Background Subtraction for Temporally Irregular Dynamic Textures

Gerald Dalley, Joshua Migdal,
and W. Eric L. Grimson

Workshop on Applications
of Computer Vision (WACV) 2008

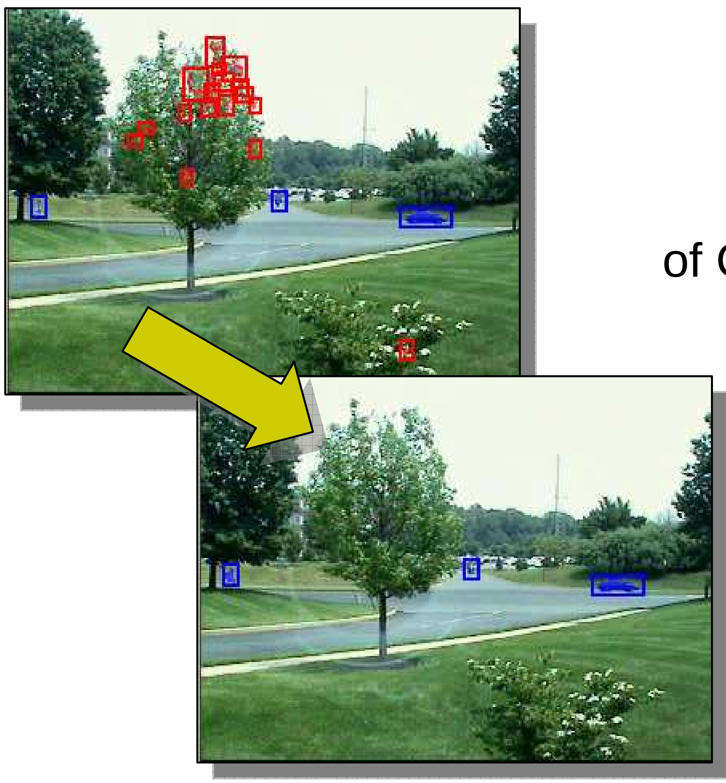

Key Prior Approaches

- MoG – learn modes for all pixels & track changes
  - Works as long as dynamic textures are regular
- Per-pixel kernels (Mittal & Paragios)
  - Can pay attention to higher-level features like optical flow
  - How many kernels to keep?
- Spatial kernels (Sheikh & Shah)
  - Handles temporal irregularity in dynamic textures
  - How many kernels to keep?

Results with Mittal & Paragios Traffic Clip

Ground truth, frame 150

Mittal and Paragios: 1 TP, 0 FP, 2 FN

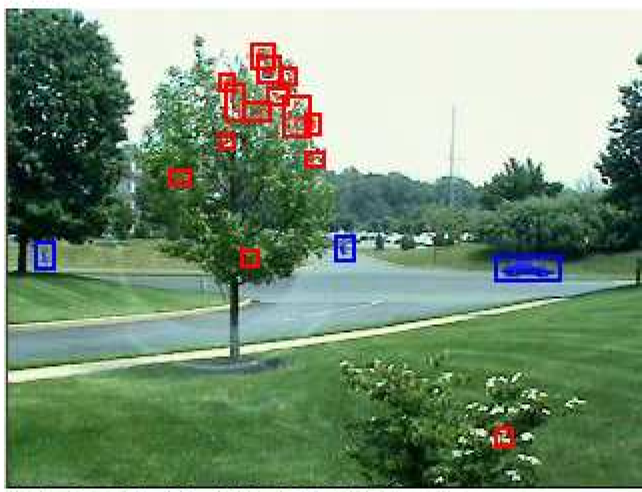

MoG: 3 TP, 14 FP, 0 FN

Ours: 3 TP, 0 FP, 0 FN

Key:

- **bboxes from ground truth**
  - **True positive**
  - **False positive**
  - **False negative**
- **No ground truth available**

Our Goals

- Compactness
  - Speed
  - Complexity growth
    - Grow with scene complexity, *not* with the amount of time to see the complexity
      - If a leaf blows in front of a pixel once a minute, we don't want to have to keep 1800 kernels
- Usable in MoG frameworks
  - MoG variants are very widely used (original Stauffer & Grimson paper has 1000+ citations)

Our Model

- Mixture of Gaussians (MoG)

$$p(c_i|\Phi) \propto \sum_{\underline{j \in N_i}} w_j \mathcal{N}(c_i; \mu_j, \Sigma_j)$$

c_i the observed color at pixel i

Φ the model $\{w_j, \mu_j, \Sigma_j\}_j$

$\boxed{N_i \quad \text{neighborhood of pixel } i}$

Pixels Affected by a Given Mixture Component

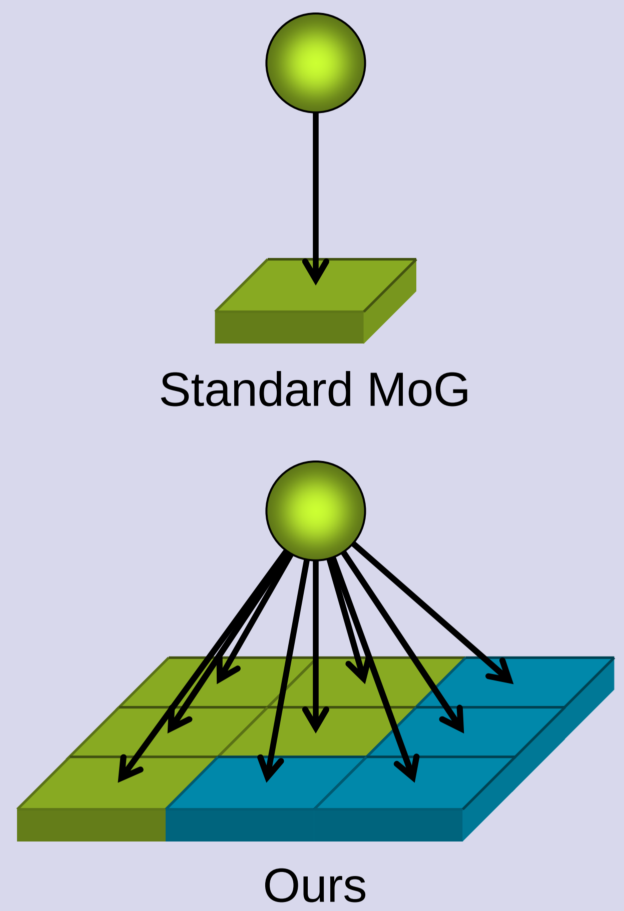

Foreground/Background Classification

- Find best matching background Gaussian, j
  - Use neighborhood
- Squared Mahalanobis Distance

$$d_{ij} = (c_i - \mu_j)^T \Sigma_j^{-1} (c_i - \mu_j)$$

-

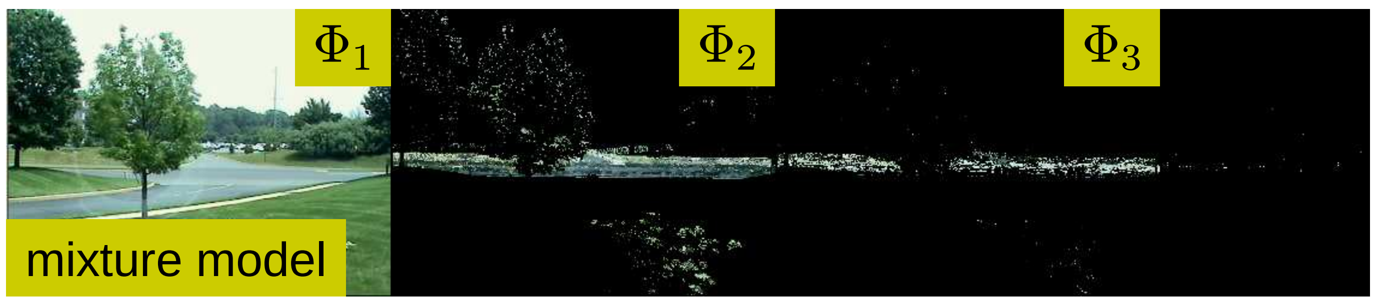

=

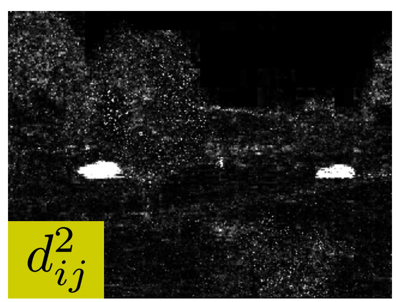

Model Update Options

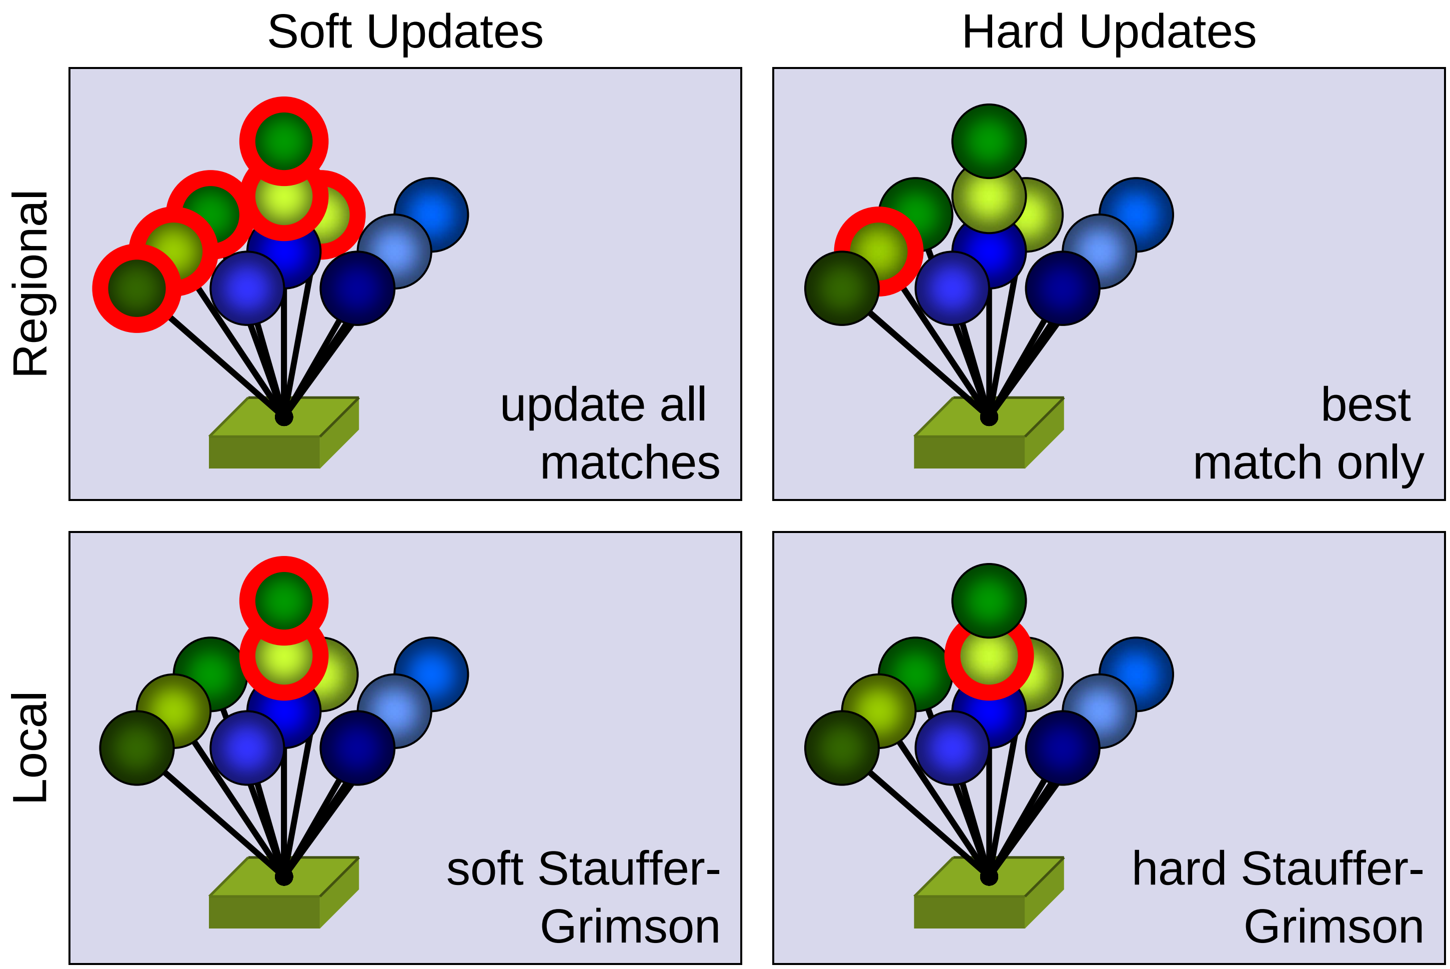

Results on a Classic Dataset: Wallflower WavingTrees

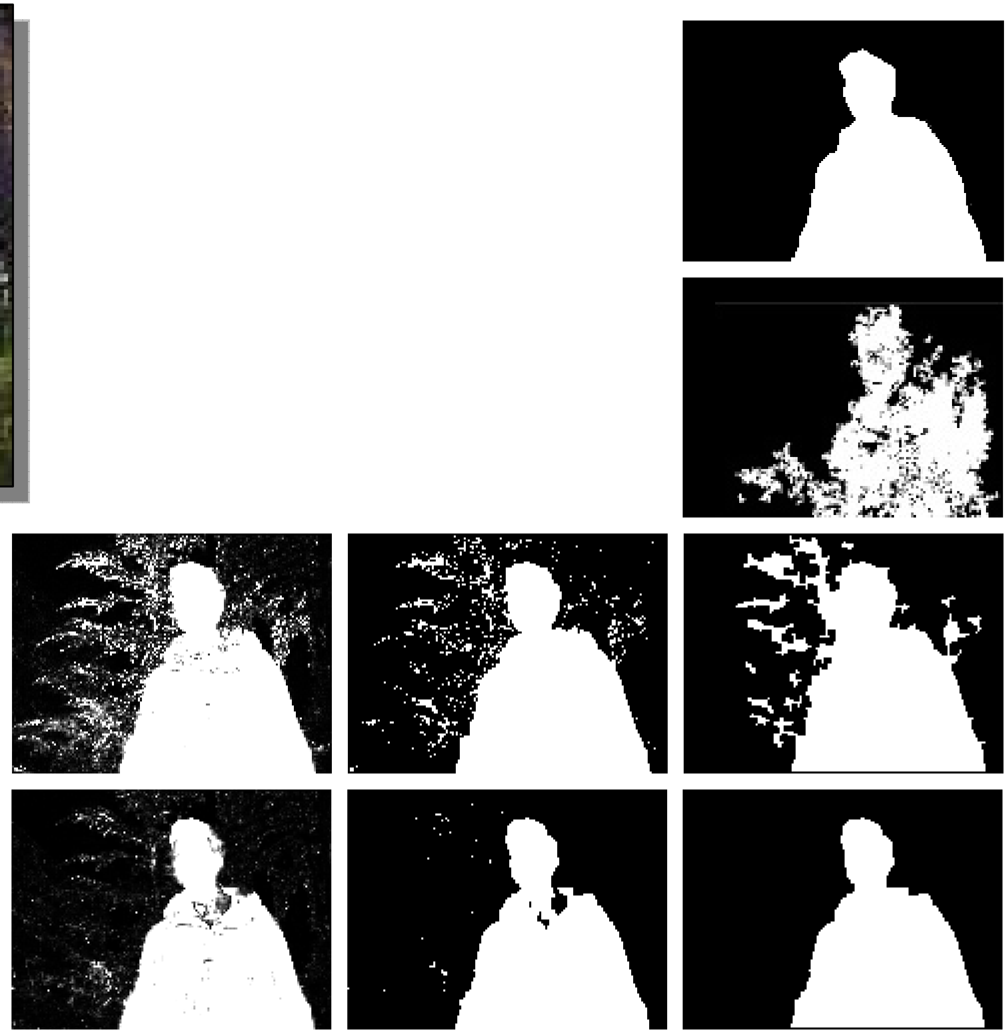

ROC (Wallflower)

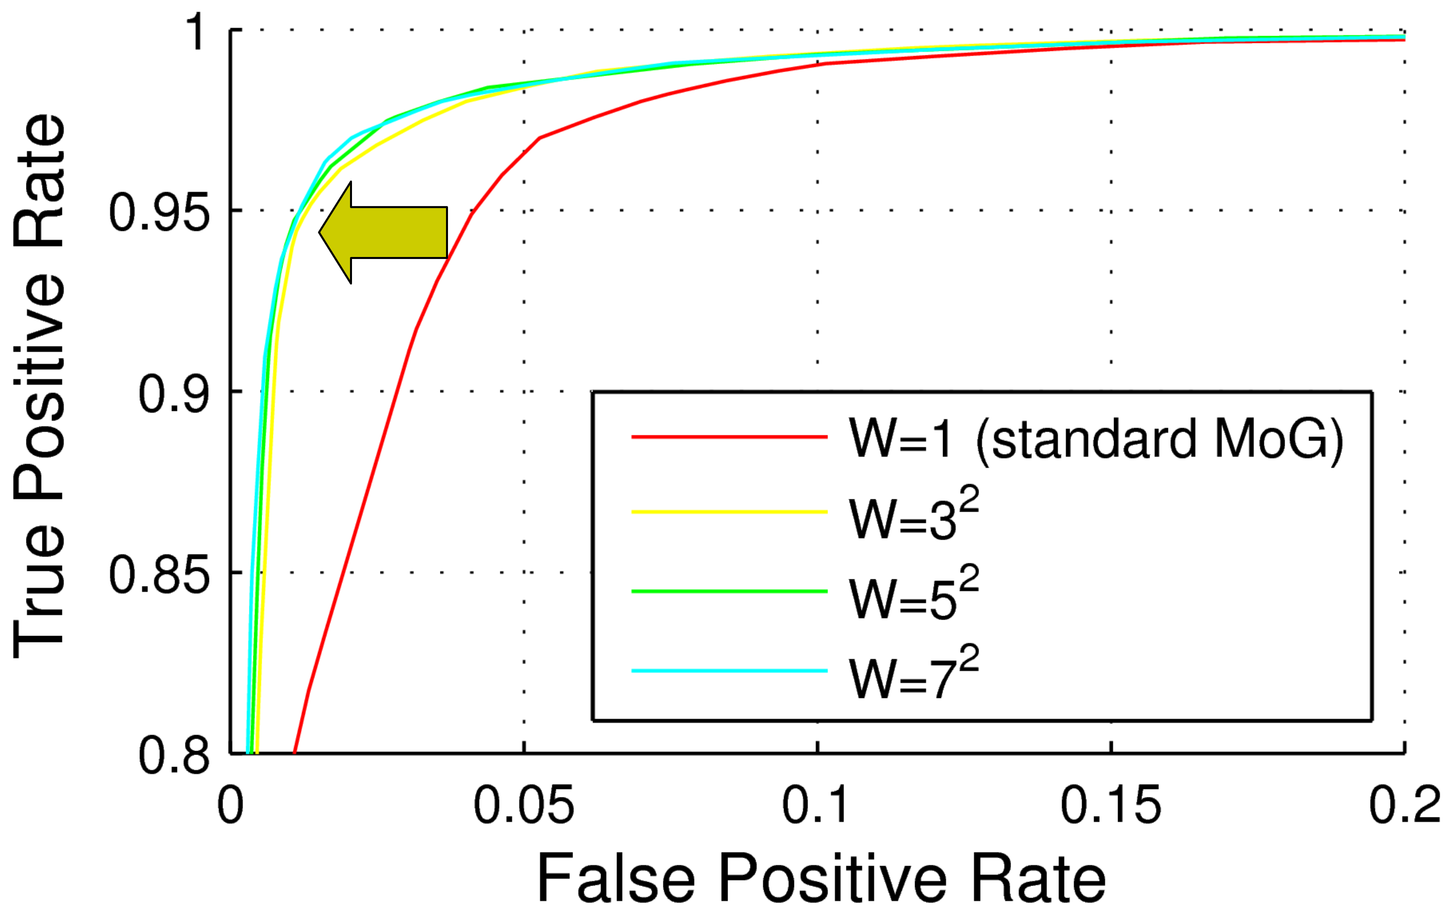

Traffic #410

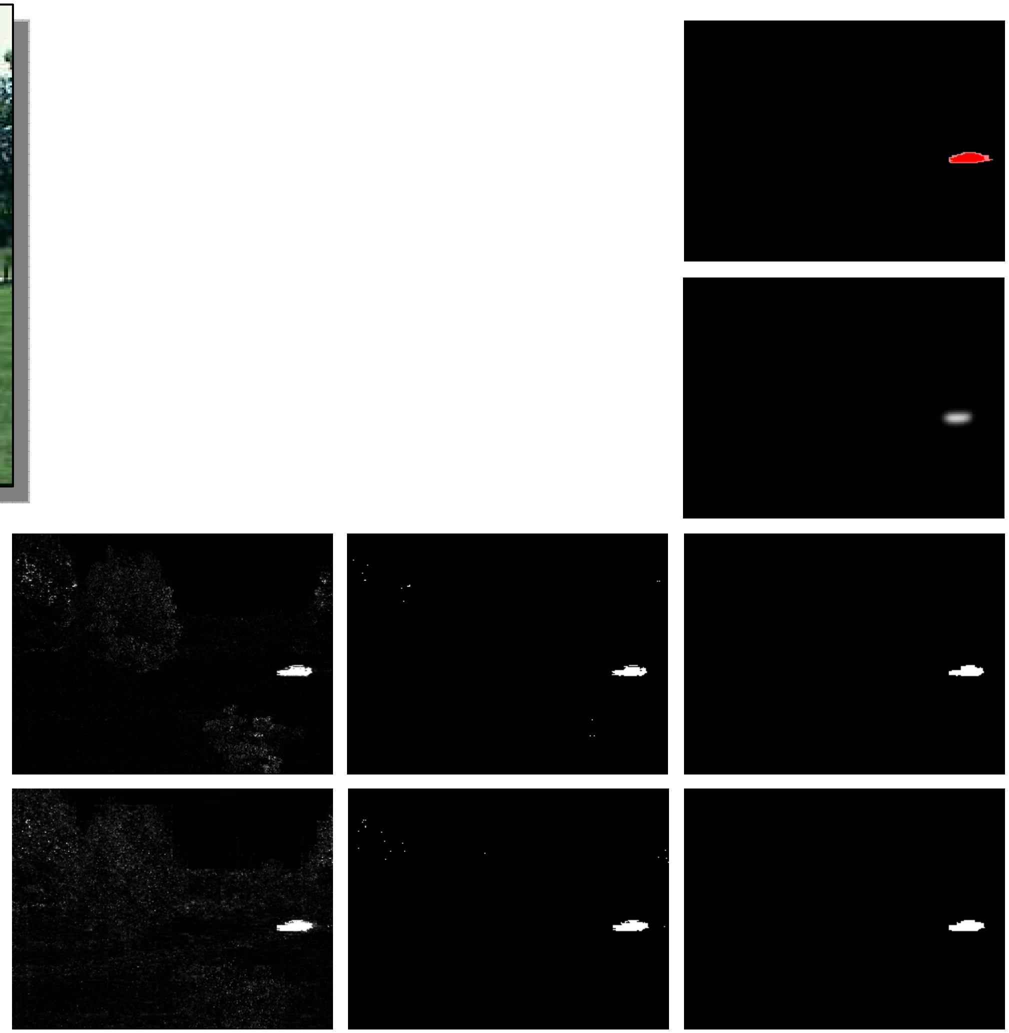
Ground Truth
Mittal & Paragios
Best MoG
Ours
Mahalanobis Distances
After MRF
Final Result

Traffic #150

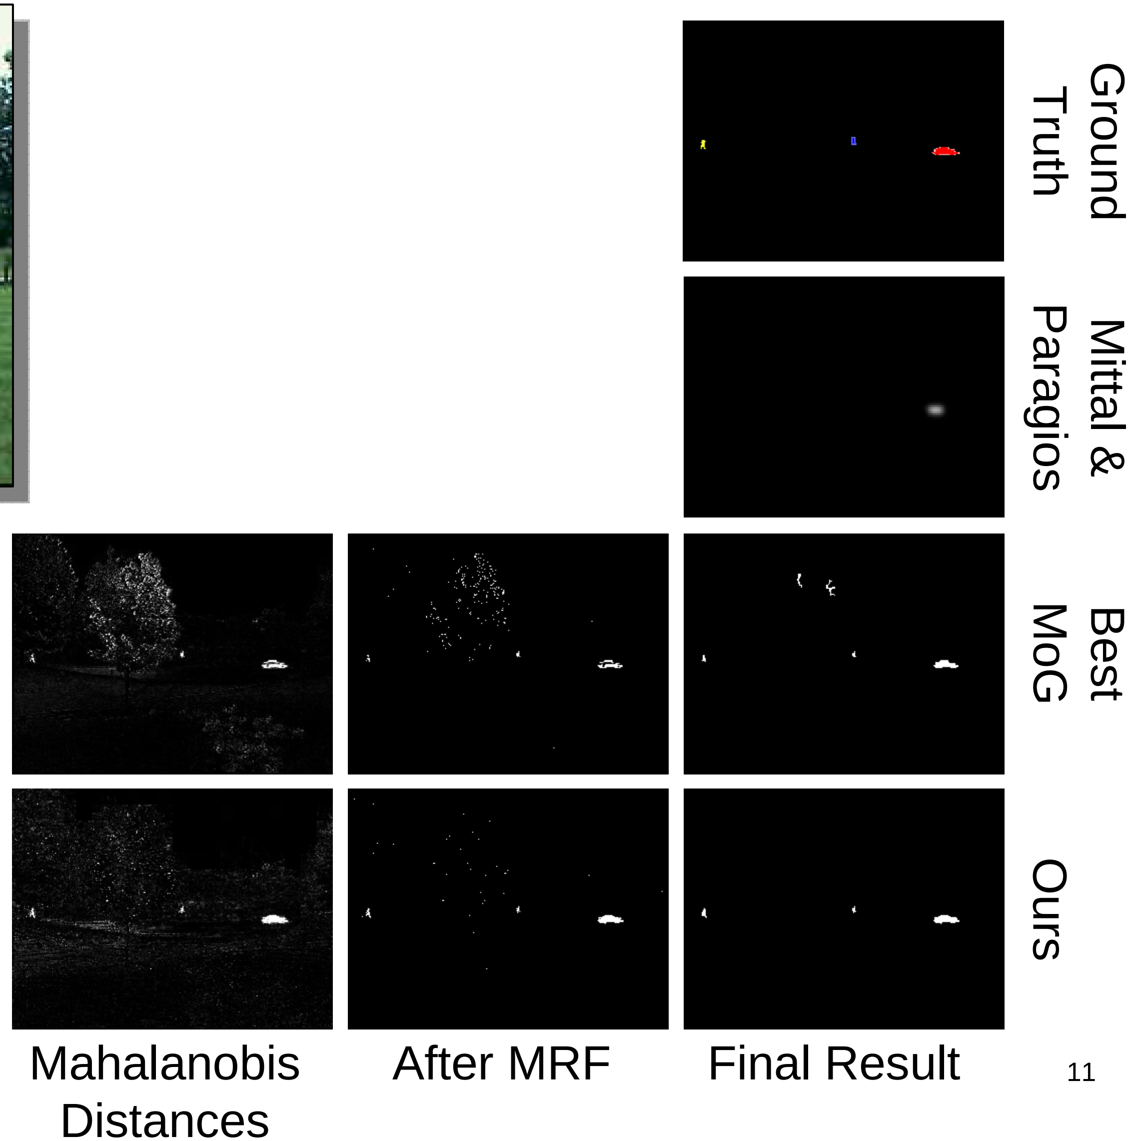

Ground Truth
Mittal & Paragios
Best MoG
Ours
Mahalanobis Distances
After MRF
Final Result

Pushing the Limits: Ducks

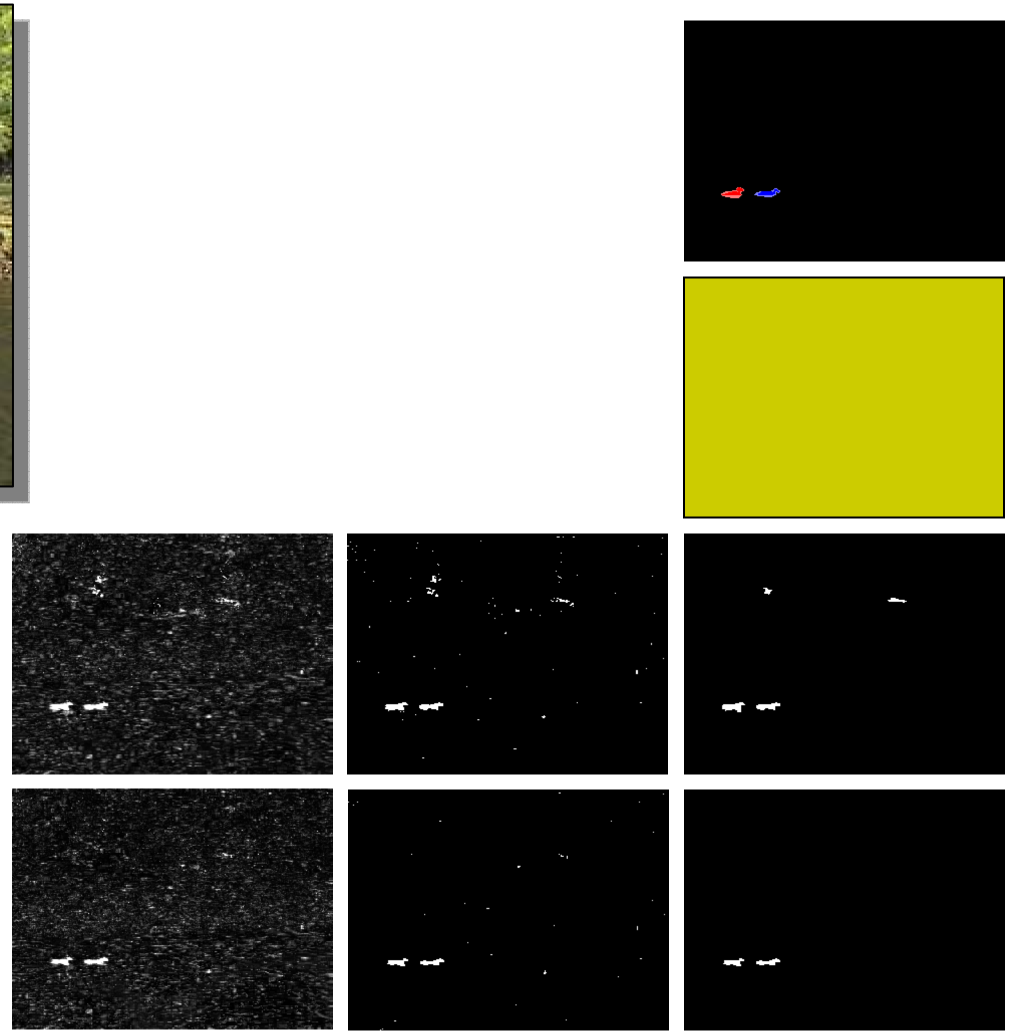

Ground Truth

Best MoG

Ours

Mahalanobis Distances

After MRF

Final Result

Pushing the Limits: Jug

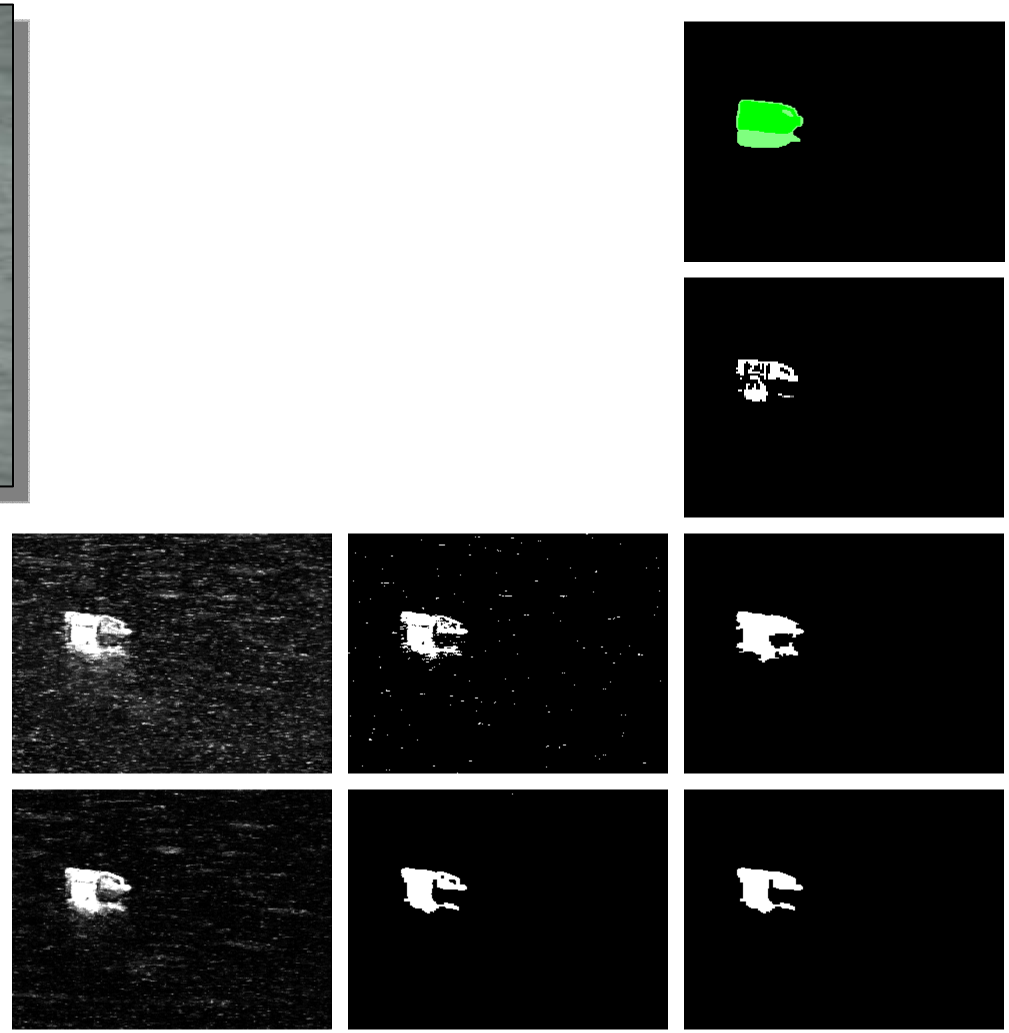

Ground Truth

Zhong & Sclaroff

Best MoG

Ours

Mahalanobis Distances | After MRF | Final Result

Parameter Sensitivity

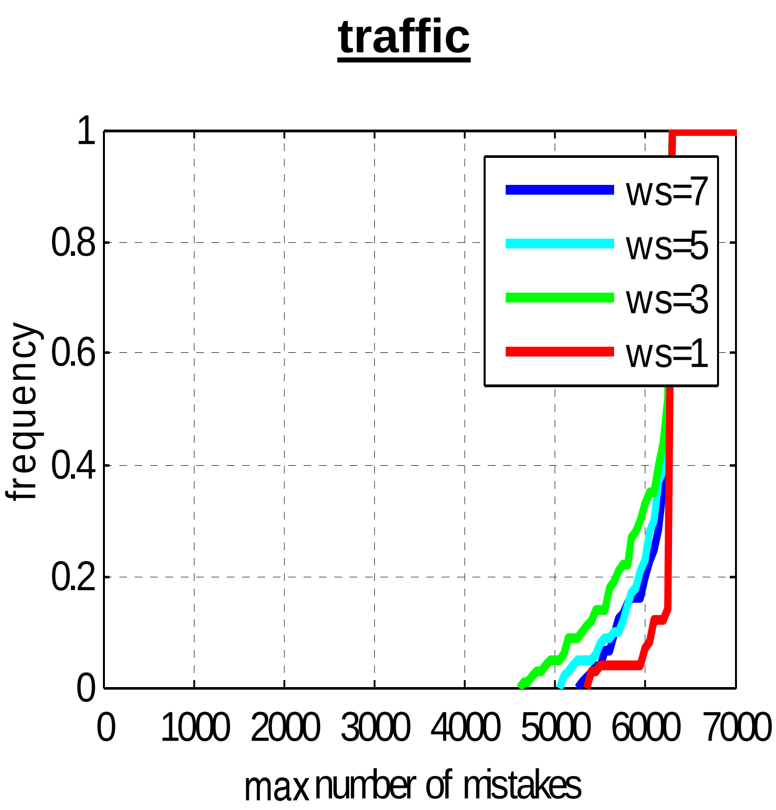

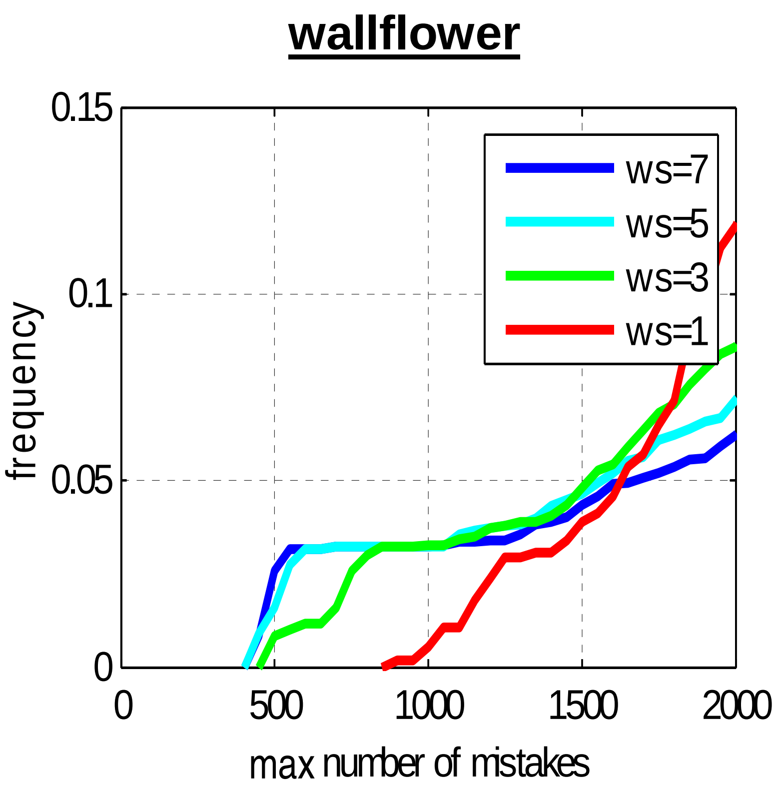

Summary

- Model representation
  - Same as MoG
- FG/BG Classifier Input
  - Add local neighborhood loop
- Update
  - Several options, including standard MoG updates
- Results
  - Reduces foreground false positives
  - Reduces need for exotic input features (e.g. optical flow)
  - Cost $\propto$ window size

Also in the Paper…

- Analysis
  - Timing experiments
  - More ROC analysis
- Implementation Details
  - MRF choice
  - Update formulas

Thank You

Questions...

Backup Slides

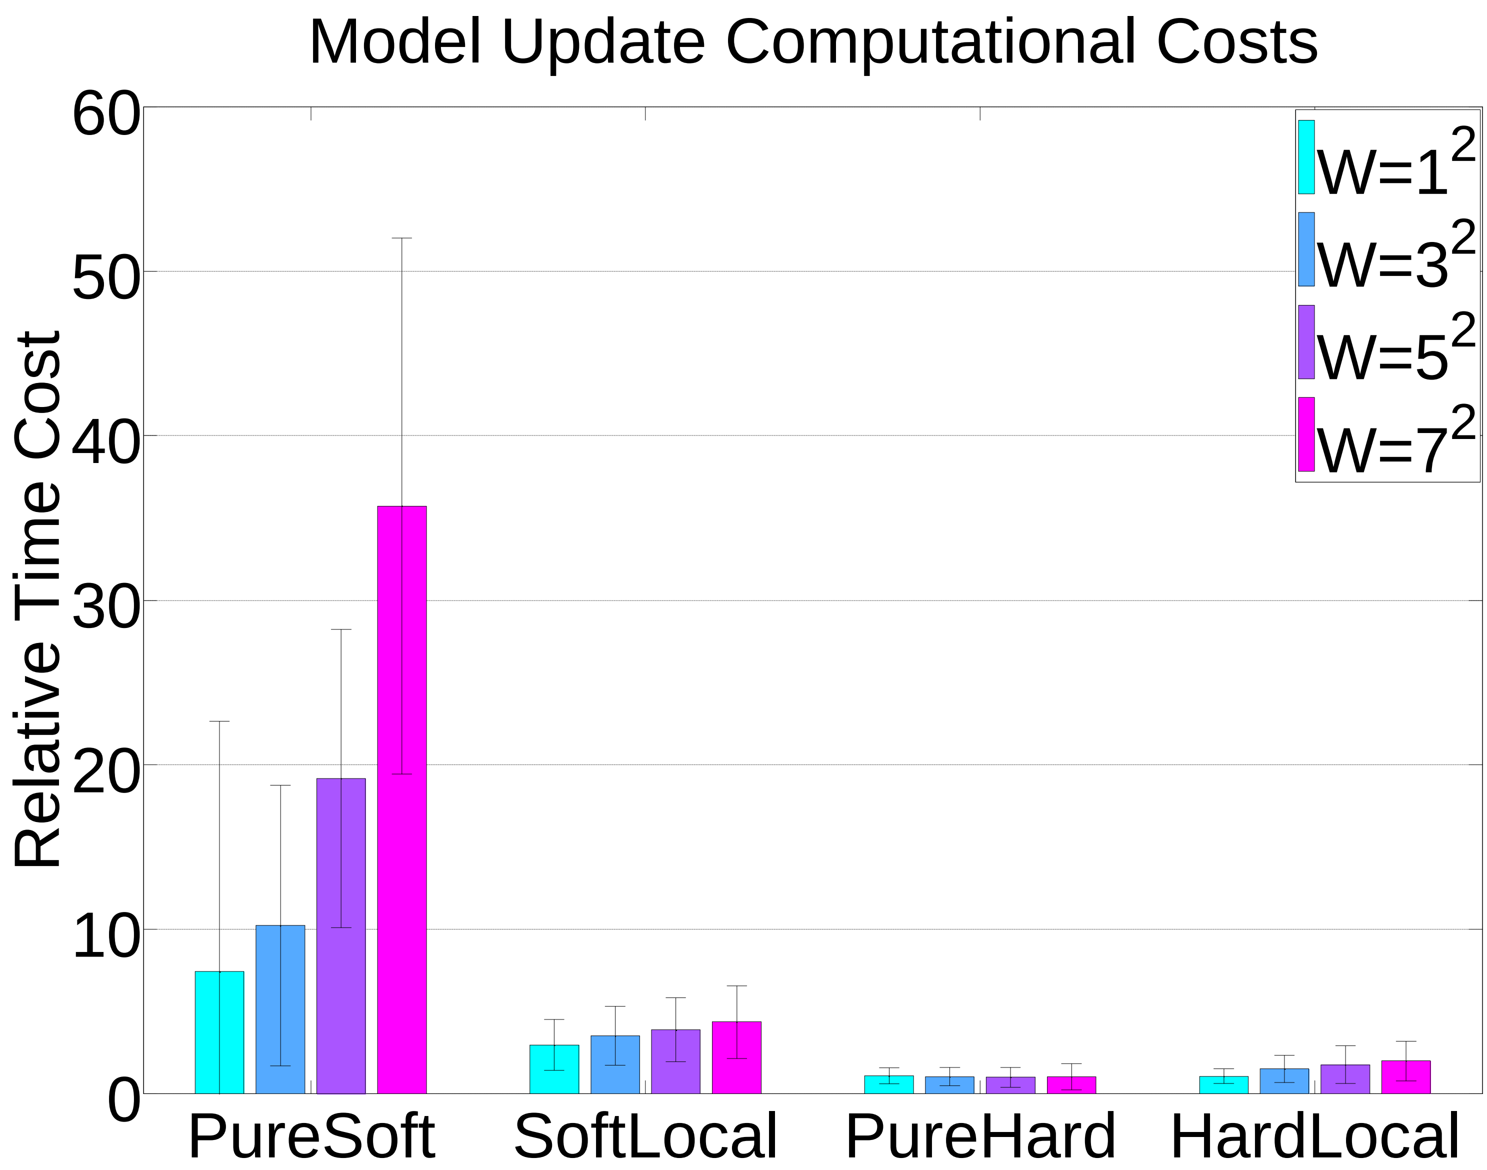

Model Update Computational Costs
Relative Time Cost
0
10
20
30
40
50
60
PureSoft
SoftLocal
PureHard
HardLocal
W=1^2
W=3^2
W=5^2
W=7^2

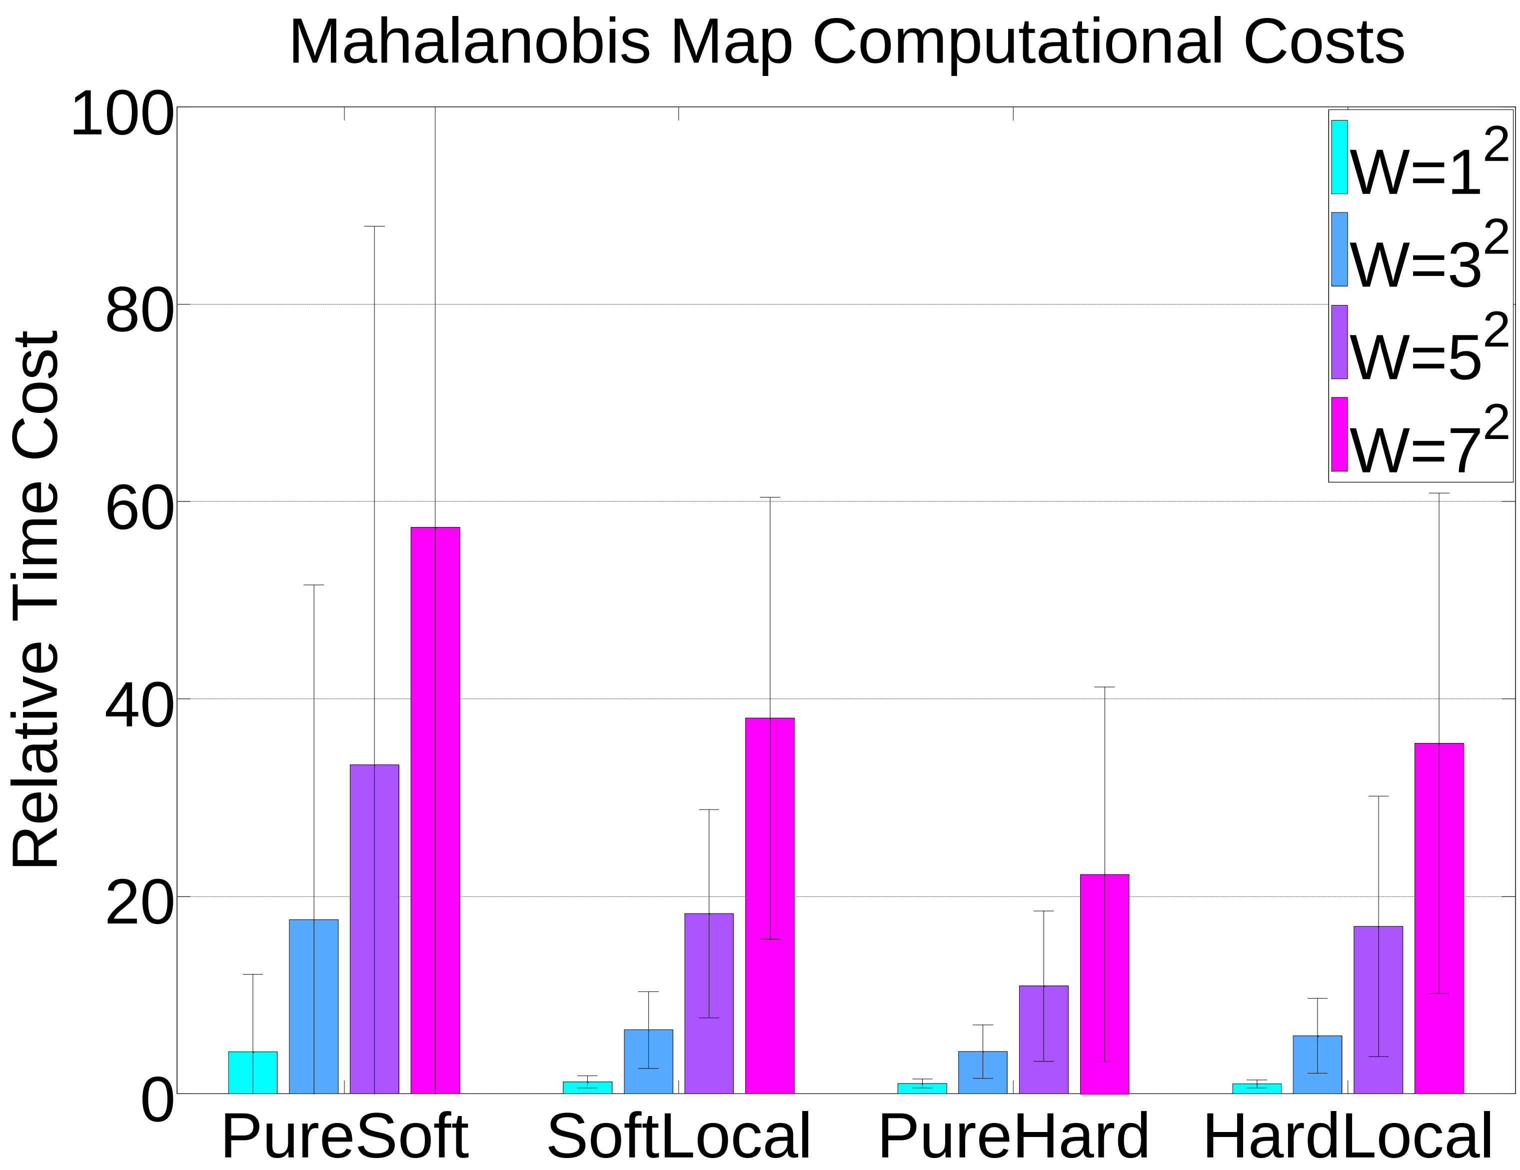
Mahalanobis Map Computational Costs
Relative Time Cost
100
80
60
40
20
0
PureSoft
SoftLocal
PureHard
HardLocal
W=1^2
W=3^2
W=5^2
W=7^2

Small Windows are Still Useful

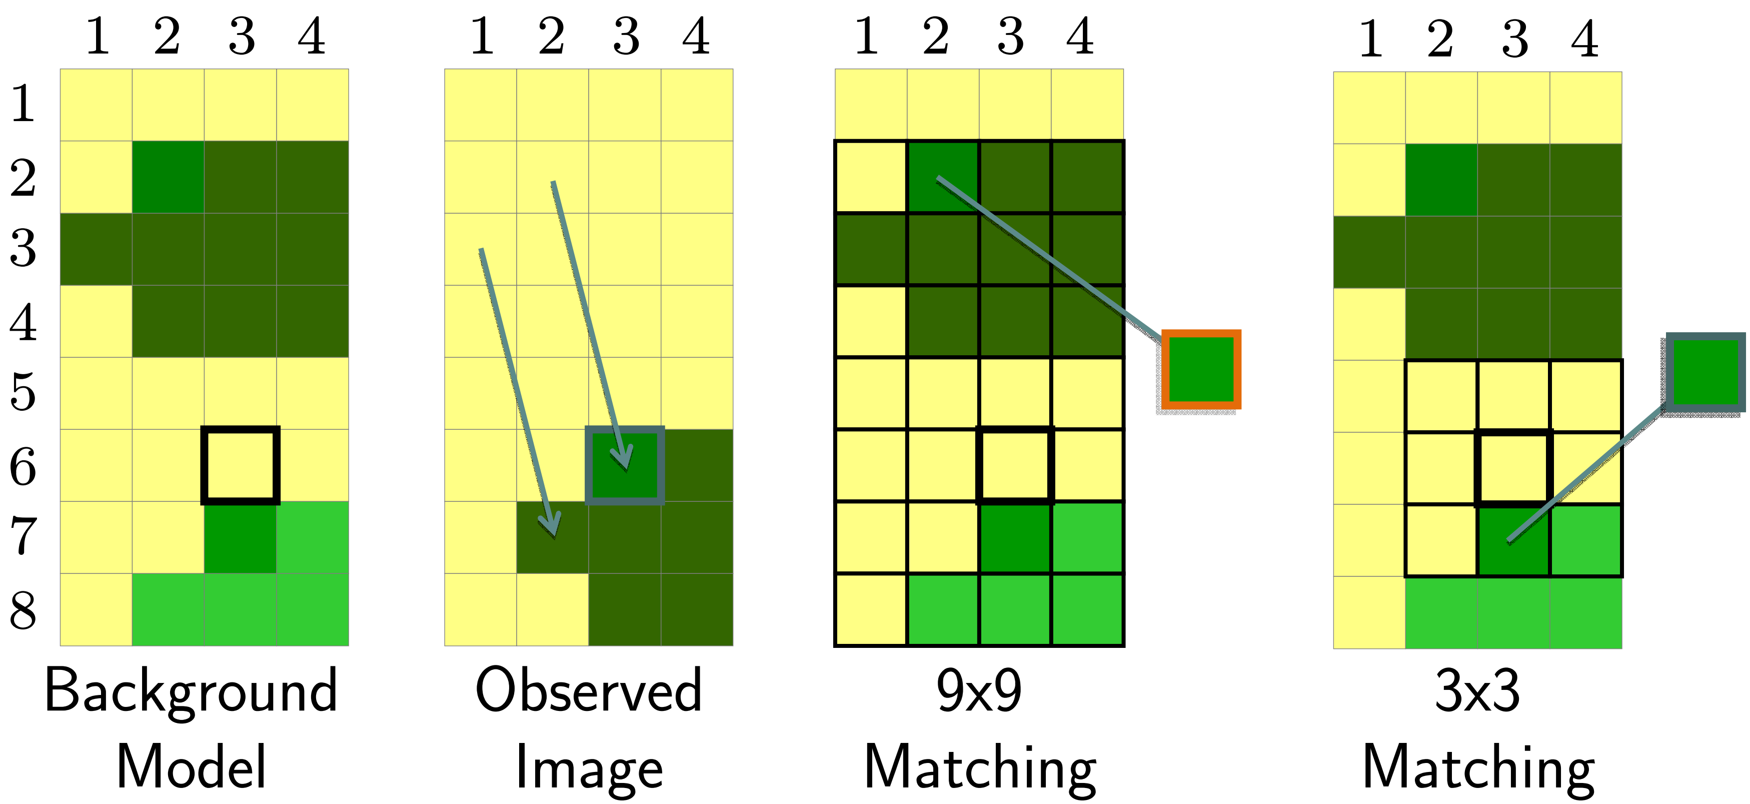

Generative Model

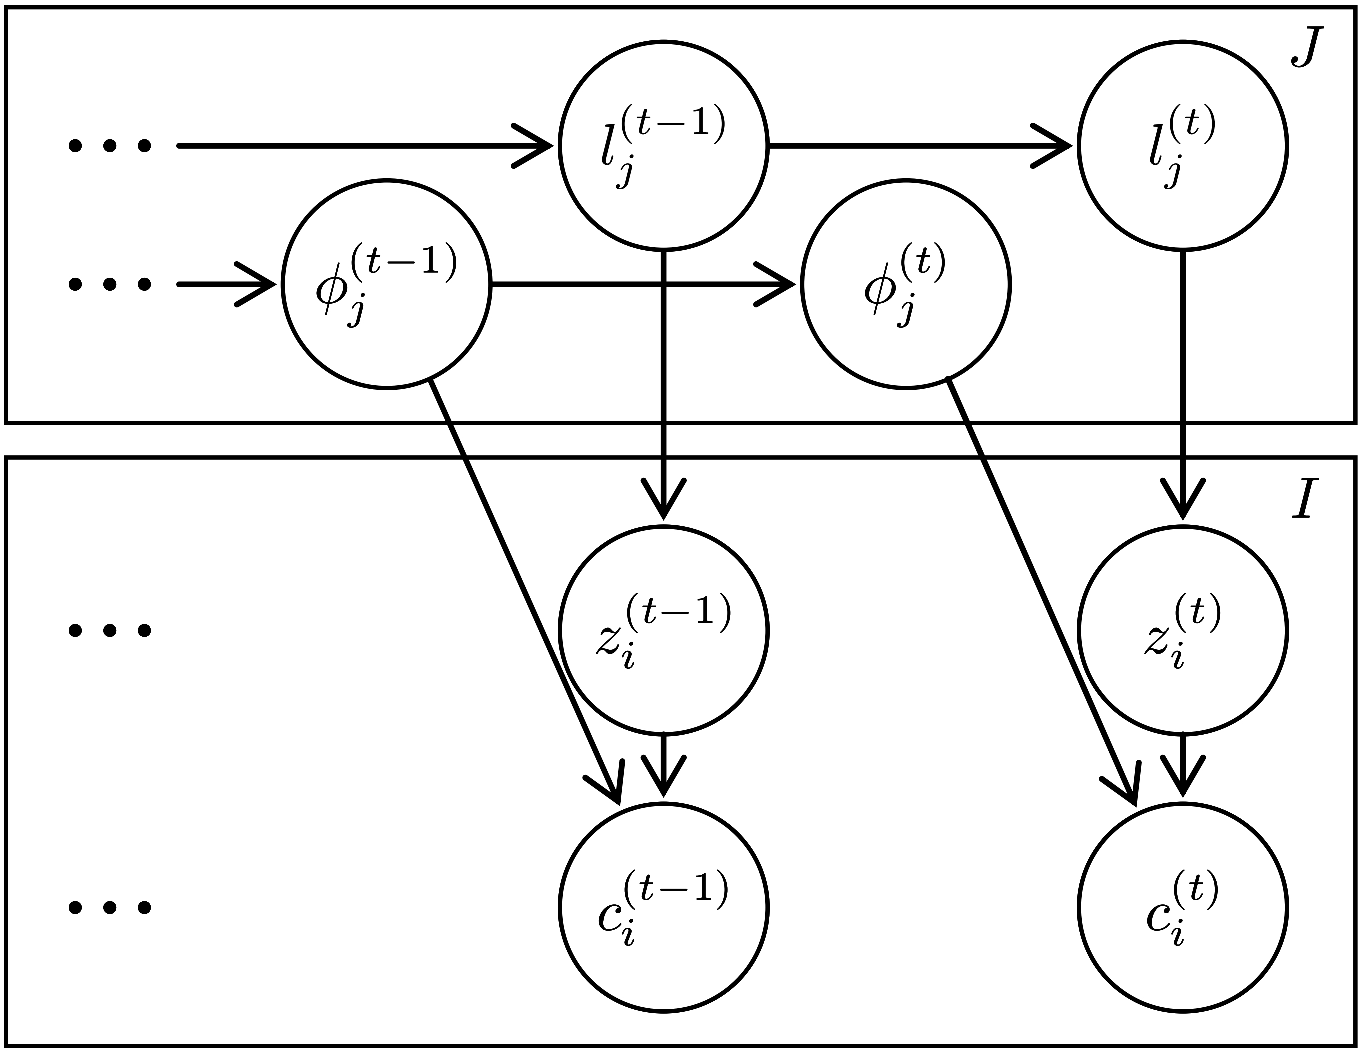

Particle locations

Particle appearances

Pixel-to-particle assignments

Pixel appearance

Independent Assignments