CRV, June 2010

MMTrack:

Max-Margin Offline Pedestrian Tracking with Multiple Cues

Bahman Yari Saeed Khanloo[1], Ferdinand Stefanus[1], Mani Ranjbar[1], Ze-Nian Li[1], Nicolas Saunier[2], Tarek Sayed[3], Greg Mori[1]

[1]School of Computing Science
Simon Fraser University

[3]Dept. of Civil Engineering
University of British Columbia

[2]Dept. of Civil, Geological and Mining Engineering
Ecole Polytechnique de Montreal

Outline

1. Overview and Previous Works
2. MMTrack
3. Learning Procedure
4. Features
5. Experimental Results

Single Object Tracking

How to combine the features?

Prior Works

- Online selection of discriminative tracking features [Collins et al, PAMI05]
- MILTrack [Babenko et al, CVPR09]

○ Single type of features
○ Appearance model

- Tracking with multiple observers [Stenger et al, CVPR09]

○ Combine complete trackers
○ Independent learning of relative weightings

Outline

1. Overview and Previous Works
2. MMTrack
3. Learning Procedure
4. Features
5. Experimental Results

MMTrack

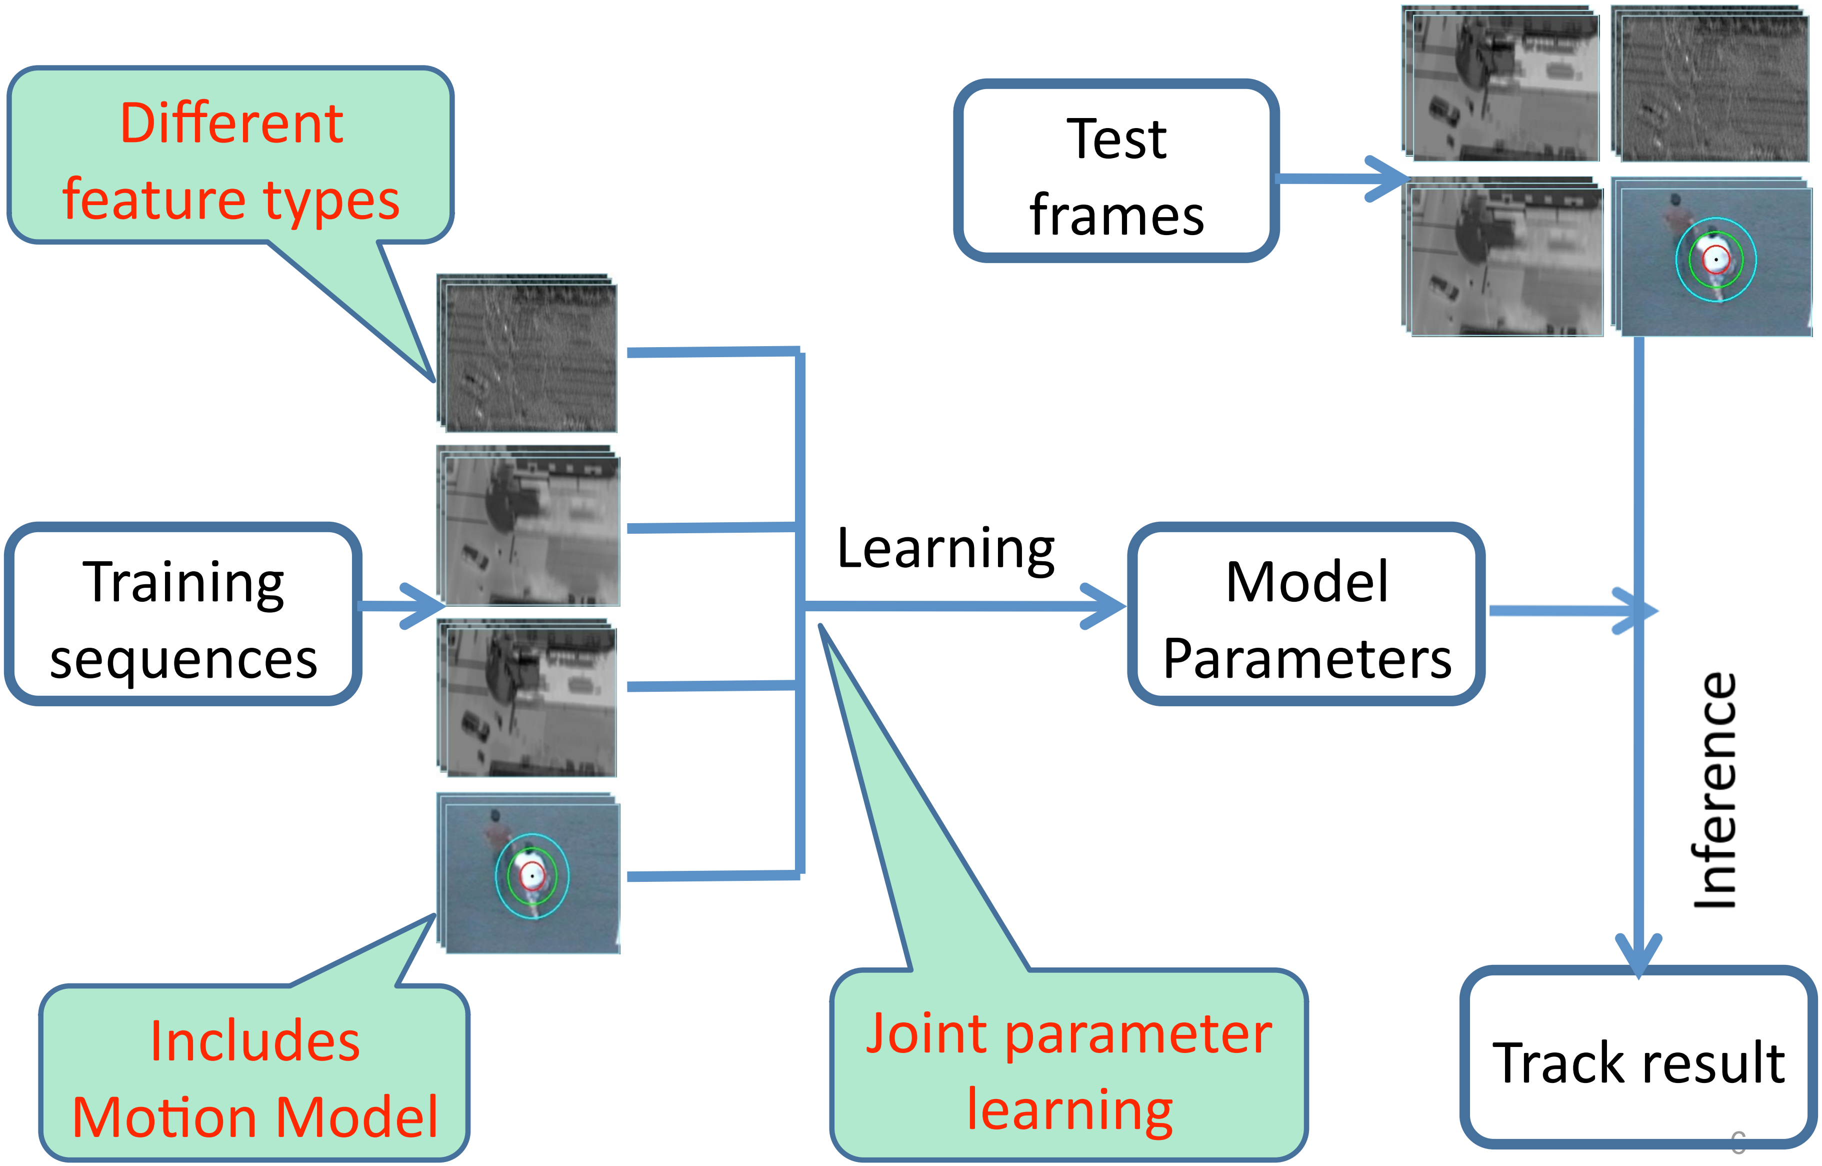
Different feature types
Test frames
Training sequences
Learning
Model Parameters
Inference
Includes Motion Model
Joint parameter learning
Track result

Learned Appearance Parameters

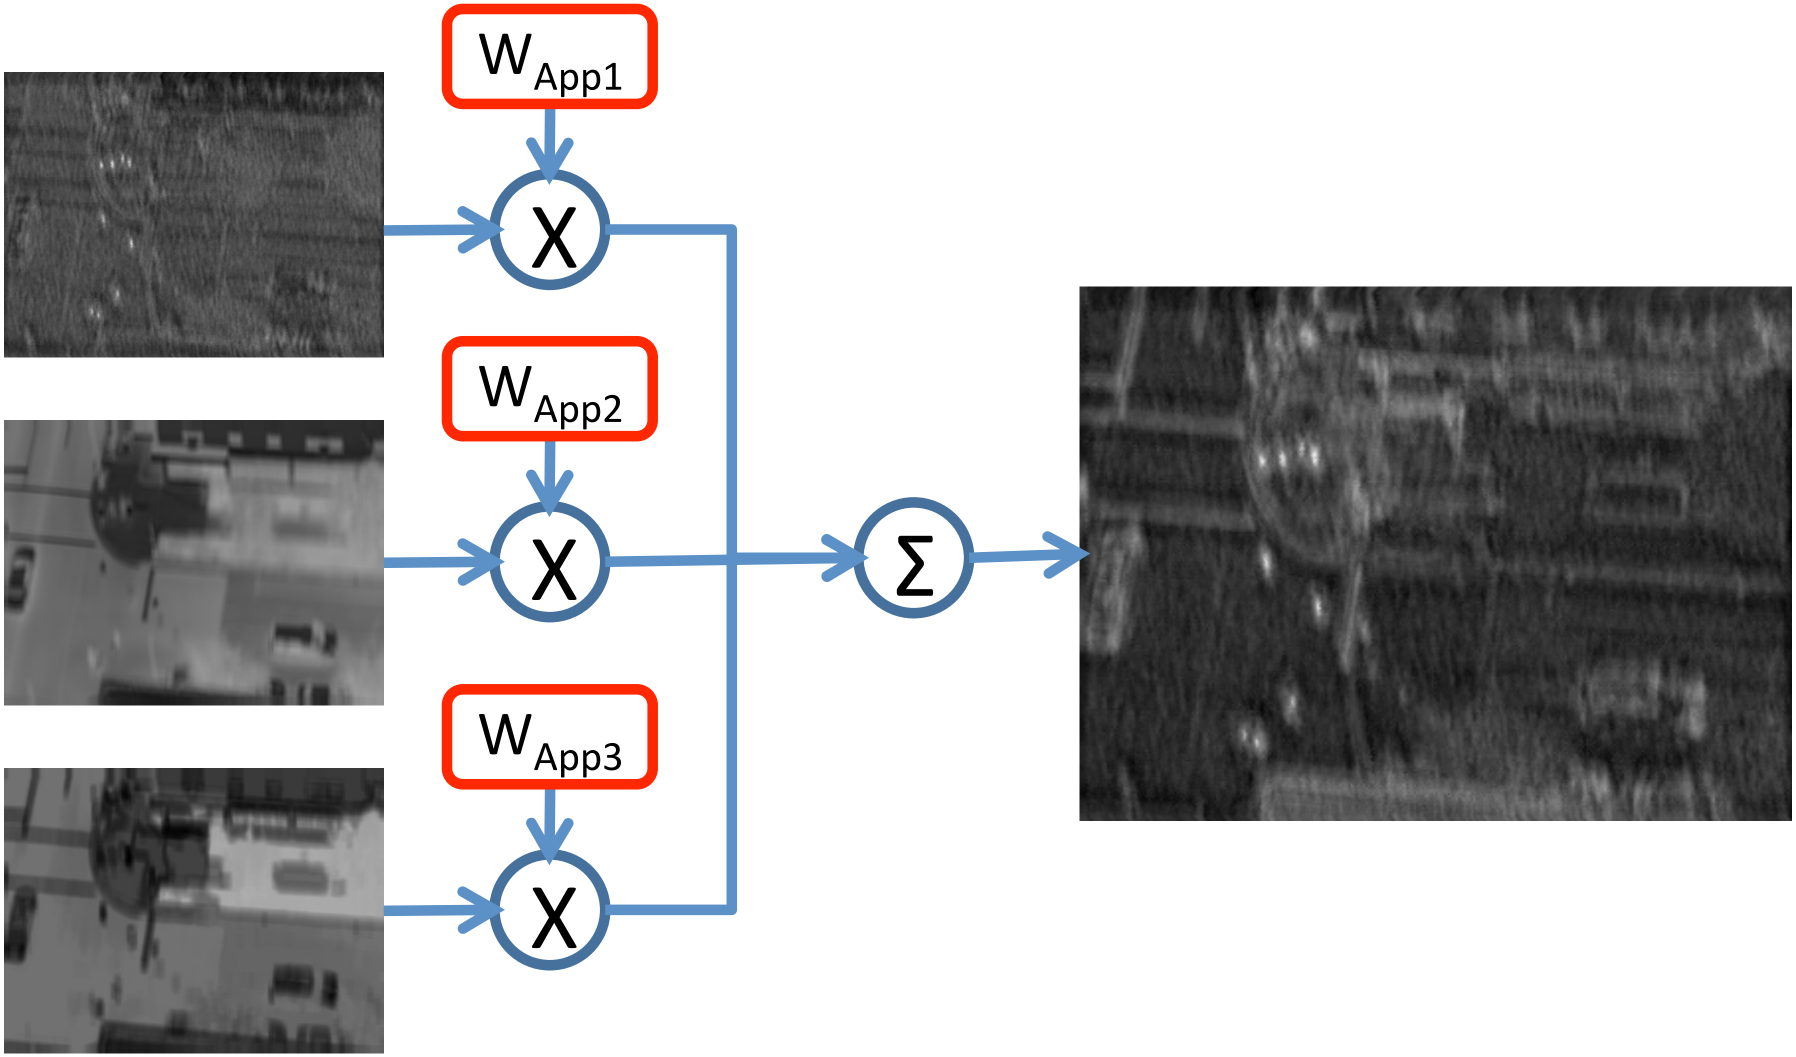

Learned Motion Parameters

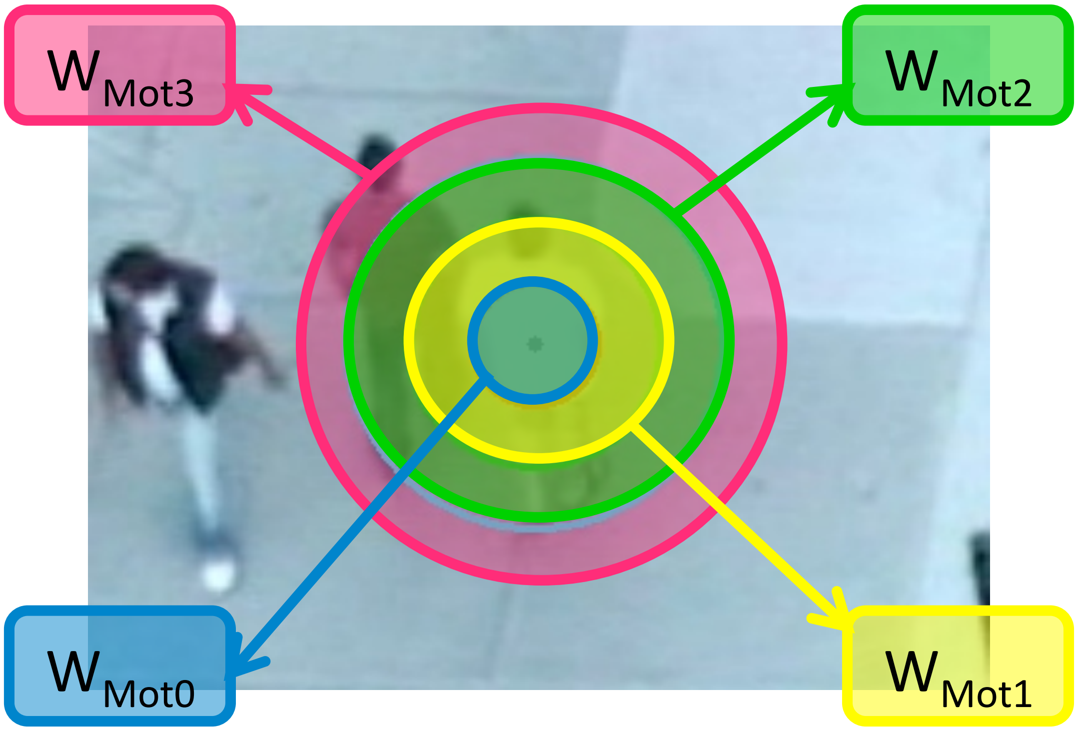

First Order Model

Trajectory Scoring

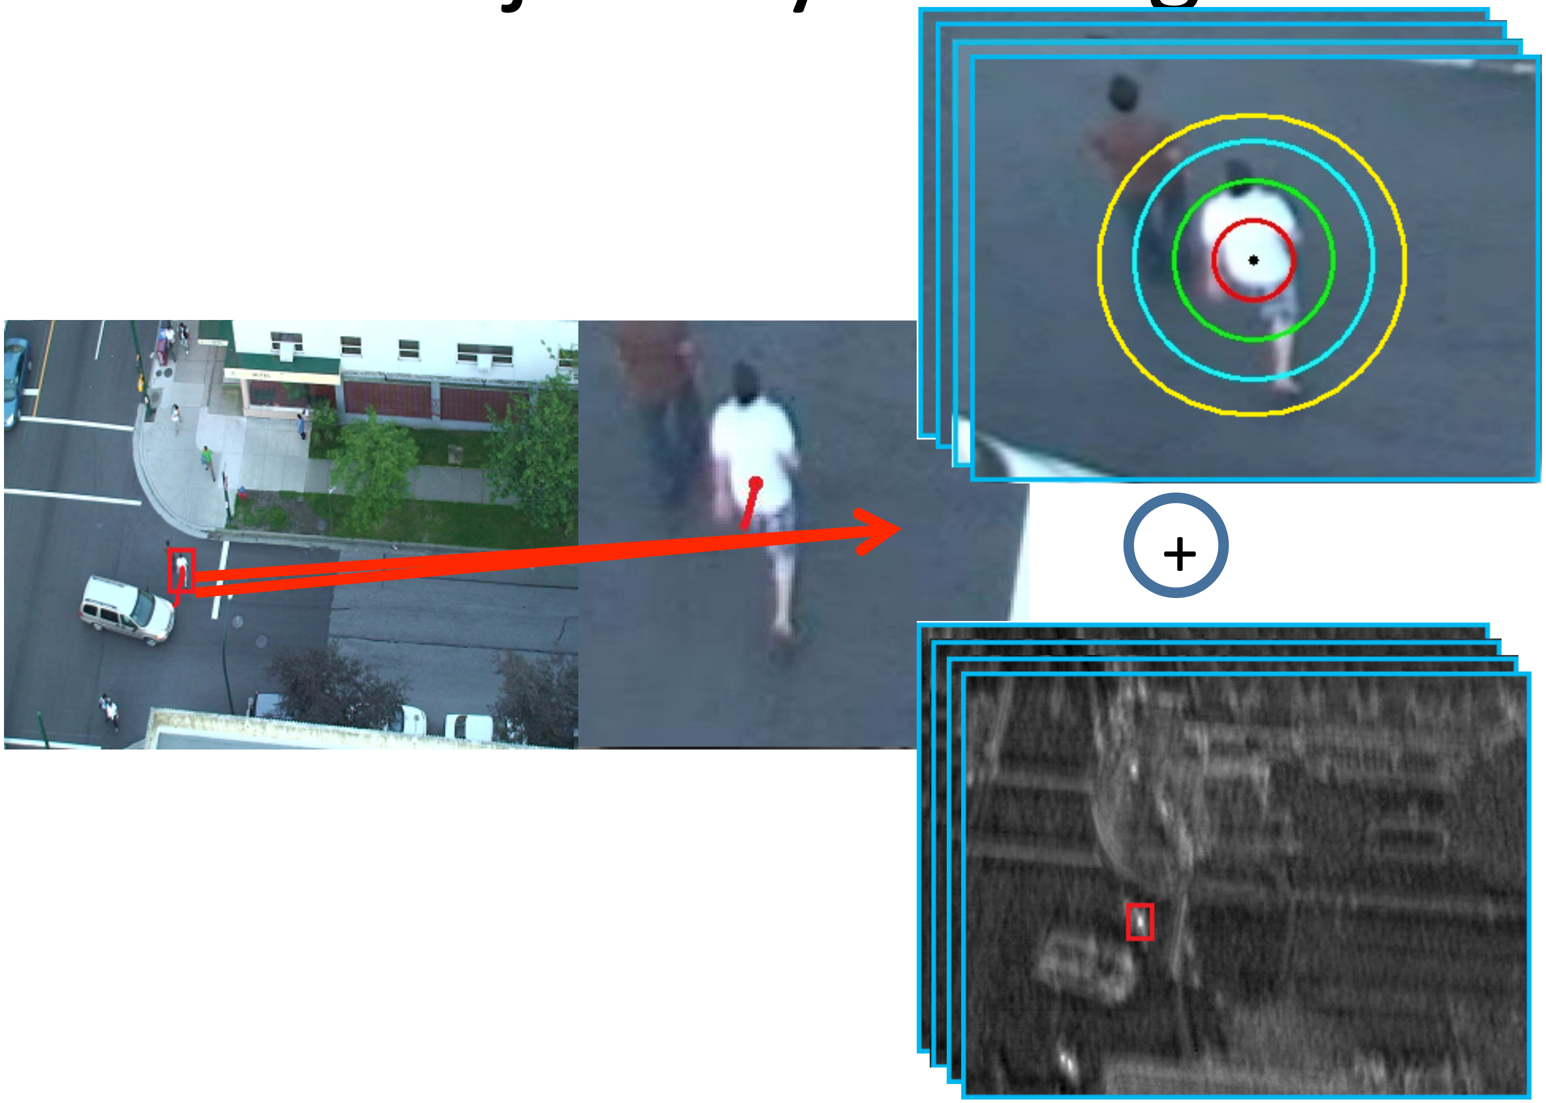

Outline

1. Overview and Previous Works
2. MMTrack
3. Learning Procedure
4. Features
5. Experimental Results

Structural SVM Learning

- Objective: find parameters that produce ‘good’ tracks
- How to measure ‘goodness’ of a track? Loss function $\Delta(y_i, y)$
- We use total Euclidean distance

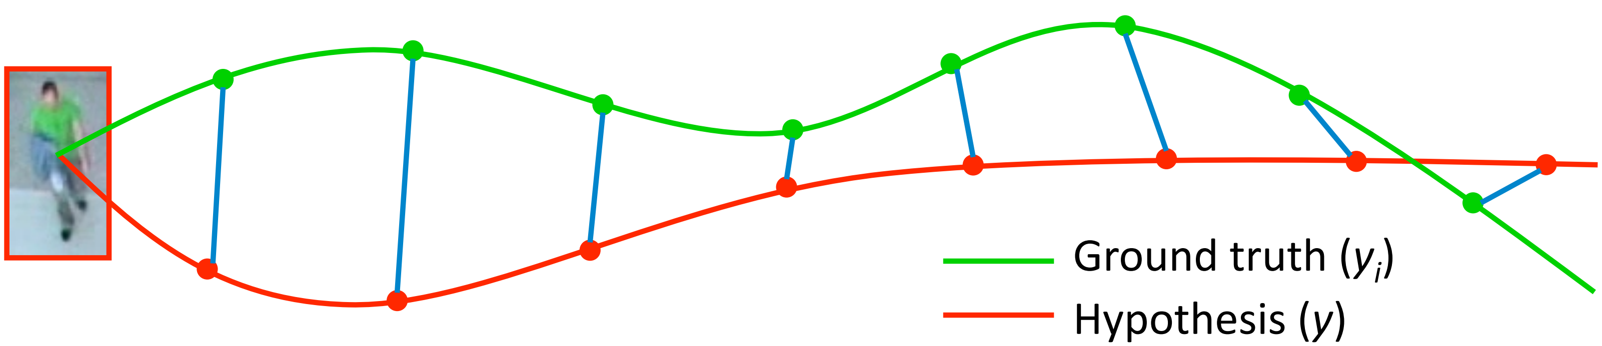

Structural SVM Formulation

$$\min_{\boldsymbol{w},\boldsymbol{\varepsilon}} \boxed{\frac{1}{2}\|\boldsymbol{w}\|^2} + \frac{C}{N}\sum_{i=1}^{N} \boxed{\varepsilon_i}, \mathrm{s.t.} \forall_i, \varepsilon_i \geq 0$$

Regularizer

Slack variables

$$\forall i, \boxed{\forall y \backslash y_i} : \boxed{\langle \phi(x_n, y_i), w \rangle} - \boxed{\langle \phi(x_n, y), w \rangle} \geq \Delta(y_i, y) - \varepsilon_i$$

Score for ground truth y_i

Exponentially large number of constraints!

Score for hypothesized track y

[Tsochantaridis et al, ICML04]

Structural SVM Learning

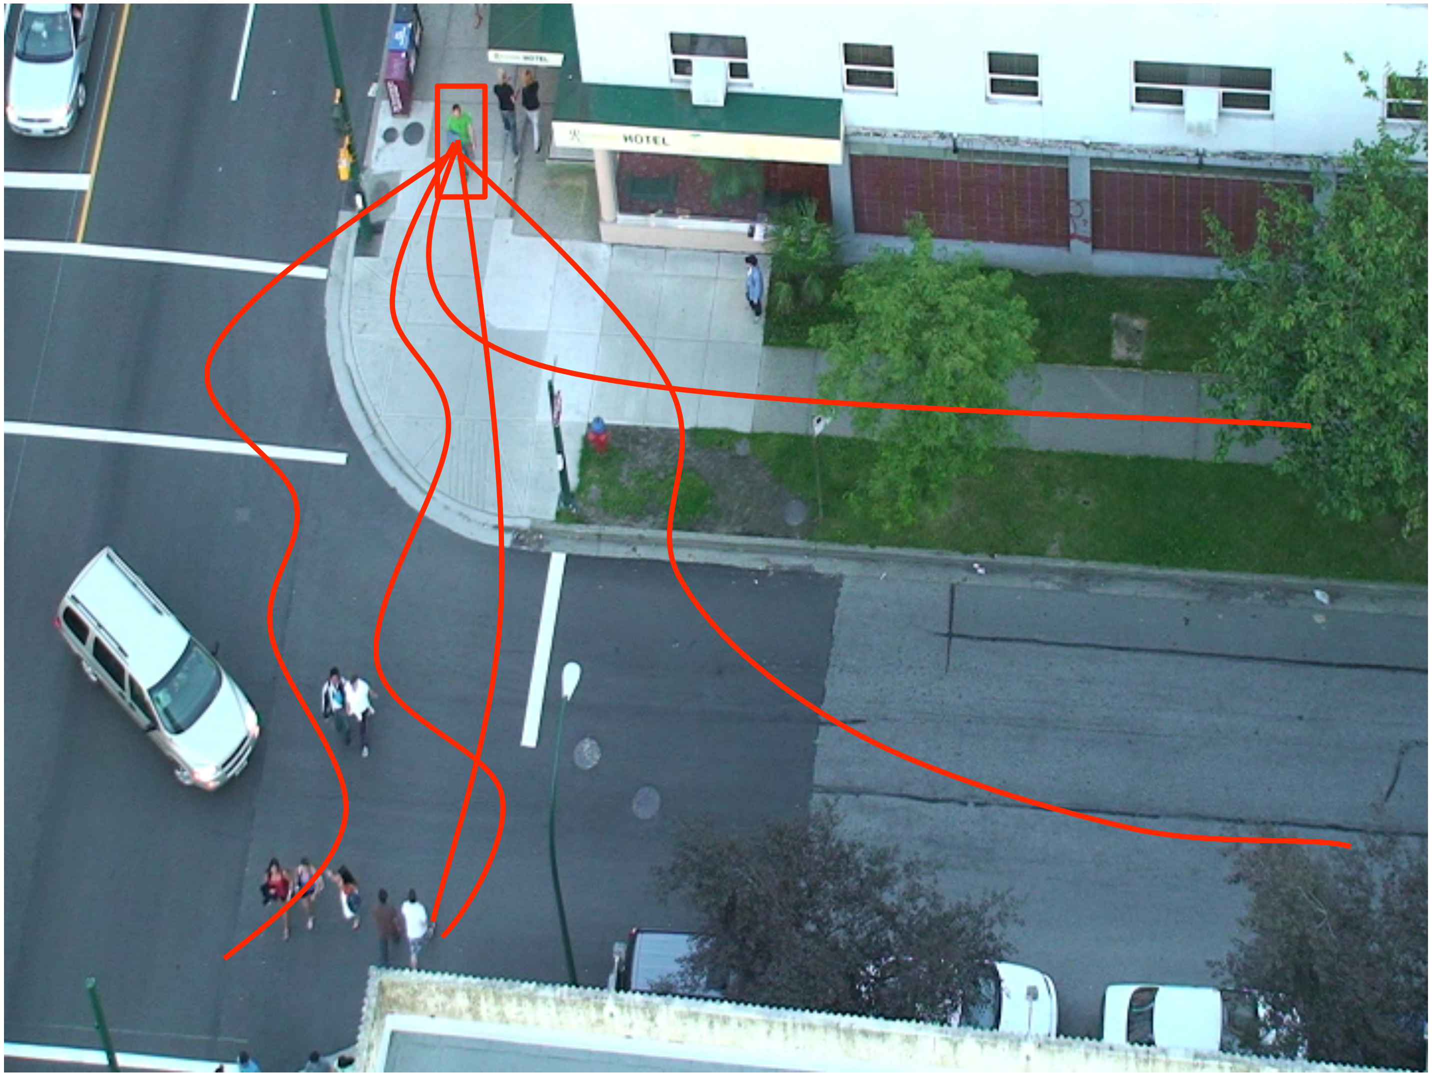

Structural SVM Learning

$$\min_{w,\varepsilon} \boxed{\frac{1}{2}\|w\|^2} + \frac{C}{n}\sum_{i=1}^{N} \boxed{\varepsilon_i}, \text{s.t.} \forall_i, \varepsilon_i \geq 0$$

Regularizer

Slack variables

$$\forall i, \boxed{\forall y \backslash y_i} : \langle \phi(x_n, y_i), w \rangle - \boxed{\langle \phi(x_n, y), w \rangle \geq \Delta(y_i, y)} - \varepsilon_i$$

Large number of constraints

Solve for ‘most violated constraint’:

$$\forall i, \forall y \backslash y_i : \langle \phi(x_n, y_i), w \rangle \geq \Delta(y_i, y) + \langle \phi(x_n, y), w \rangle - \varepsilon_i$$

$$\operatorname*{argmax}_{y} \Delta(y_i, y) + \langle \phi(x_n, y), w \rangle$$

[Tsochantaridis et al, ICML04]

Structural SVM Learning

Structural SVM Learning

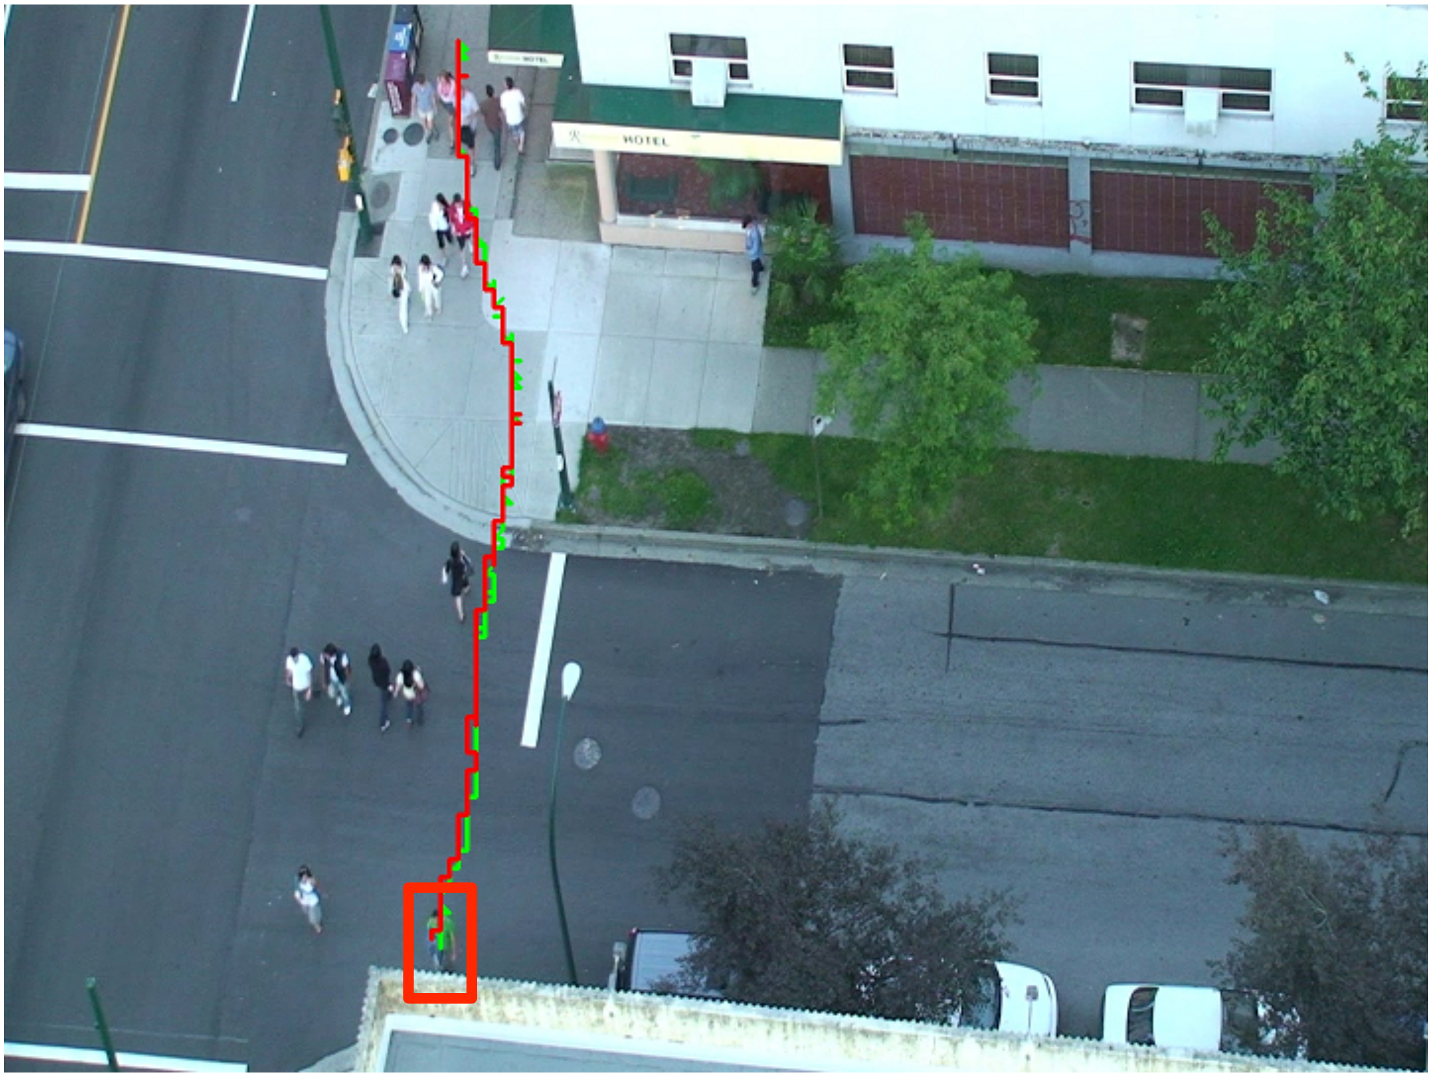

Outline

1. Overview and Previous Works
2. MMTrack
3. Learning Procedure
4. Features
5. Experimental Results

Features – HOG Score Map

- Histogram of Oriented Gradient [Dalal&Triggs, CVPR05]

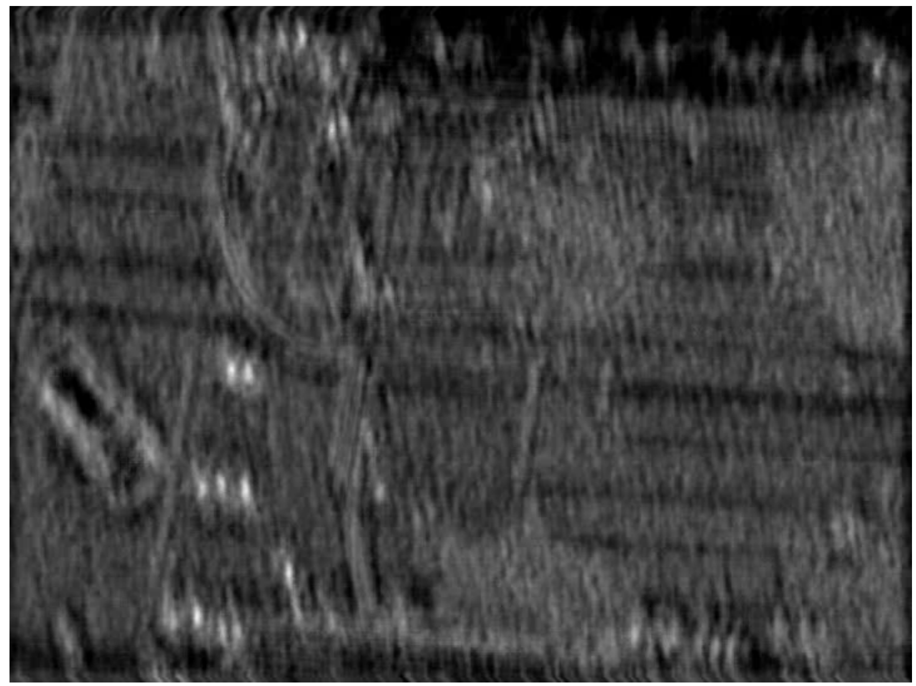

Features – Appearance Templates

- Object templates from initial and previous frames
- Compute sum of absolute difference at nearby locations

Features – Color Histogram Distance

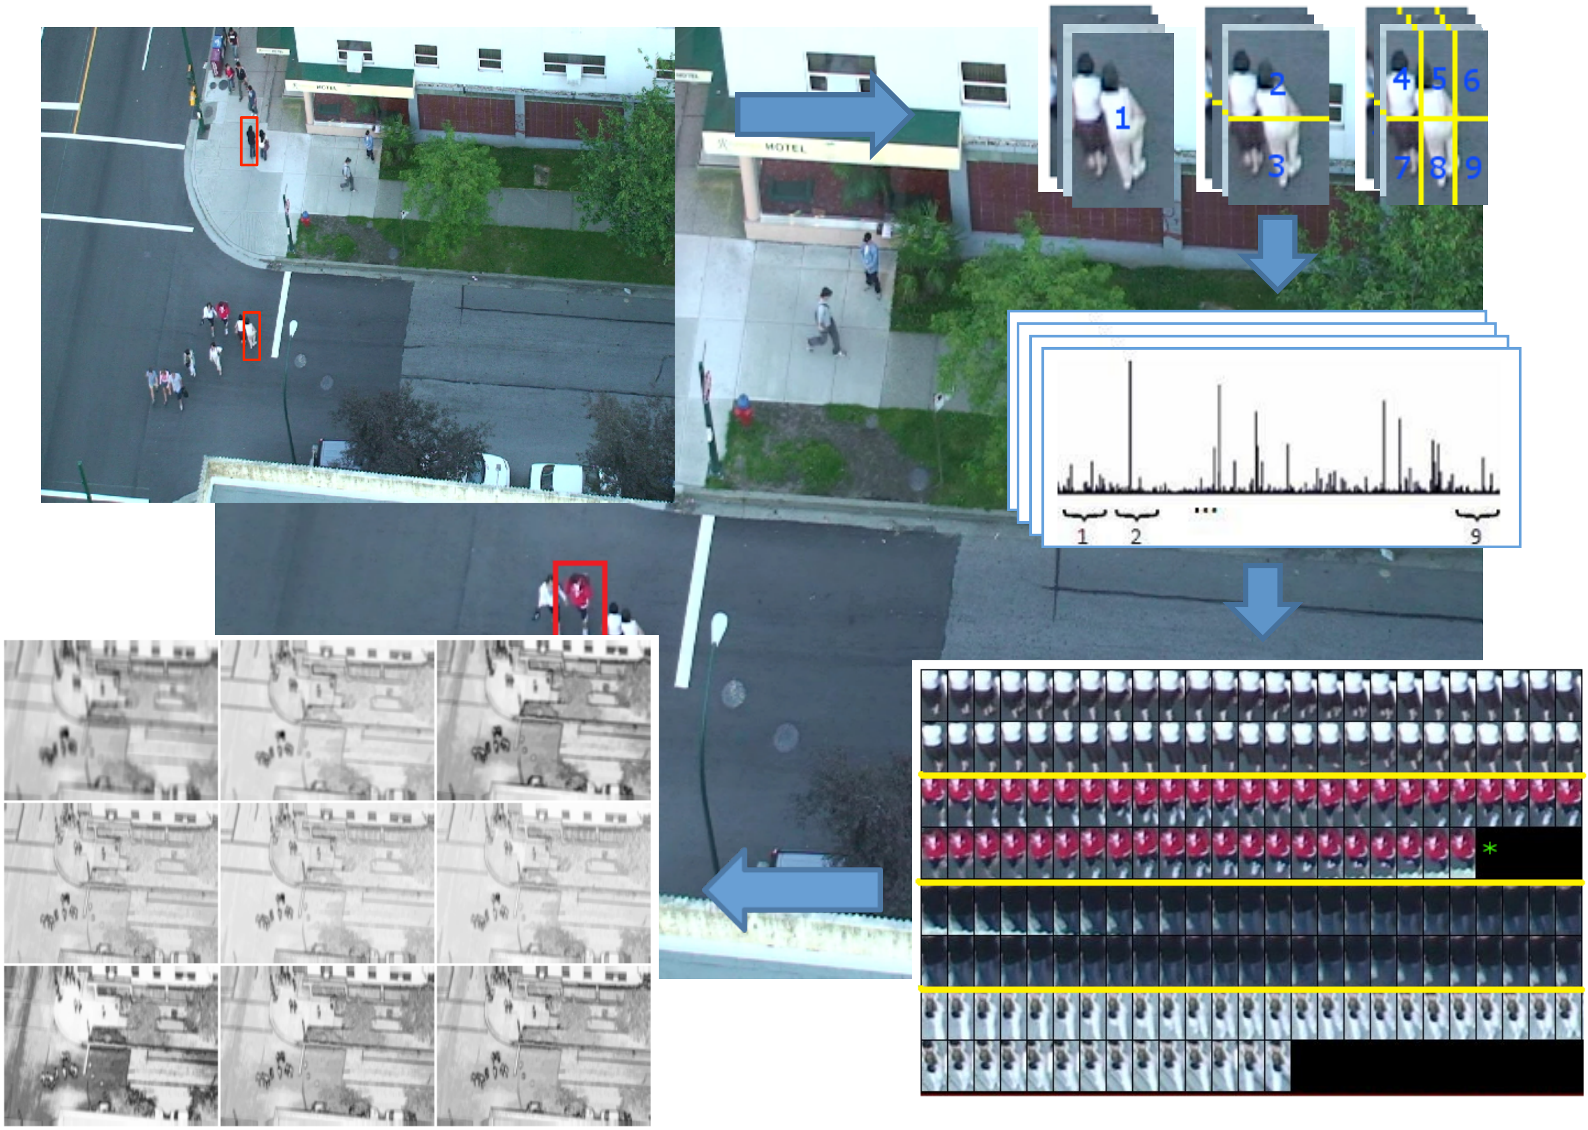

Features

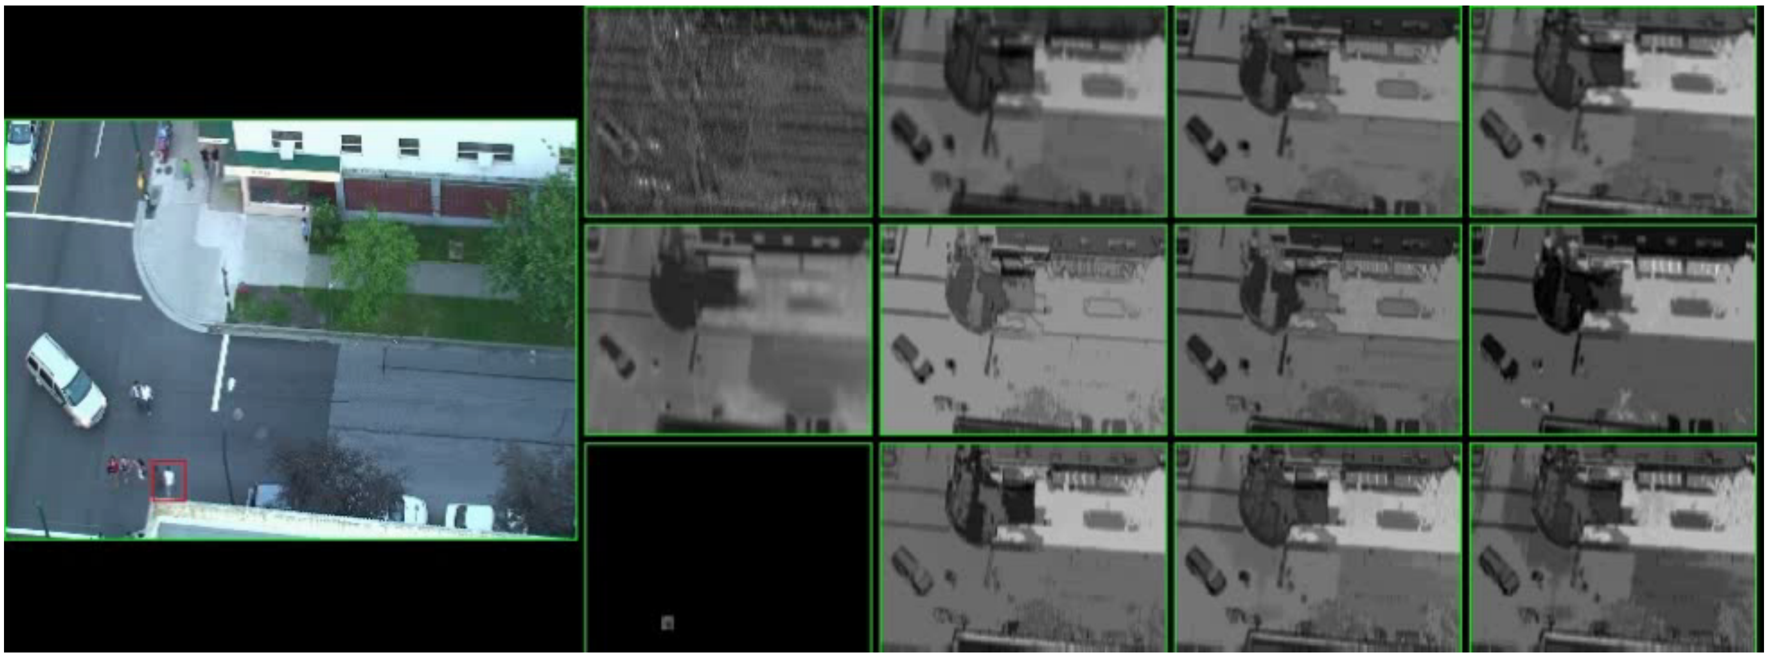

12 features: HOG + 9 histogram difference + 2 appearance template

Outline

1. Overview and Previous Works
2. MMTrack
3. Learning Procedure
4. Features
5. Experimental Results

UBC Dataset

Experimental Results

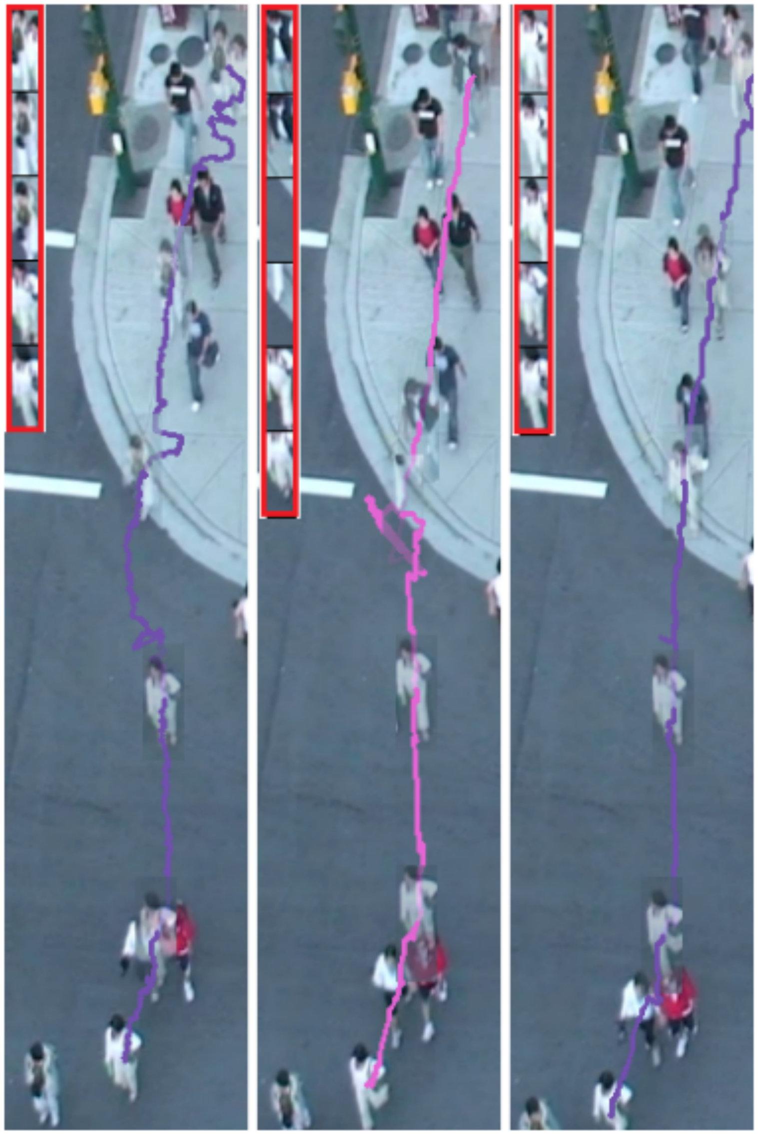

- Left: HOG + 9 histogram
- Middle: HOG + appearance templates
- Right: HOG + 9 histogram + appearance templates

Experimental Results

Tracker	#CDT (max=22)	Avg CT	Avg Error
MMTrack: all features	21	0.66	7.01
MMTrack: HOG + Hist	10	0.47	14.40
MMTrack: HOG + Template	14	0.52	22.24
MILTrack [1]	19	0.61	19.87
Collins-Liu [2]	14	0.54	21.24

[1]. B. Babenko et al, CVPR09
[2]. R. Collins et al, PAMI05

More Experimental Results

Structural SVM Learning

$$\underset{y}{\operatorname{argmax}}\, \phi(y, x^m; w) \approx y^m$$

- How to compare y with y^m? Loss function
- We use total Euclidean distance

Structural SVM Learning

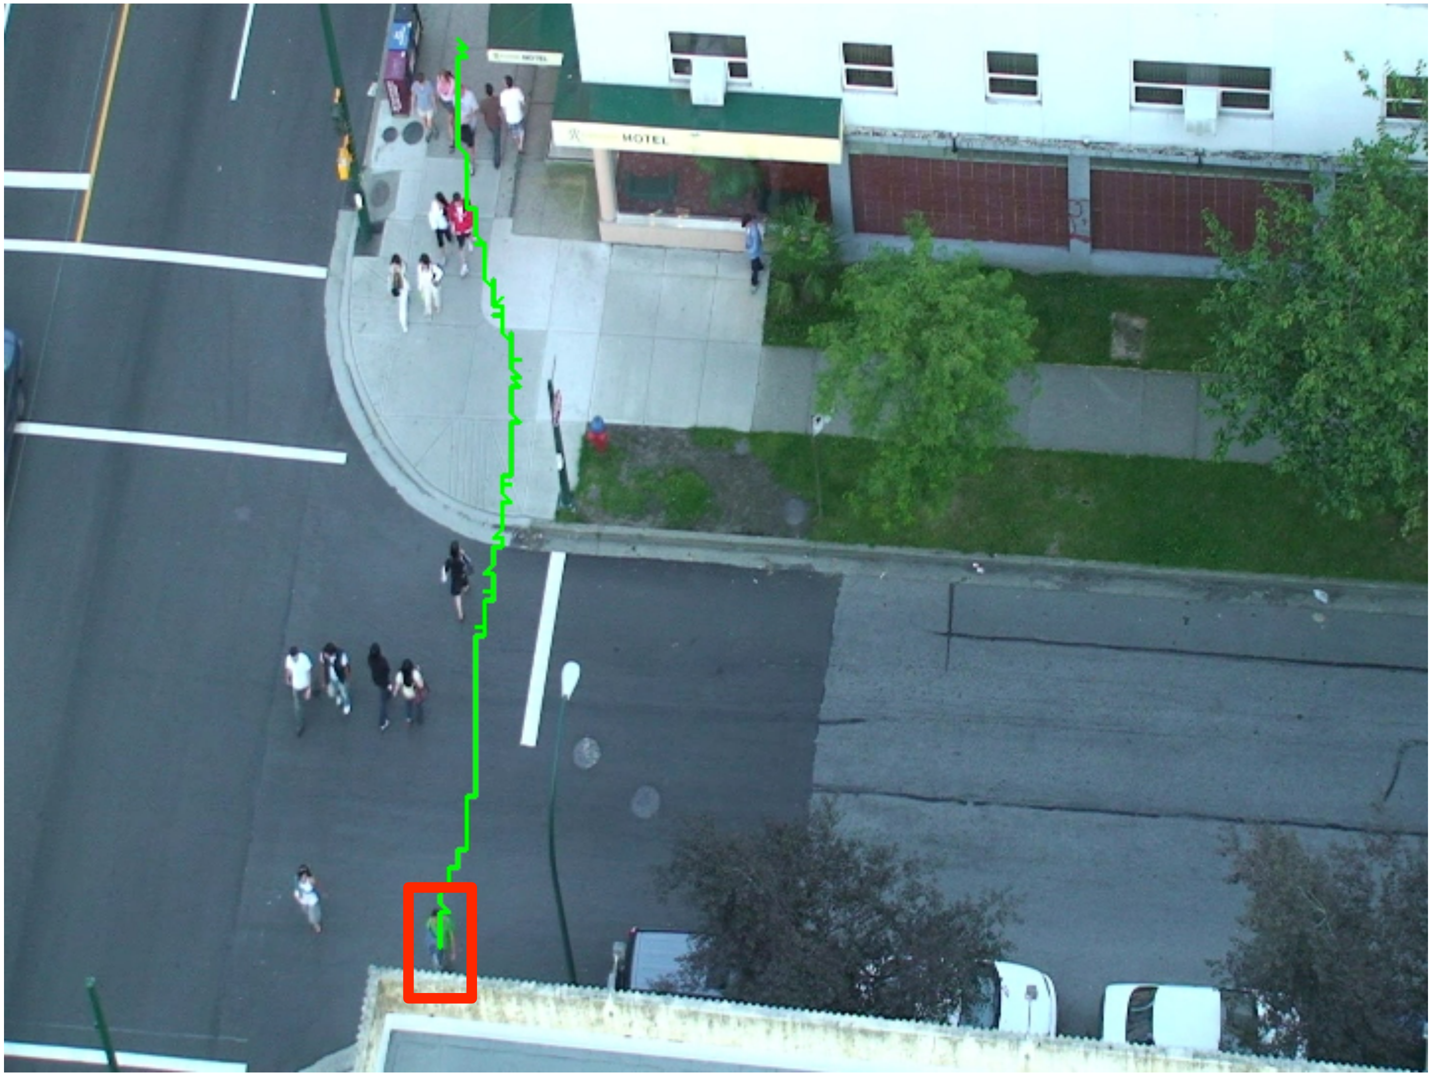

MMTrack

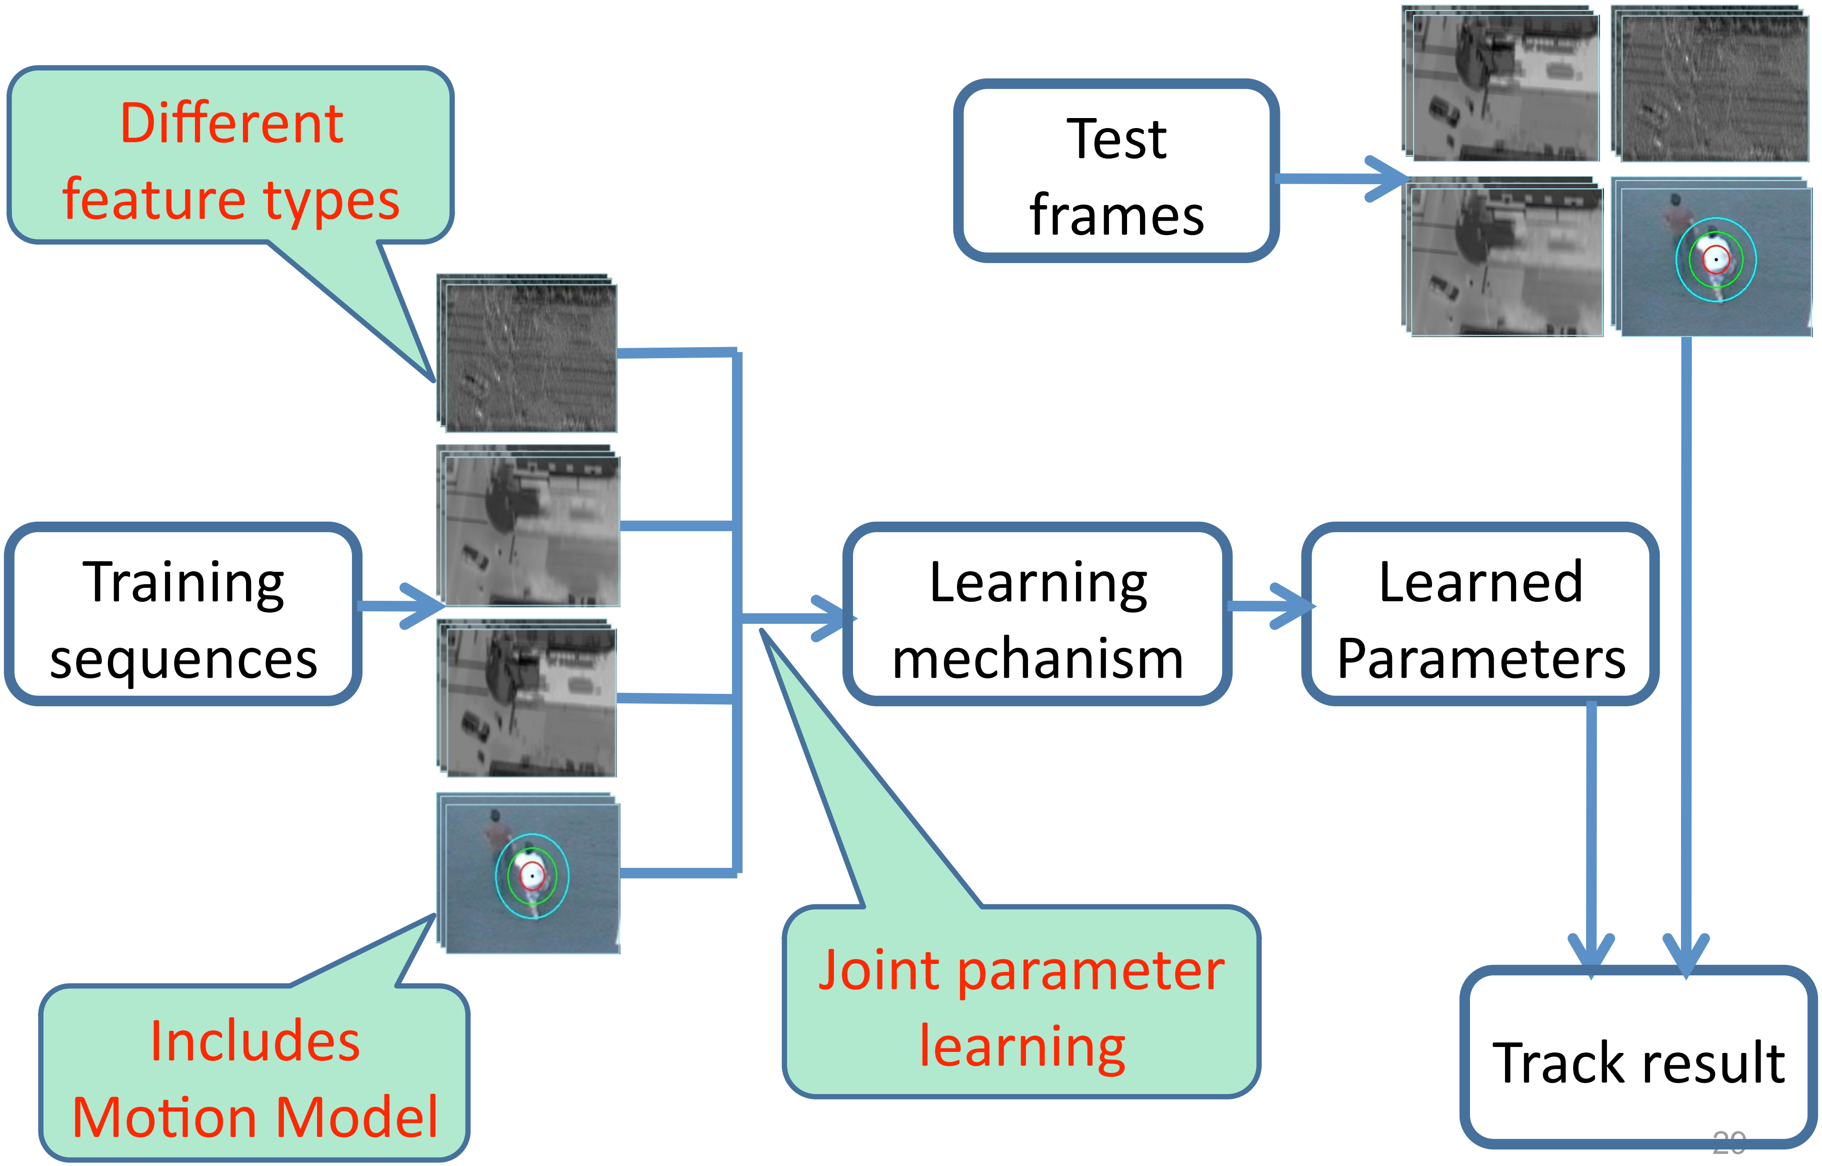
Different feature types
Test frames
Training sequences
Learning mechanism
Learned Parameters
Includes Motion Model
Joint parameter learning
Track result

Structural SVM Framework

$$\begin{cases} S_{gt(i)} - S_N \geq \Delta(gt(i), N) \\ \text{Maximize the} \end{cases}$$

Margin Criterion

$$\forall i, \forall y \backslash y_i : \langle \phi(x_n, y_i), w \rangle - \langle \phi(x_n, y), w \rangle \geq \Delta(y_i, y) - \varepsilon_i$$

Score for ground truth y_i

Exponentially large number of constraints!

Score for hypothesized track y

Slack variables

[Tsochantaridis et al, ICML04]