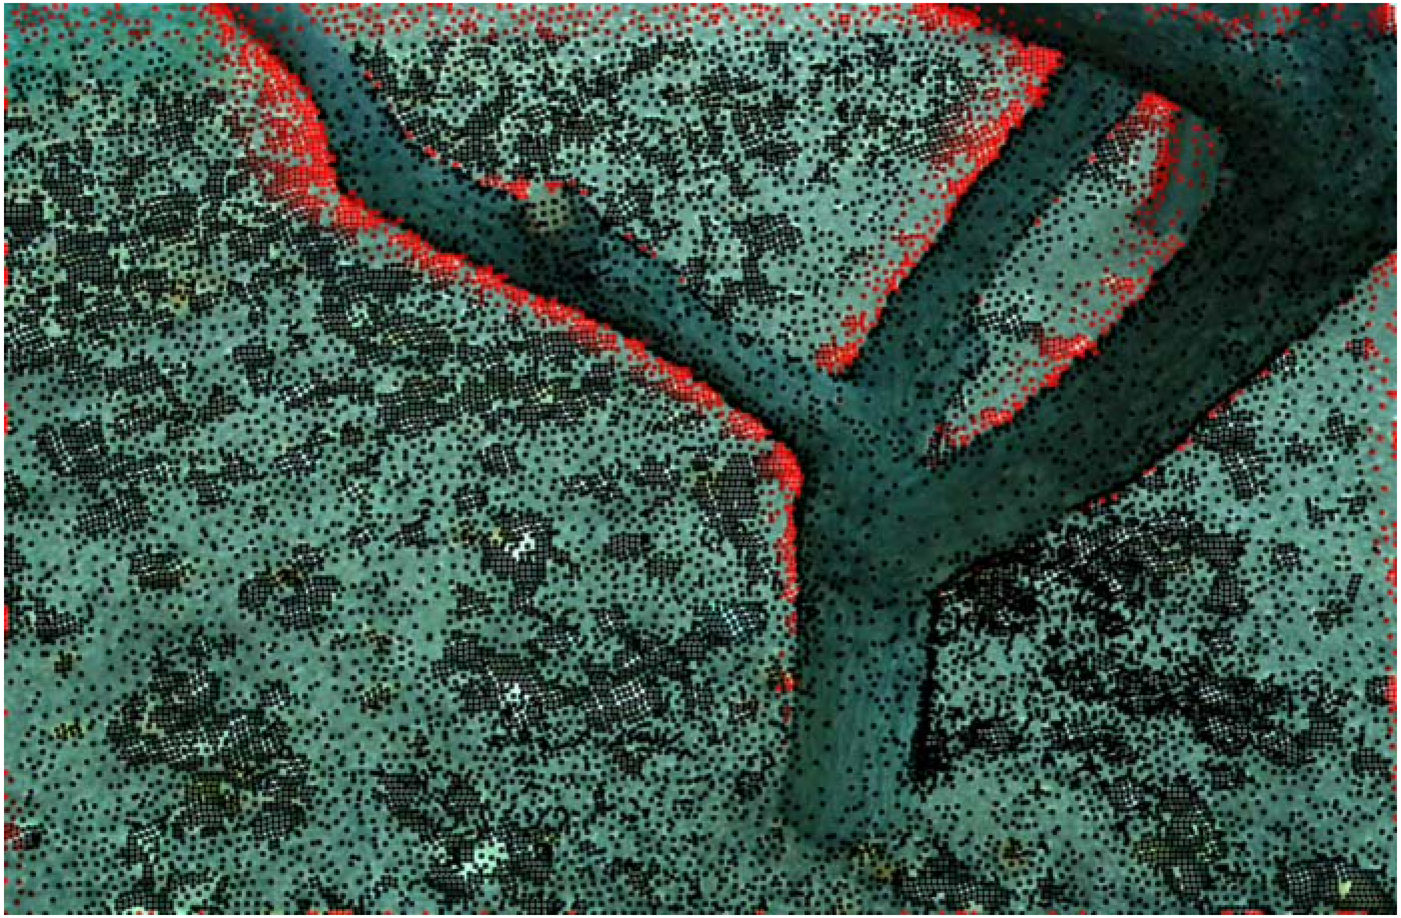

Results / Evaluation

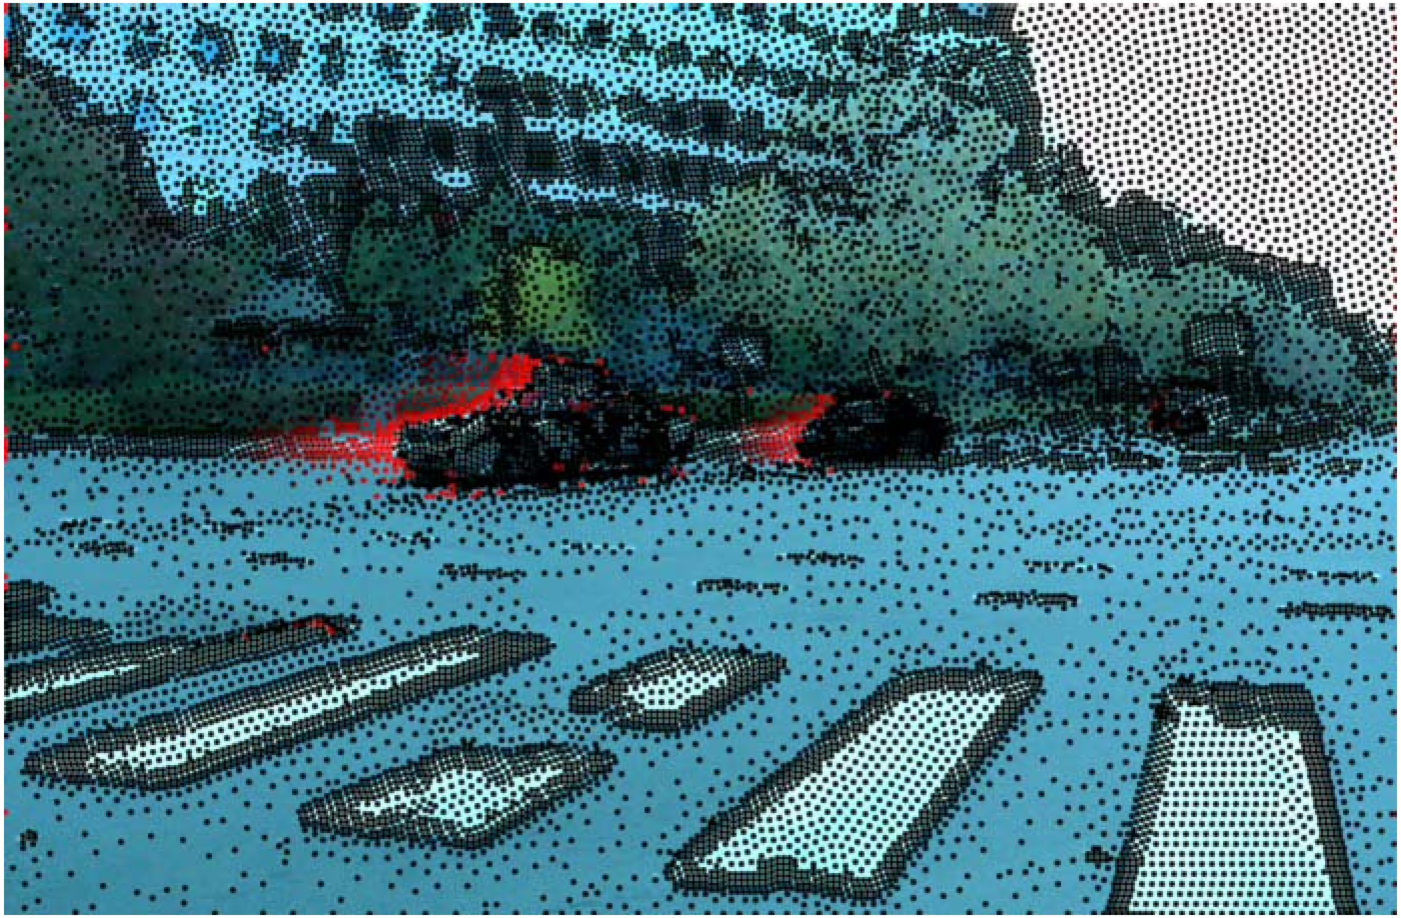

Results / Evaluation

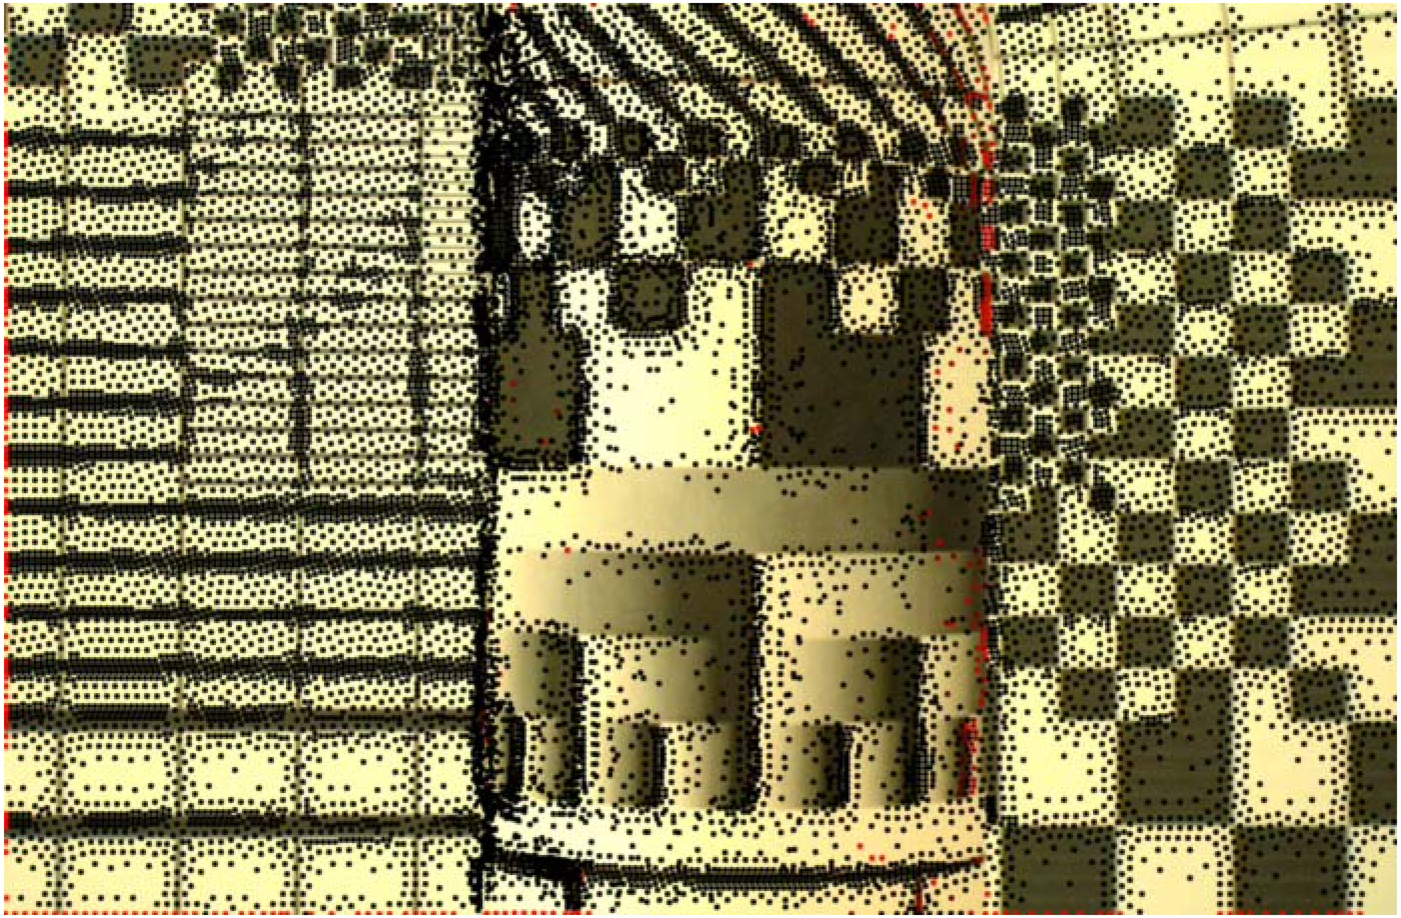

Results / Evaluation

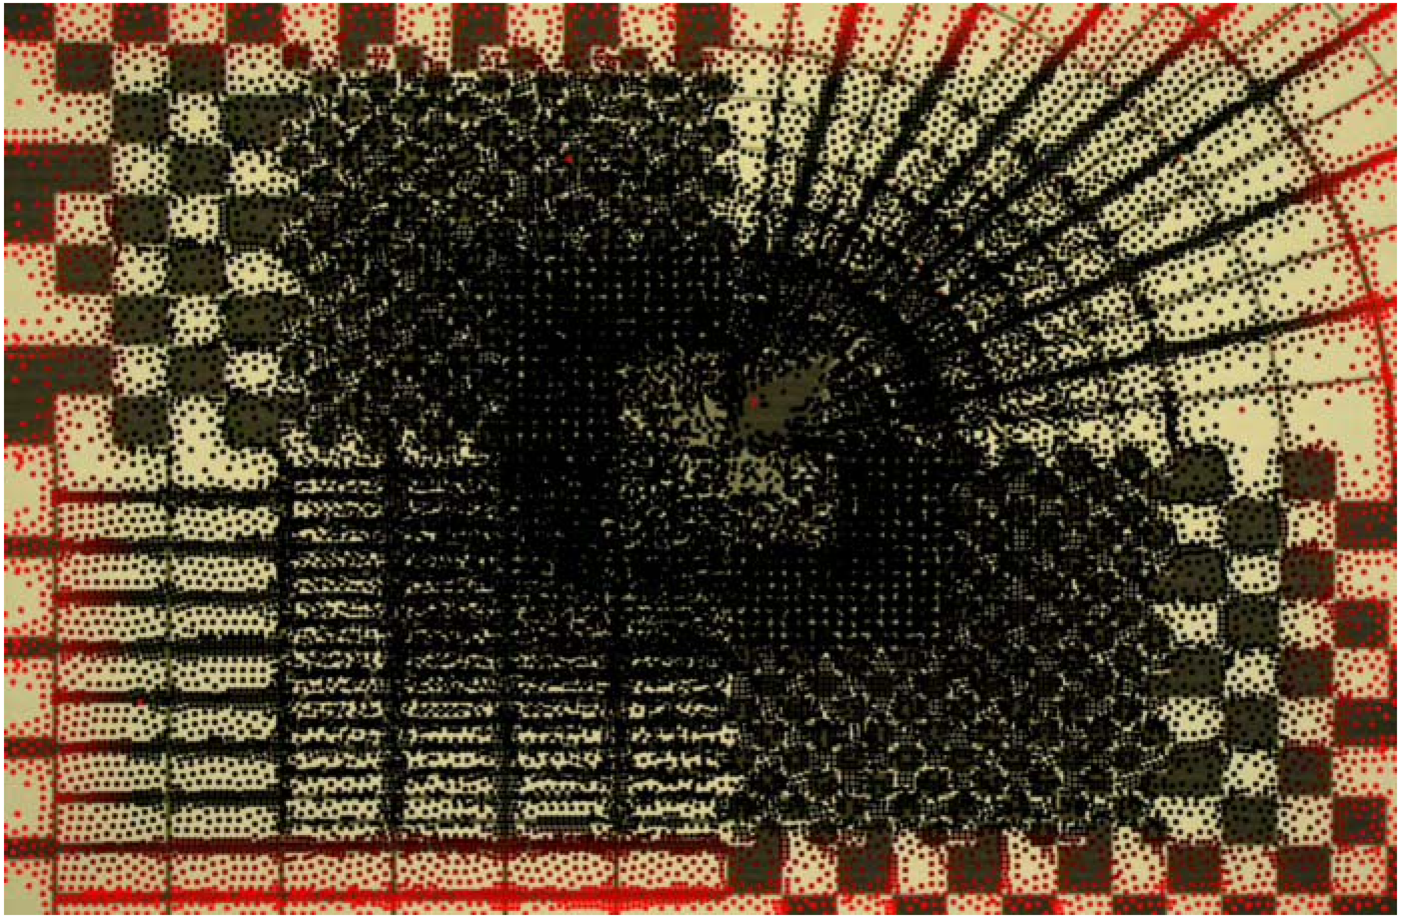

Results / Evaluation

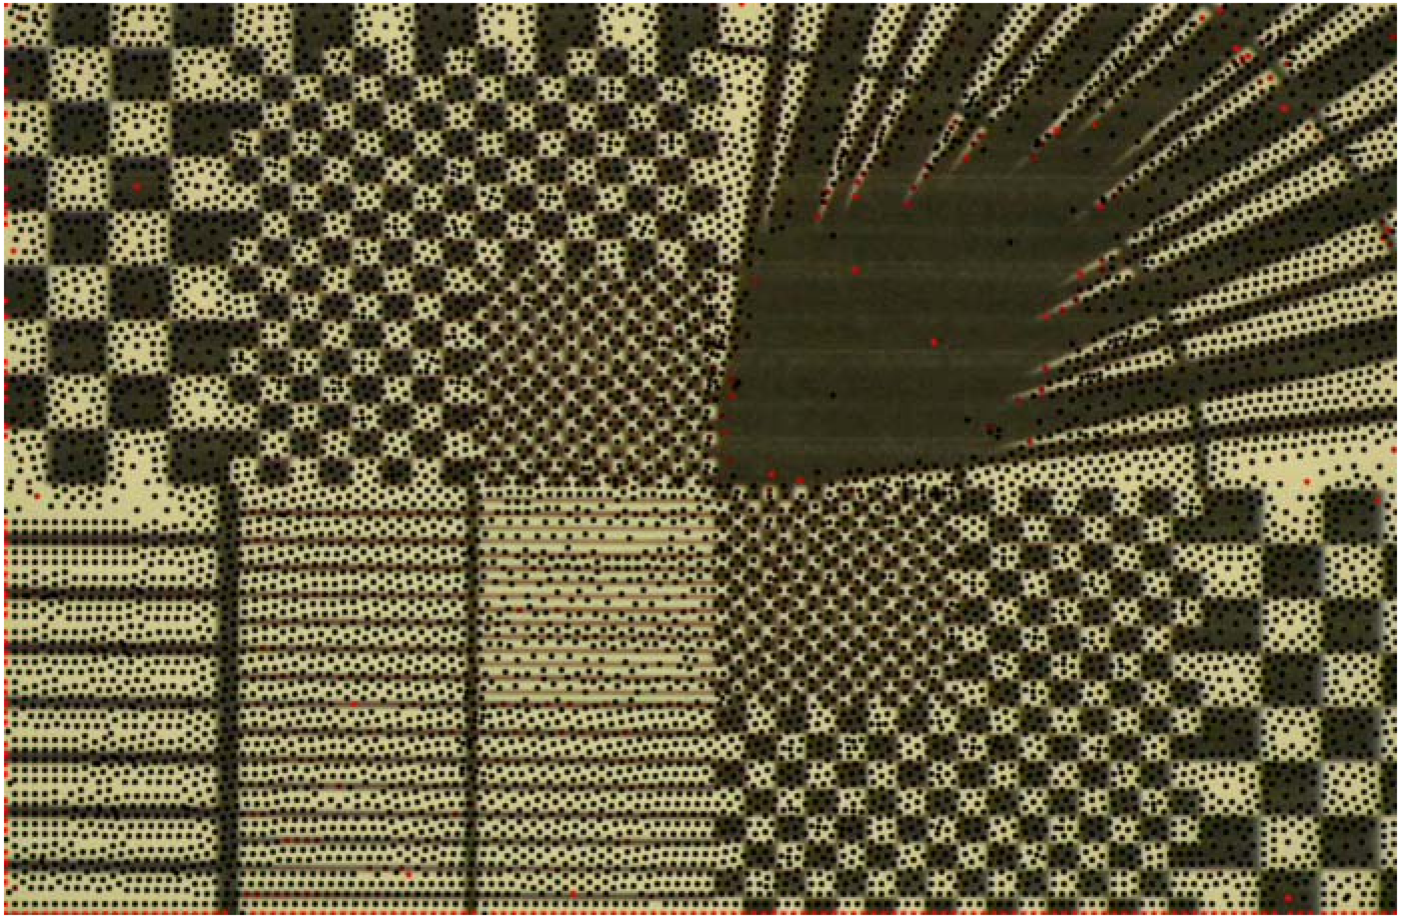

Results / Evaluation

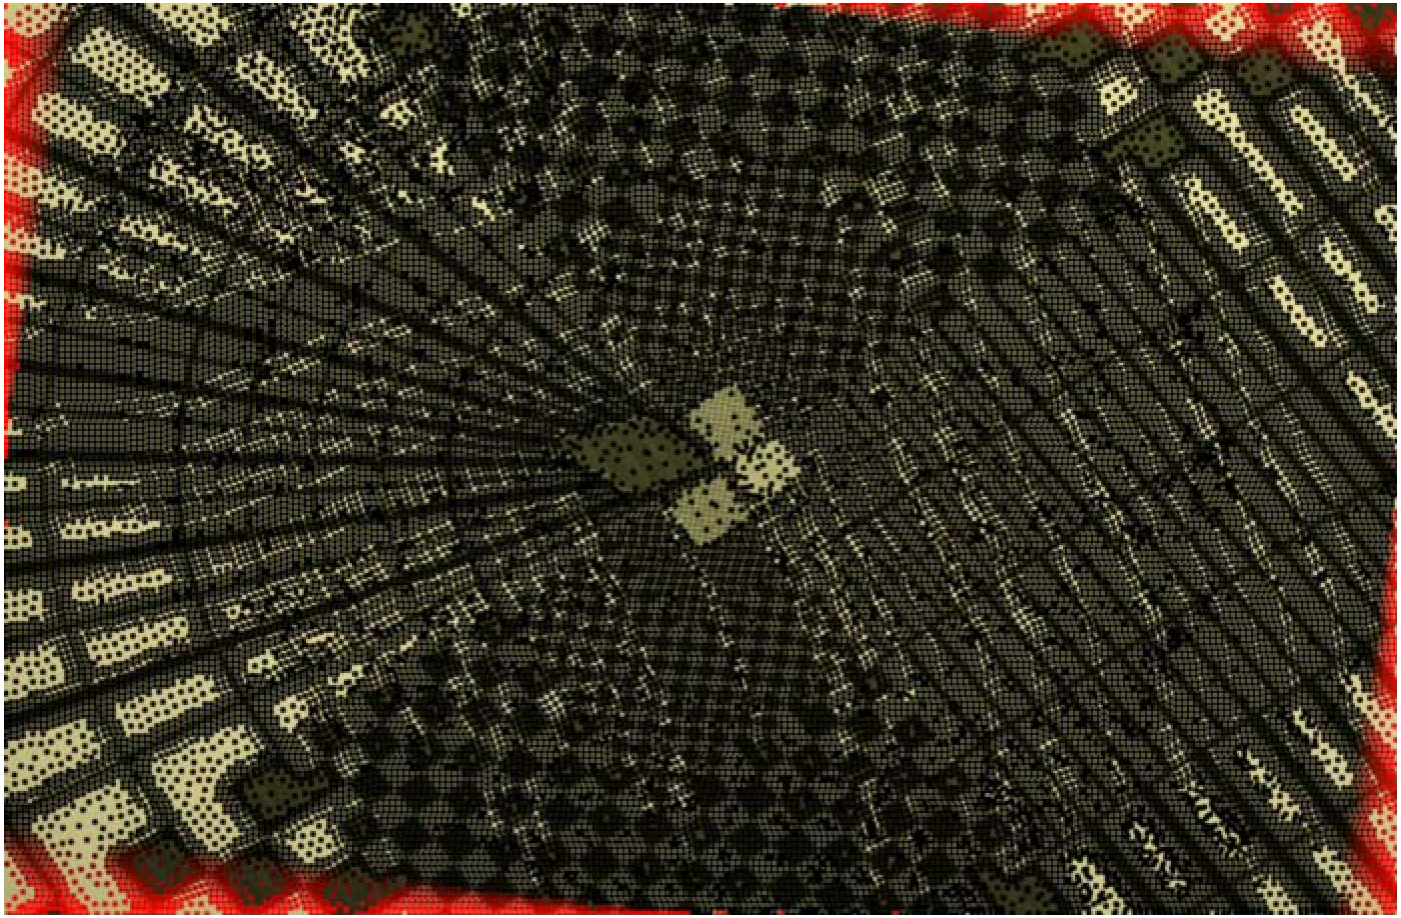

Failure Modes

Failure Modes

Limitations / Future Work

Issue: occlusion handling

Possible solution: analyze local motion histories to distinguish good/bad distortion

Issue: flow dependence

Possible solution: hybrid flow / particle optimization

Limitations / Future Work

Issue: appearance changes due to reflectance and scaling

Possible solution: invariant feature descriptors for particles away from occlusions

Other areas of exploration:

- Faster algorithms (currently 40 seconds/frame)
- Geometric constraints
- Batch particle positioning
- Evaluation on synthetic sequences

Summary

- Particles can represent complex motion and geometry
- Particle representation is useful for application algorithms
- Different from other representations (vector fields, rigid components, layers, tracked feature patches)

More Info

http://rvsn.csail.mit.edu/pv

Peter Sand and Seth Teller. *Particle Video: Long-Range Motion Estimation using Point Trajectories*, CVPR 2006.

Thanks: William Freeman, Berthold Horn, Matt Brand, Bryt Bradley, Tom Buehler, Frédo Durand, Jovan Popović, the Graphics Group, CSAIL, EECS, MIT