

Action recognition in videos III

Cordelia Schmid
INRIA Grenoble



Action recognition - goal

- Short actions, i.e. answer phone, shake hands



answer phone



hand shake

Action recognition - goal

- Activities/events, i.e. making a sandwich, doing homework

Making sandwich



Doing homework



TrecVid Multi-media event detection dataset

Action recognition - goal

- Activities/events, i.e. birthday party, parade

Birthday party



Parade



TrecVid Multi-media event detection dataset

Action recognition - tasks

- Action classification: assigning an action label to a video clip



Making sandwich: present
Feeding animal: not present
...


Action recognition - tasks

- Action classification: assigning an action label to a video clip



Making sandwich: present
Feeding animal: not present
...

- Action localization: search locations of an action in a video



Outline

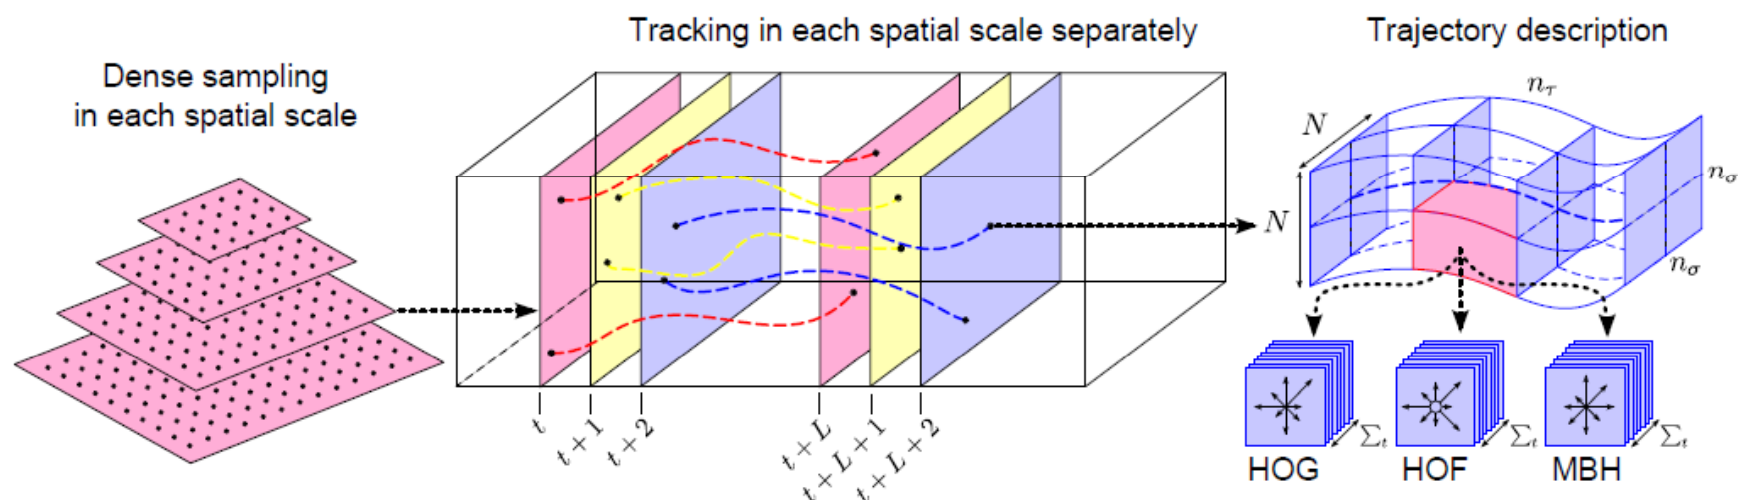
- *Improved video description*
 - *Dense trajectories and motion-boundary descriptors*
- Adding temporal information to the bag of features
 - Actom sequence model for efficient action detection
- Modeling human-object interaction

Dense trajectories - motivation

- Dense sampling improves results over sparse interest points for image classification [Fei-Fei'05, Nowak'06]
 - Recent progress by using feature trajectories for action recognition [Messing'09, Sun'09]
 - The 2D space domain and 1D time domain in videos have very different characteristics
- ➔ Dense trajectories: a combination of dense sampling with feature trajectories [Wang, Klaeser, Schmid & Lui, CVPR'11]

Approach

- Dense multi-scale sampling
- Feature tracking over L frames with optical flow
- Trajectory-aligned descriptors with a spatio-temporal grid



Approach

Dense sampling

- remove untrackable points
- based on the eigenvalues of the auto-correlation matrix




Feature tracking

- by median filtering in dense optical flow field

$$P_{t+1} = (x_{t+1}, y_{t+1}) = (x_t, y_t) + (M * \omega_t)|_{(\bar{x}_t, \bar{y}_t)}$$

- length is limited to avoid drifting



Feature tracking



KLT tracks



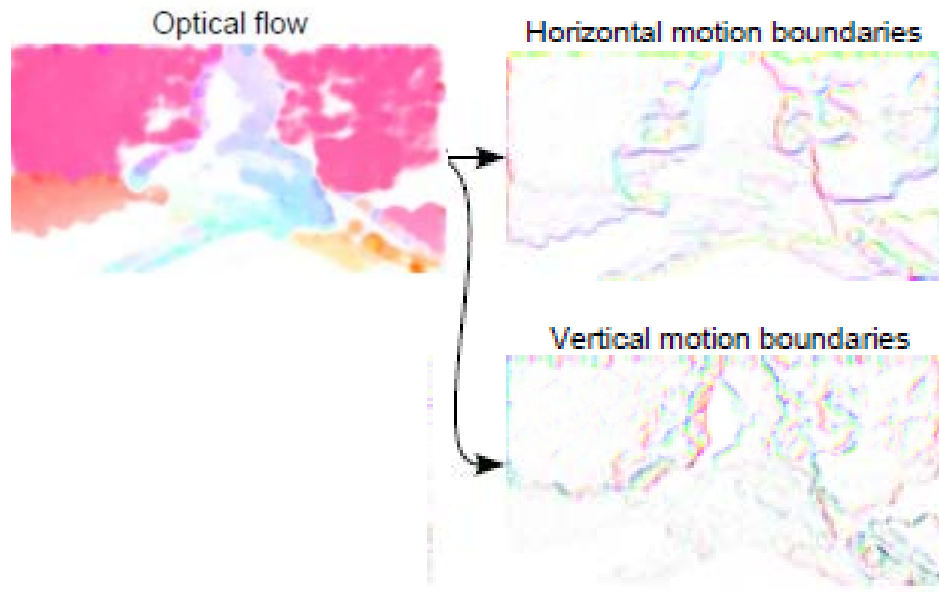
SIFT tracks



Dense tracks

Trajectory descriptors

- Motion boundary descriptor
 - spatial derivatives are calculated separately for optical flow in x and y , quantized into a histogram
 - relative dynamics of different regions
 - suppresses constant motions as appears for example due to background camera motion




Trajectory descriptors

- Trajectory shape described by normalized relative point coordinates

$$S = \frac{(\Delta P_t, \dots, \Delta P_{t+L-1})}{\sum_{j=t}^{t+L-1} \|\Delta P_j\|}$$

- HOG, HOF and MBH are encoded along each trajectory



Experimental setup

- Bag-of-features with 4000 clusters obtained by k-means, classification by non-linear SVM with RBF + chi-square kernel
 - confirmed by recent results with Fisher vector + linear SVM
- Descriptors are combined by addition of distances
- Evaluation on two datasets: UCFSport (classification accuracy) and Hollywood2 (mean average precision)
- Two baseline trajectories: KLT and SIFT

Comparison of descriptors

	Hollywood2	UCFSports
Trajectory	47.8%	75.4%
HOG	41.2%	84.3%
HOF	50.3%	76.8%
MBH	55.1%	84.2%
Combined	58.2%	88.0%

- Trajectory descriptor performs well
- HOF >> HOG for Hollywood2, dynamic information is relevant
- HOG >> HOF for sports datasets, spatial context is relevant
- MBH consistently outperforms HOF, robust to camera motion

Comparison of trajectories

	Hollywood2	UCFSports
Dense trajectory + MBH	55.1%	84.2%
KLT trajectory + MBH	48.6%	78.4%
SIFT trajectory + MBH	40.6%	72.1%

- Dense >> KLT >> SIFT trajectories

Comparison to state of the art

	Hollywood2 (SPM)	UCFSports (SPM)
Our approach (comb.)	58.2% (59.9%)	88.0% (89.1%)
[Le'2011]	53.3%	86.5%
other	53.2% [Ullah'10]	87.3% [Kov'10]

- Improves over the state of the art with a simple BOF model

Excellent results in TrecVid MED'12

- Combination of MBH SIFT, audio and text recognition
- Second in the know event challenge, first in the adhoc event challenge

Making sandwich – results



Rank 1 (pos)




Rank 20 (pos)



Rank 21 (neg)

Excellent results in TrecVid MED'12

FlashMob gathering – results



Rank 1 (pos)



Rank 18 (pos)



Rank 19 (neg)

Conclusion

- Dense trajectory representation for action recognition outperforms existing approaches
- Motion boundary histogram descriptors perform very well, they are robust to camera motion
- Efficient algorithm, on-line available at https://lear.inrialpes.fr/people/wang/dense_trajectories
- Recent excellent results in the TrecVID MED 2012 challenge

Outline

- Improved video description
 - Dense trajectories and motion-boundary descriptors
- *Adding temporal information to the bag of features*
 - *Actom sequence model for efficient action detection*
- Modeling human-object interaction

Approach for action modeling

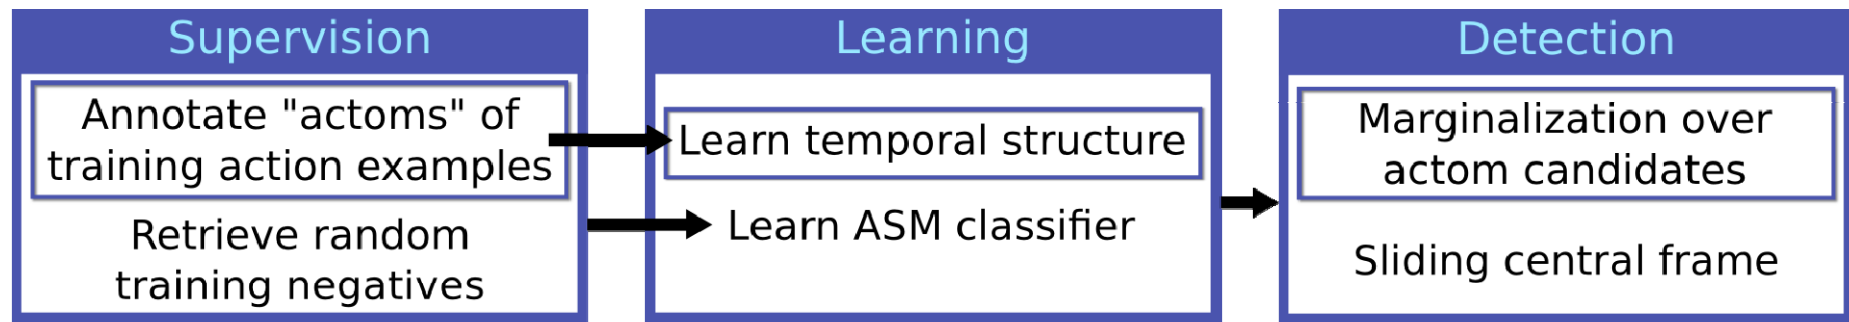
- Model of the temporal structure of an action with a sequence of “action atoms” (actoms)
- Action atoms are action specific short key events, whose sequence is characteristic of the action



Related work

- Temporal structuring of video data
 - Bag-of-features with spatio-temporal pyramids [Laptev'08]
 - Loose hierarchical structure of latent motion parts [Niebles'10]
 - Facial action recognition with action unit detection and structured learning of temporal segments [Simon'10]

Approach for action modeling



- Actom Sequence Model (**ASM**):
histogram of time-anchored visual features


Actom annotation

- Actoms for training actions are obtained manually (3 actoms per action here)
- Alternative supervision to clips annotation (beginning and end frames) with similar cost and smaller annotation variability
- Automatic detection of actoms at test time

Actom descriptor

- An actom is parameterized by:

- central frame location
- time-span
- temporally weighted feature assignment mechanism

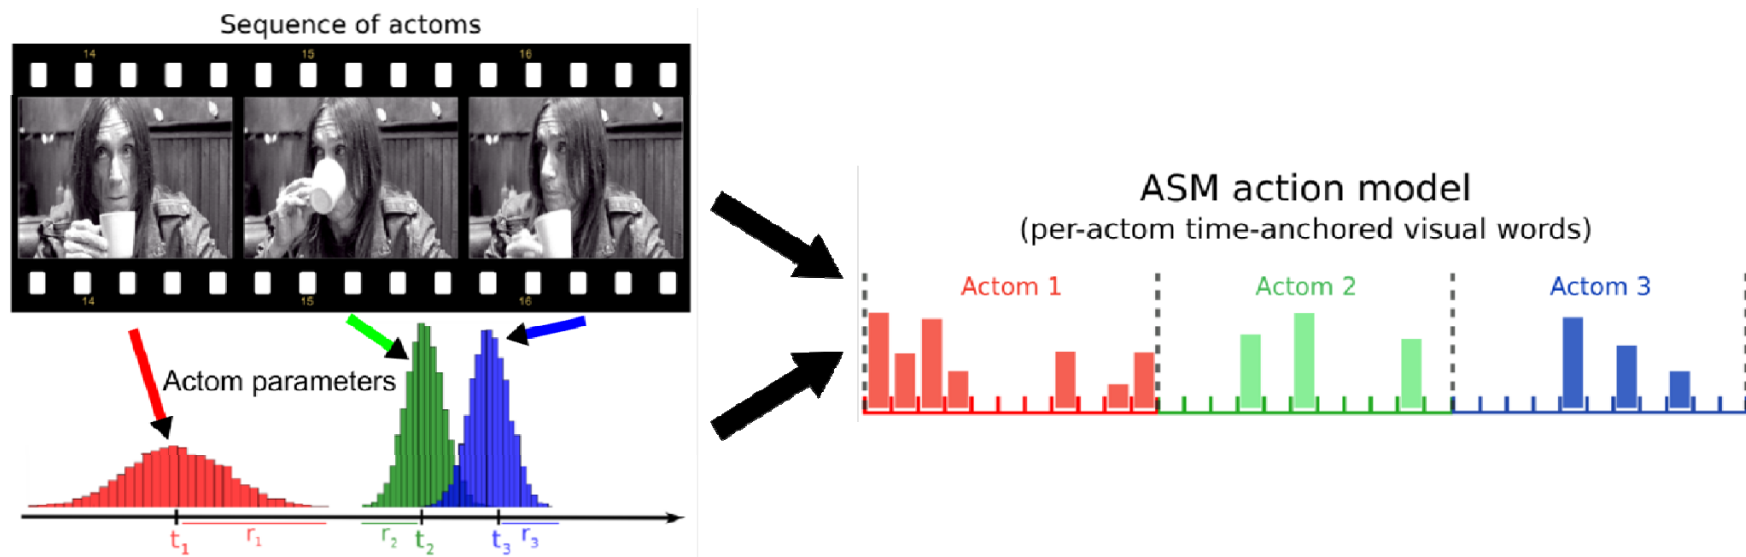


- Actom descriptor:

- histogram of quantized visual words in the actom's range
- contribution depends on temporal distance to actom center (using temporal Gaussian weighting)

Actom sequence model (ASM)

- ASM: concatenation of actom histograms

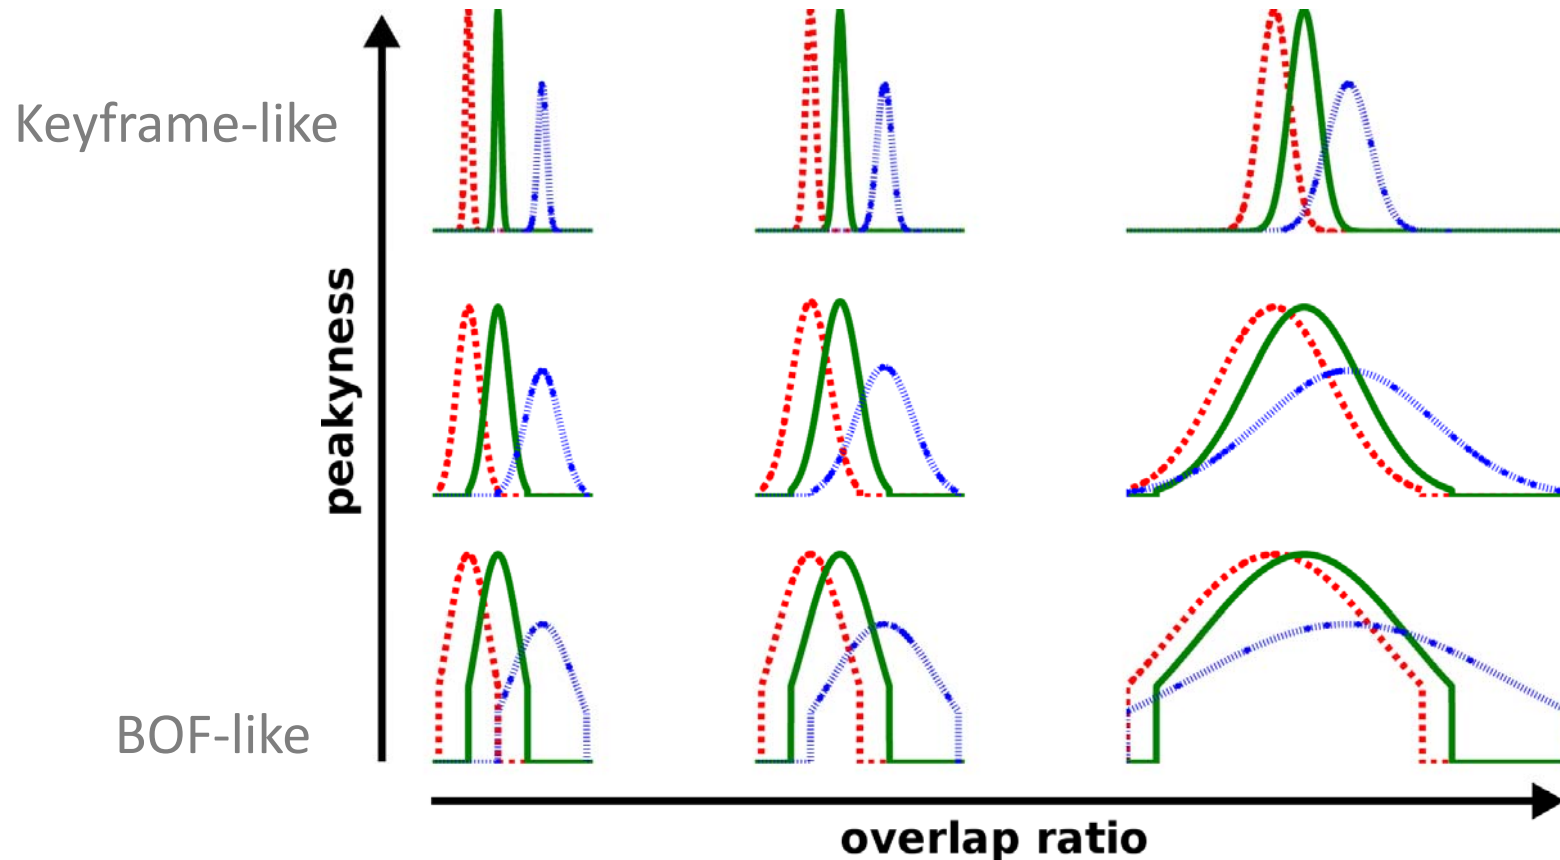


- Temporally structured extension of BOF
- Action represented by a single sparse sequential model

Actom Sequence Model (ASM)

Parameters

- ASM model has two parameters: *overlap* between actoms (controls radius) and soft-voting “*peakyness*” (controls profile)



Automatic temporal detection - training

- ASM classifier:
 - non-linear SVM on ASM representations with intersection kernel, random training negatives, probability outputs
 - estimates posterior probability of an action knowing the temporal location of its actoms
- Actoms unknown at test time:
 - use training examples to learn prior on temporal structure of actom candidates

Training

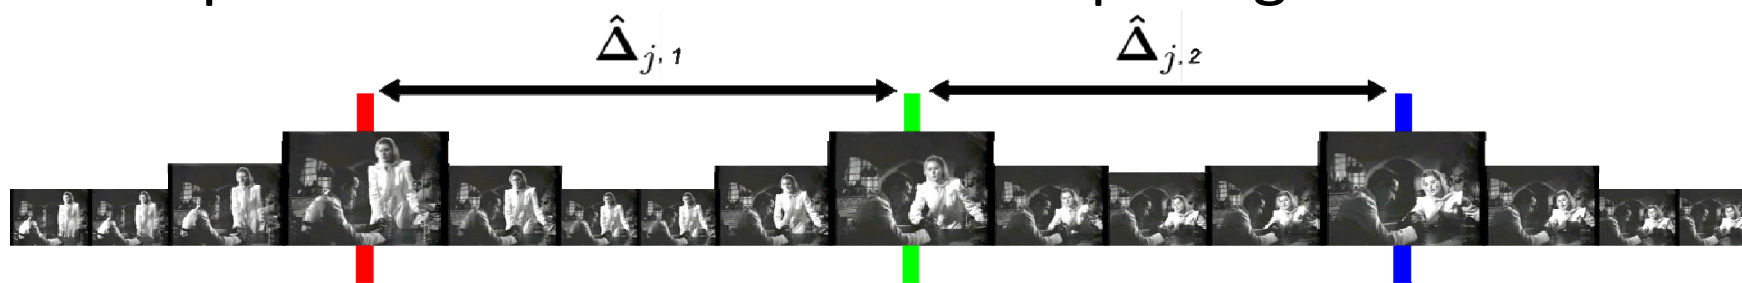
Action classifier

- **ASM classifier:** non-linear SVM on ASM representations
 - intersection kernel: $K(x, x') = \sum_{j=1}^N \min(x_j, x'_j)$
 - random training negatives
 - class-balancing
 - probability outputs

➔ **estimates posterior probability of an action knowing the temporal location of its actoms**
- Actoms unknown at test time:
 - use training examples to **learn actom candidates**

Prior on temporal structure

- Temporal structure: inter-actom spacings




- Non-parametric model of the temporal structure
 - kernel density estimation over inter-actom spacings from training action examples
 - discretize it to $\hat{\mathcal{D}} = \{(\hat{\Delta}_j, \hat{p}_j) \mid j = 1 \cdots K\}$, $\hat{p}_j = \mathbf{P}(\hat{\Delta}_j)$ (small support in practice: $K \approx 10$)
 - use as prior on temporal structure during detection

Training

Example of learned candidates

- Actom models corresponding to the \hat{D} learned for “smoking”
(with the ASM parameters used in our experiments)

Smoking temporal structure ($\rho = 25\%$, $p = 70\%$)




Automatic Temporal Detection

- Probability of action at frame t_m by marginalizing over all learned candidate actom sequences:

$$\mathbf{P}(\text{action at } t_m) = \sum_{j=1}^K \mathbf{P}(\text{action at } t_m | \hat{\Delta}_j) \mathbf{P}(\hat{\Delta}_j)$$

- Sliding central frame: detection in a long video stream by evaluating the probability every N frames ($N=5$)




- Non-maxima suppression post-processing step

Experiments - Datasets

- Coffee & Cigarettes: localize drinking , smoking in 36k frames [Laptev'07]



- DLSBP: localize opening a door, sitting down in 443k frames [Duchenne'09]



- Evaluation: average precision (AP) computed wrt 20% overlap with ground truth test actions

Quantitative Results

Coffee & Cigarettes

Method	“Drinking”	“Smoking”
matching criterion: OV20		
DLSBP [9]	40	NA
LP-T [12]	49	NA
KMSZ-T [78]	59	33
BOF	36 (± 1)	17 (± 2)
BOF T3	44 (± 2)	20 (± 3)
ASM	63 (± 3)	40 (± 4)

DLSBP

Method	“Open Door”	“Sit Down”
matching criterion: OV20		
DLSBP [9]	14	14
BOF	8 (± 3)	14 (± 3)
BOF T3	8 (± 1)	17 (± 3)
ASM	14 (± 3)	22 (± 2)

- ASM method outperforms BOF
- ASM improves over rigid temporal structure, BOF T3
(BOF T3: concatenation of 3 BOF: beginning, middle and end of the action)
- More accurate detections with ASM compared to the state of the art

Qualitative Results

Central frames

Frames of the top 5 actions detected with ASM for
drinking and opening a door
(only #2 of opening a door is a false positive)



Qualitative Results


Actoms

Frames of automatically detected actom sequences for 4 actions

Open Door



Drinking




Smoking




Sitting Down



Localization results for action drinking



Localization results for action smoking



Conclusion

- ASM: efficient model of actions with a flexible sequence of key semantic sub-actions (actoms)
- Principled multi-scale action detection using a learned prior on temporal structure
- ASM outperforms bag-of-features, rigid temporal structures and state of the art

Outline

- Improved video description
 - Dense trajectories and motion-boundary descriptors
- Adding temporal information to the bag of features
 - Actom sequence model for efficient action detection
- *Modeling human-object interaction*

Action recognition

- Action recognition is person-centric



Movies




TV




YouTube

Action recognition

- Action recognition is person-centric



Movies




TV



YouTube

Action recognition

- Description of the human pose
 - Silhouette description [Sullivan & Carlsson, 2002]
 - Histogram of gradients (HOG) [Dalal & Triggs 2005]



- Human body part estimation [Felzenszwalb & Huttenlocher 2005]



Importance of action objects



- Human pose often not sufficient by itself
- Objects define the actions

Action recognition from still images

- Supervised modeling interaction between human & object
[Gupta et al. 2009, Yao & Fei-Fei 2009]
- Weakly-supervised learning of objects [Prest, Schmid & Ferrari 2011]



Results on PASCAL VOC 2010 Human action classification dataset

Importance of temporal information




- Video/temporal information necessary to disambiguate actions
- Temporal context describes the action/activity
- Key frames provide significant less information

Modeling temporal human-object interactions



Describing human and object tracks and their relative motion

Tracking humans and objects



Fully automatic human tracks: state of the art detector + Brox tracks

Object tracks: detector learnt from annotated training examples + Brox tracks


Extraction of a large number of human-object track pairs

Action descriptors

- Interaction descriptor: relative location, area and motion between human and object tracks

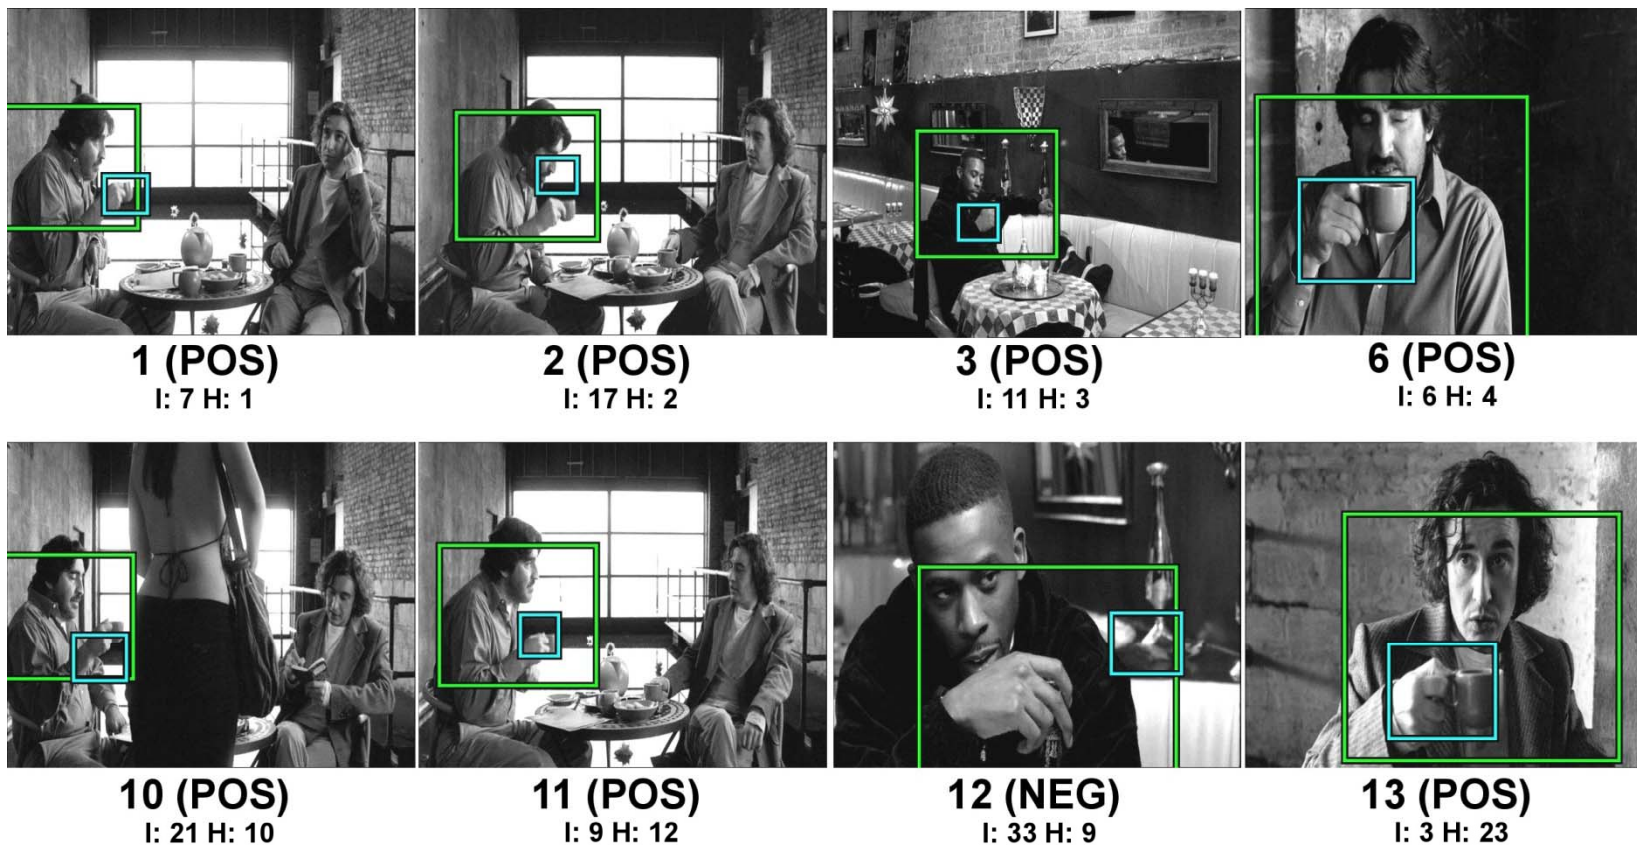


- Human track descriptor: 3DHOG-track [Klaeser et al.'10]



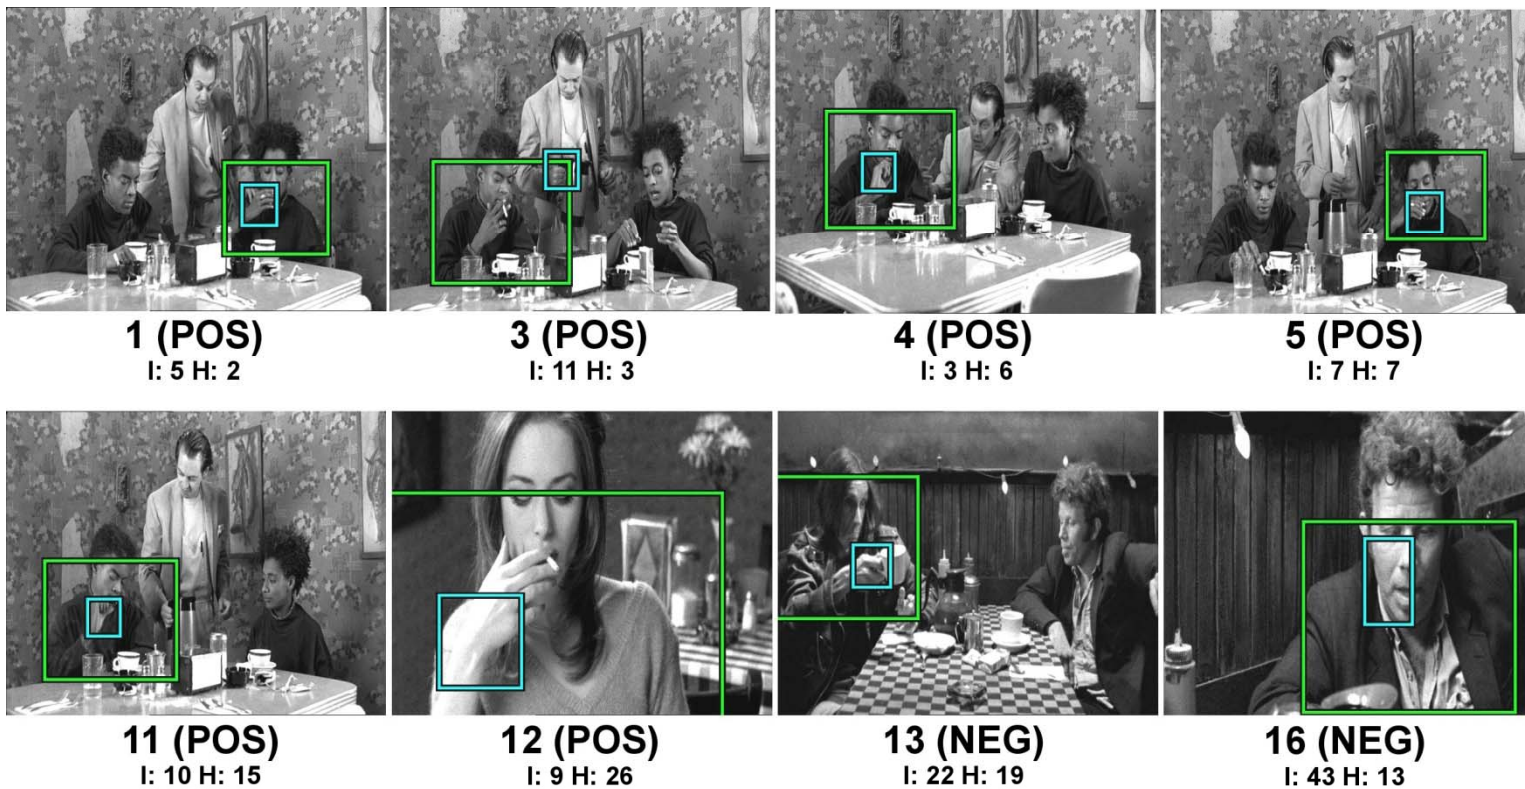
Experimental results on C&C

Drinking





Experimental results on C&C

Smoking



Experimental results on C&C



Comparison to the state of the art

	Drinking	Smoking
Interaction classifier	31.60	16.20
Object classifier	4.30	5.50
3DHOG-track classifier	52.20	21.50
Combination	62.10	32.80
Laptev et al. [22]	43.40	-
Willems et al. [35]	45.20	-
Klaeser et al. [20]	54.10	24.50

Experimental results on Gupta dataset

Answering the
phone



Making a phone call



Drinking



Using a light torch



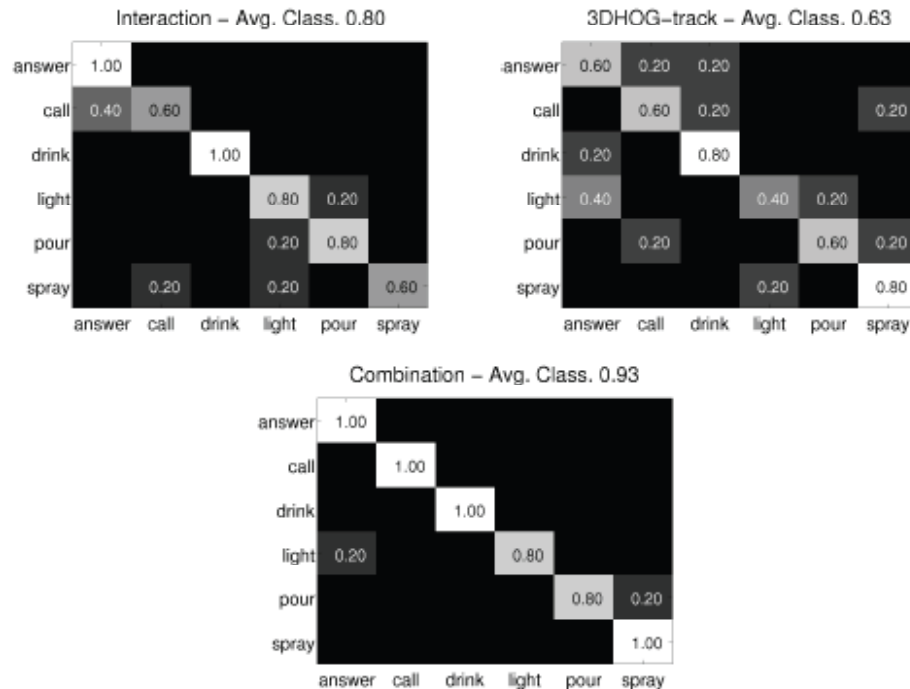
Pouring water from
a cup



Using a spray bottle



Experimental results on Gupta dataset

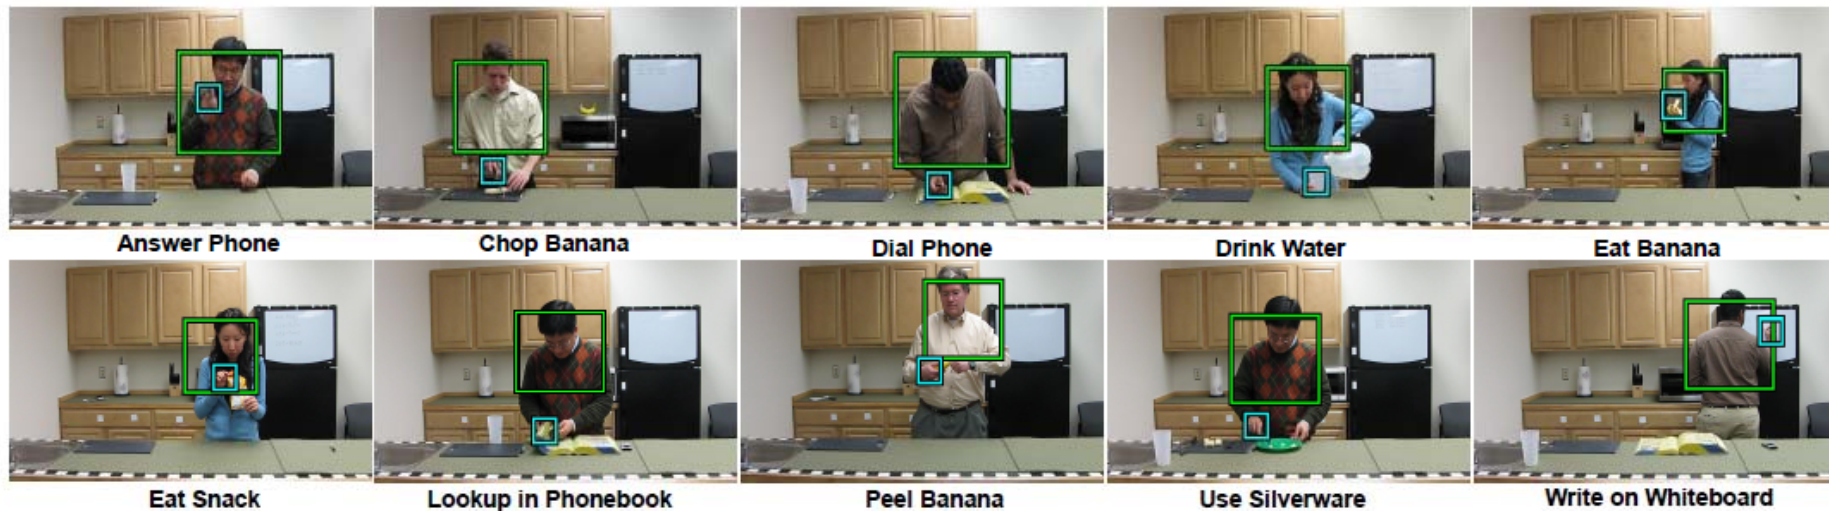


	Gupta video
Interaction classifier	80.00
Object classifier	36.60
3DHOG-track classifier	63.30
Combination	93.30
Gupta et al. [17]	93.00

- Interactions achieve the best performance alone
- Combination improves results further: only 2 misclassified samples
- Comp. state of the art: Gupta use significantly more training information

Experimental results on Rochester dataset

- Rochester daily activities dataset
 - 150 videos of 5 persons
 - leave-one-person-out test scenario



Experimental results on Rochester dataset

	Rochester Daily Activities
Interaction classifier	74
Combination (our full method)	92
Messing et al. (full method) [29]	89
Messing et al. (point tracks) [29]	67
Matikainen et al. (point tracks) [28]	70

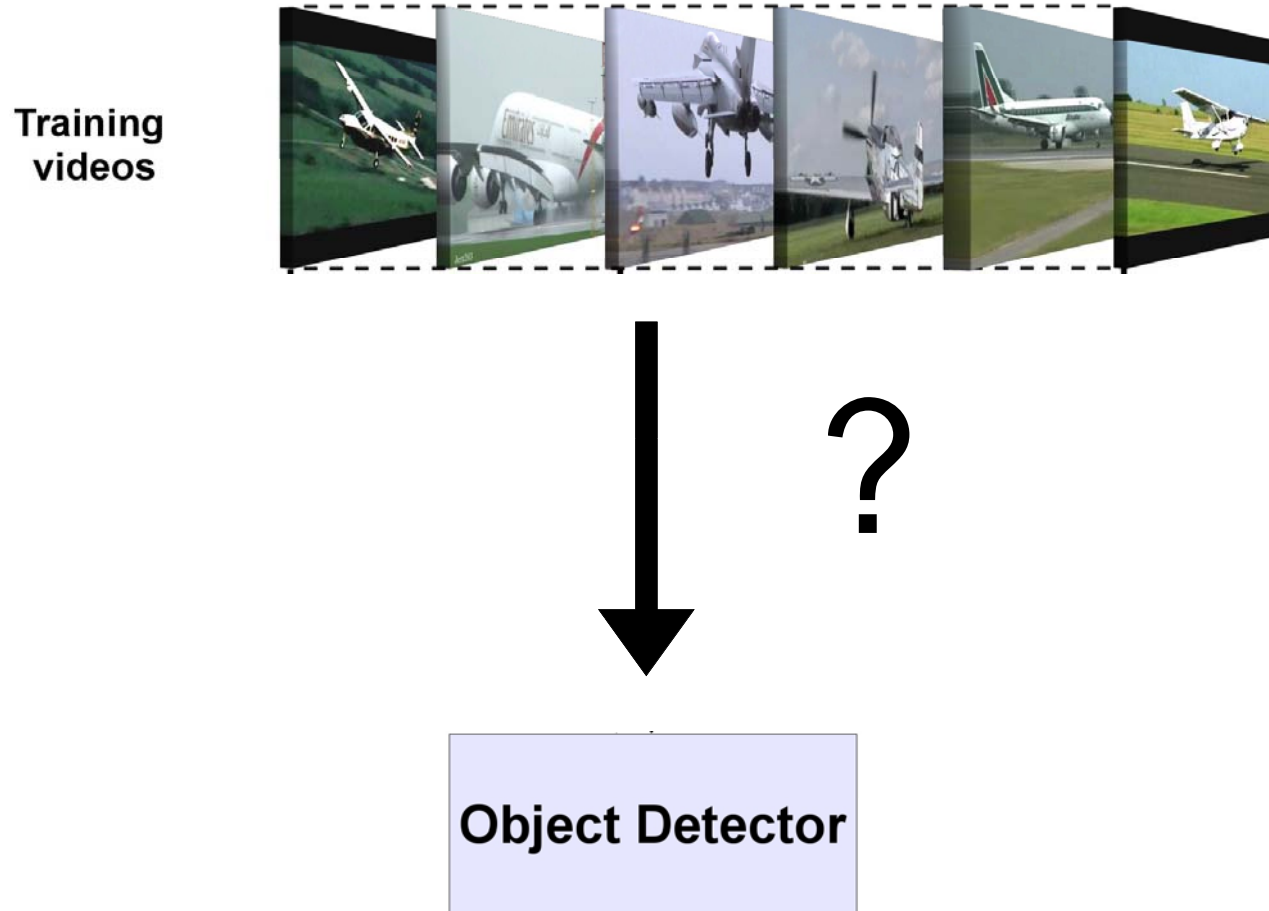
Experimental results on Rochester dataset



Conclusion

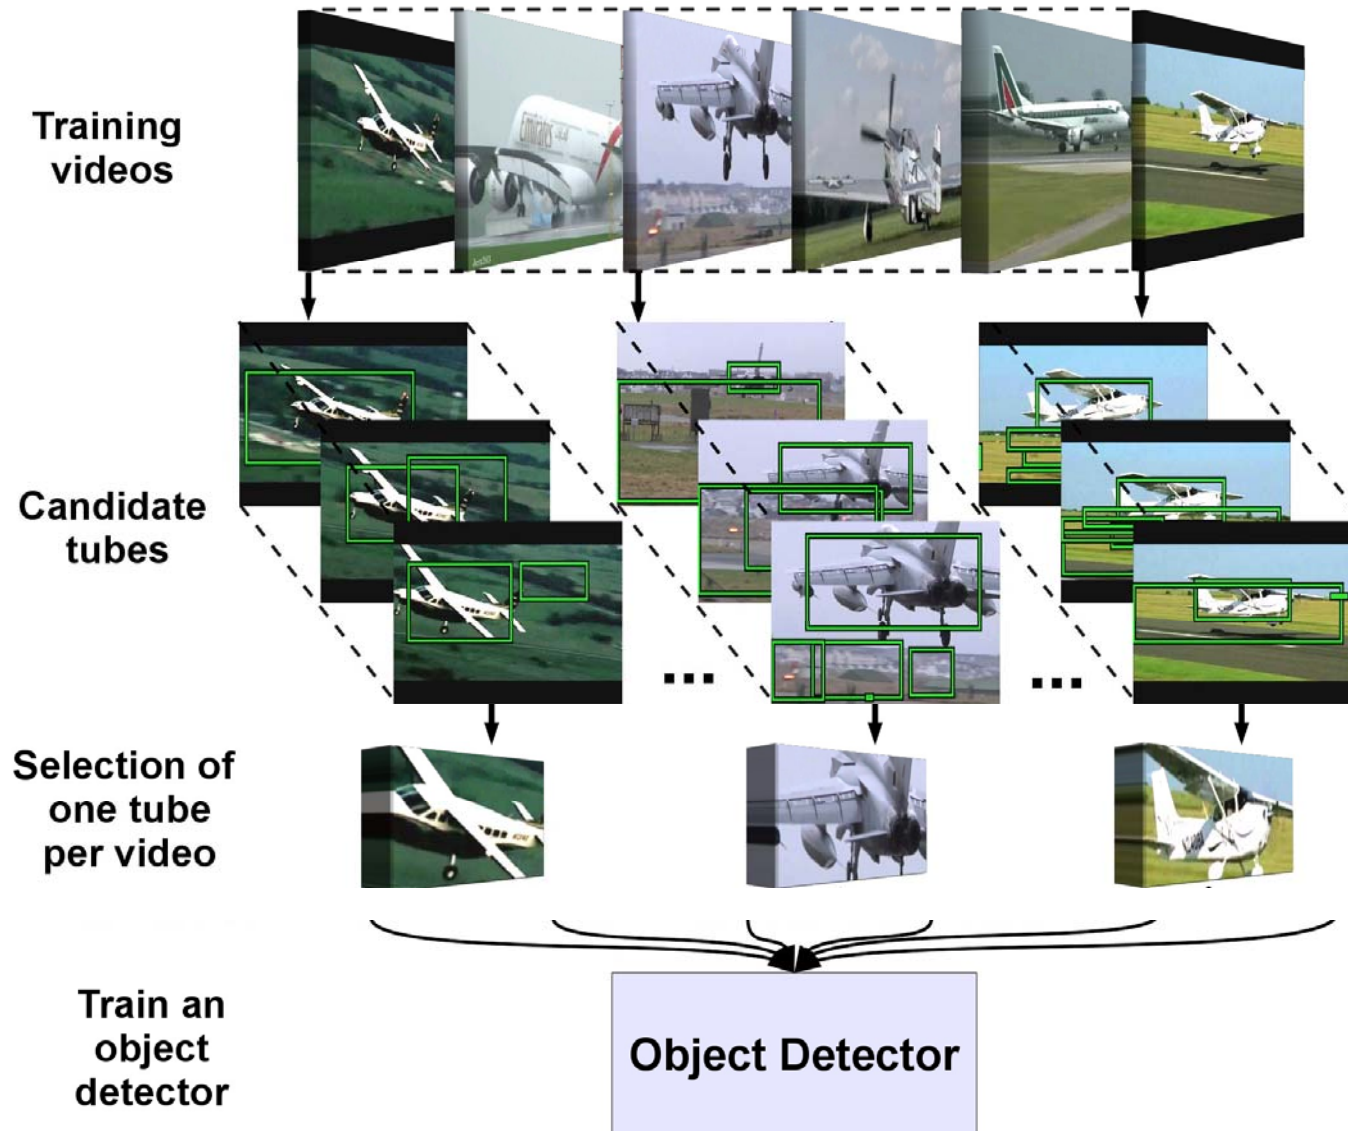
- Human-object interaction descriptor obtains state-of-the-art performance
- Complementary to 3DHOG-track descriptor
- Combination obtains excellent performance
- Automatic extraction of objects

Automatic extraction of objects




Prest, Leistner, Civera, Schmid, Ferrari CVPR 2012, *Learning object detectors from weakly annotated video*

Automatic extraction of objects



Candidate tubes


dense point tracks



N. Sundaram et al., Dense point trajectories by GPU-accelerated large displacement optical flow, ECCV 2010

Candidate tubes

motion segmentation



Candidate tubes


motion segmentation



Selecting candidate tubes



Selecting candidate tubes




Training + testing object detectors


Still images



Video



Combination



Still images from PASCAL VOC 2007

Discussion

- Need for more challenging datasets
 - Need for realistic datasets



KTH dataset



Hollywood dataset

- Scale up number of classes (today ~10 actions per dataset)
- Increase number of examples per class, possibly with weakly supervised learning (the number of examples per videos is low)
- Define a taxonomy, use redundancy between action classes to improve training
- Manual exhaustive labeling of all actions impossible

Discussion

- Make better use of the large amount of information inherent in videos
 - automatic collection of additional examples
 - improve models incrementally
 - use weak labels from associated data (text, sound, subtitles)
- Many existing techniques are straightforward extensions of methods for images
 - almost no use of 3D information
 - learn better interaction and temporal models
 - design activity models by decomposition into simple actions