Joint Audio-Visual Signatures for Web Video Analysis

Dan Ellis & Shih-Fu Chang
Wei Jiang, Courtenay Cotton, Yu-Gang Jiang, Xiaohong Zeng

Electrical Engineering, Columbia University, NY USA
dpwe@ee.columbia.edu http://labrosa.ee.columbia.edu/

1. Web Video Analysis
2. Labeled Data Gathering
3. Scene / Object Context
4. Future Work

Lab ROSA
Laboratory for the Recognition and Organization of Speech and Audio

DVMM lab
digital video | multimedia laboratary

Columbia University
IN THE CITY OF NEW YORK

1. Web Video Analysis

- Web video
  - Huge volume
  - Poor labels
  - Low quality

- The goal:
  - Automatic, efficient, human-like labeling
    - into categories
    - by events/objects

Conventional Approach: Global Features

- Train classifiers for predefined categories based on statistics of whole clip

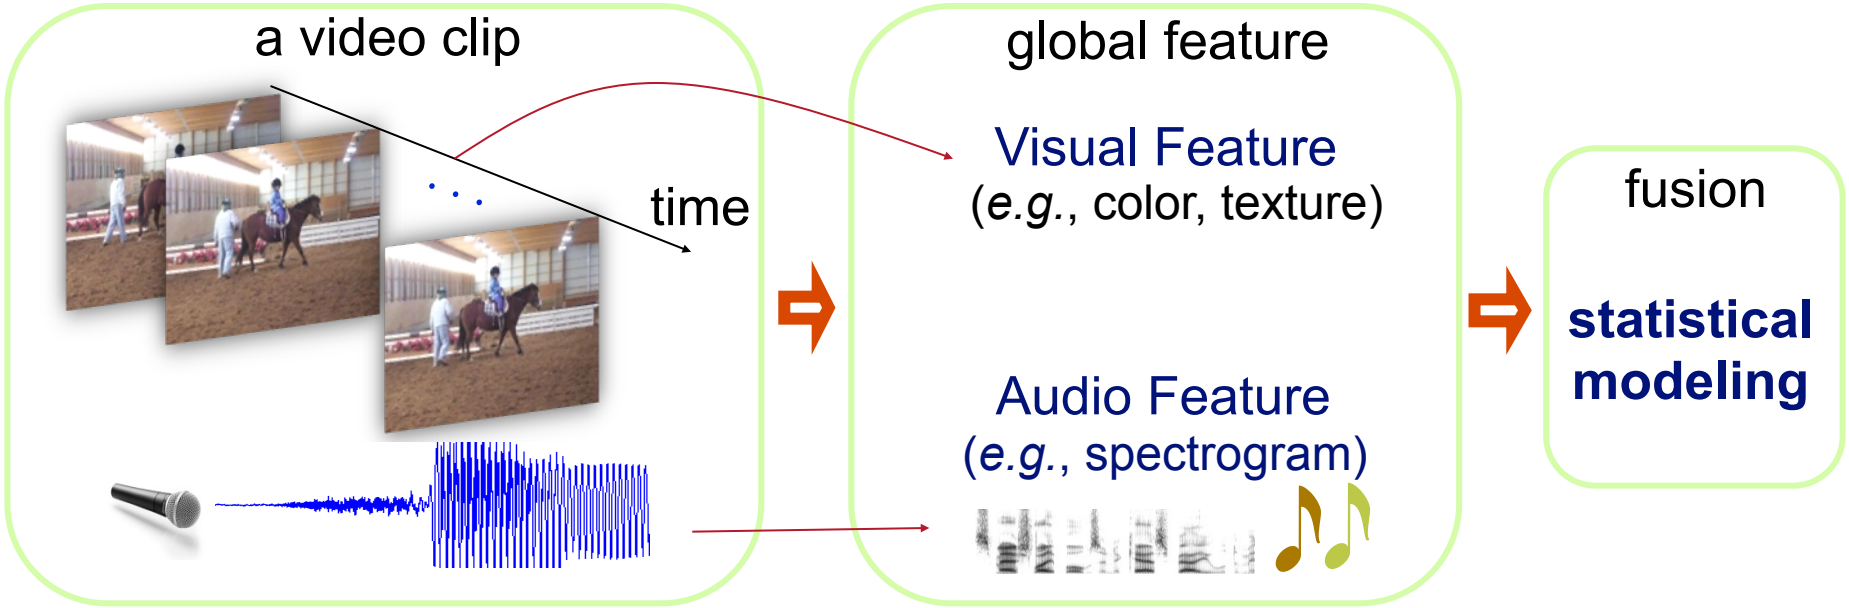

- no object-level description

[Chang et al. MIR 2007]
[Cristani et.al., TMM 2007]

Novel Approach: Audio-Visual Atoms

- Decompose aud/vid into object-like atoms
  - statistical models of their combinations

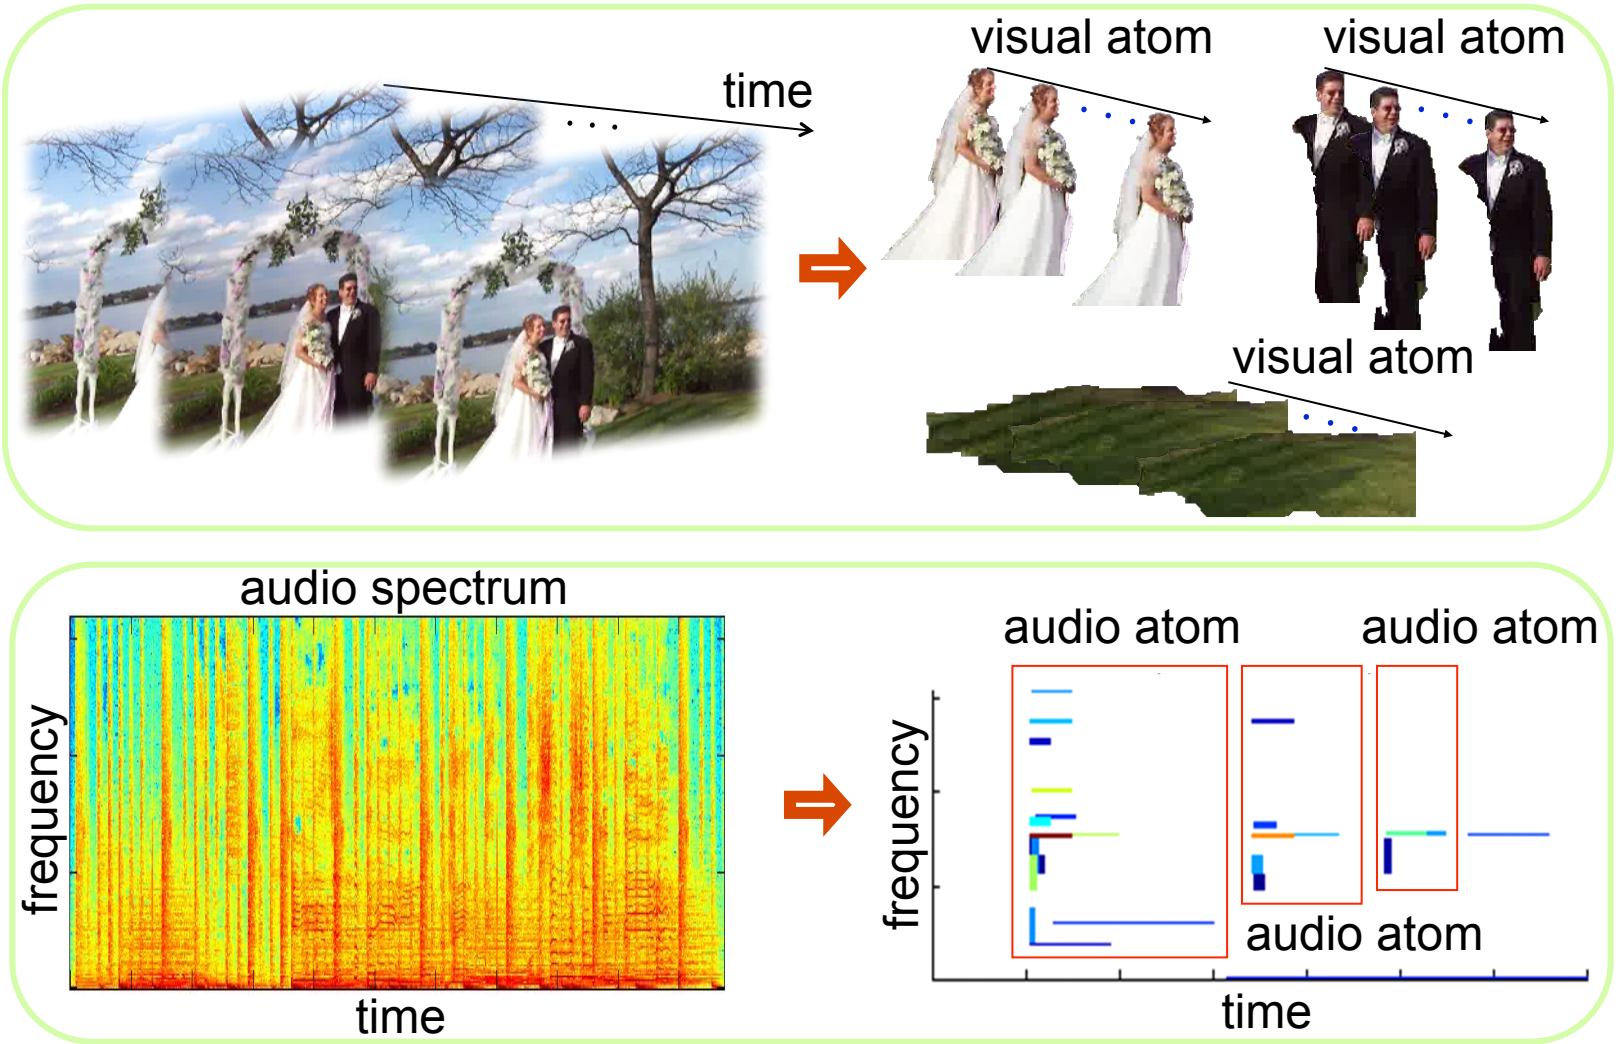

Challenges in Unconstrained Video

- Poor quality
  - focus
  - lighting
  - camera motion
  - occlusions
  - ambient noise
  - handling noise

- Poor A-V correlation
  - sounds from unobserved objects
  - sound-producing motions are slight

Visual Atom Formation

- Point tracking
  - sift points link successive frames
- + Region Segmentation
  - color / texture define regions

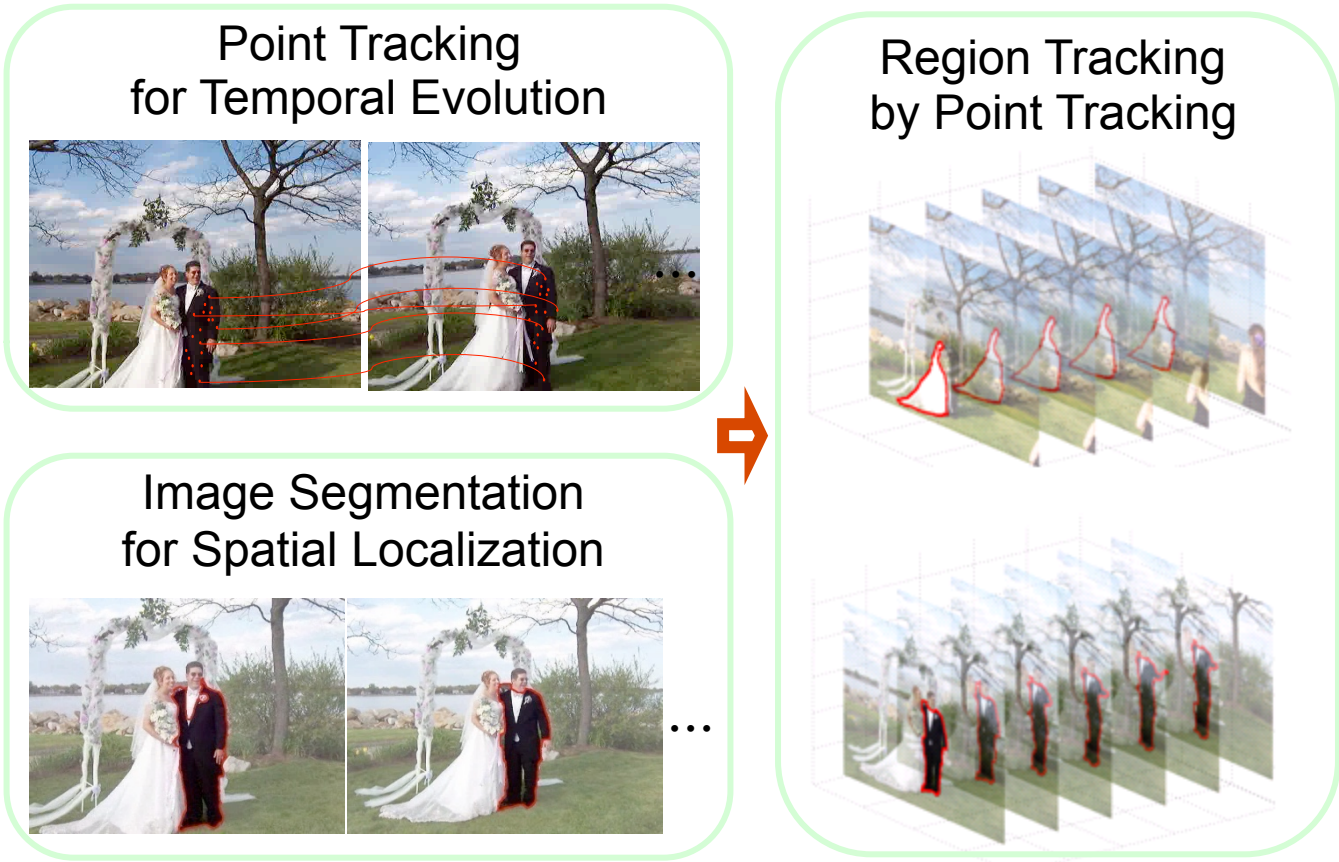

- Link shorter tracks across time

[Jiang et.al., MM 2009]

Visual Atom Examples

- A few examples out of 100+ for "wedding"
  - build codebook based on appearance, shape

Audio Atom Formation

- **Extract & describe Transients**
  - for maximum robustness against noise

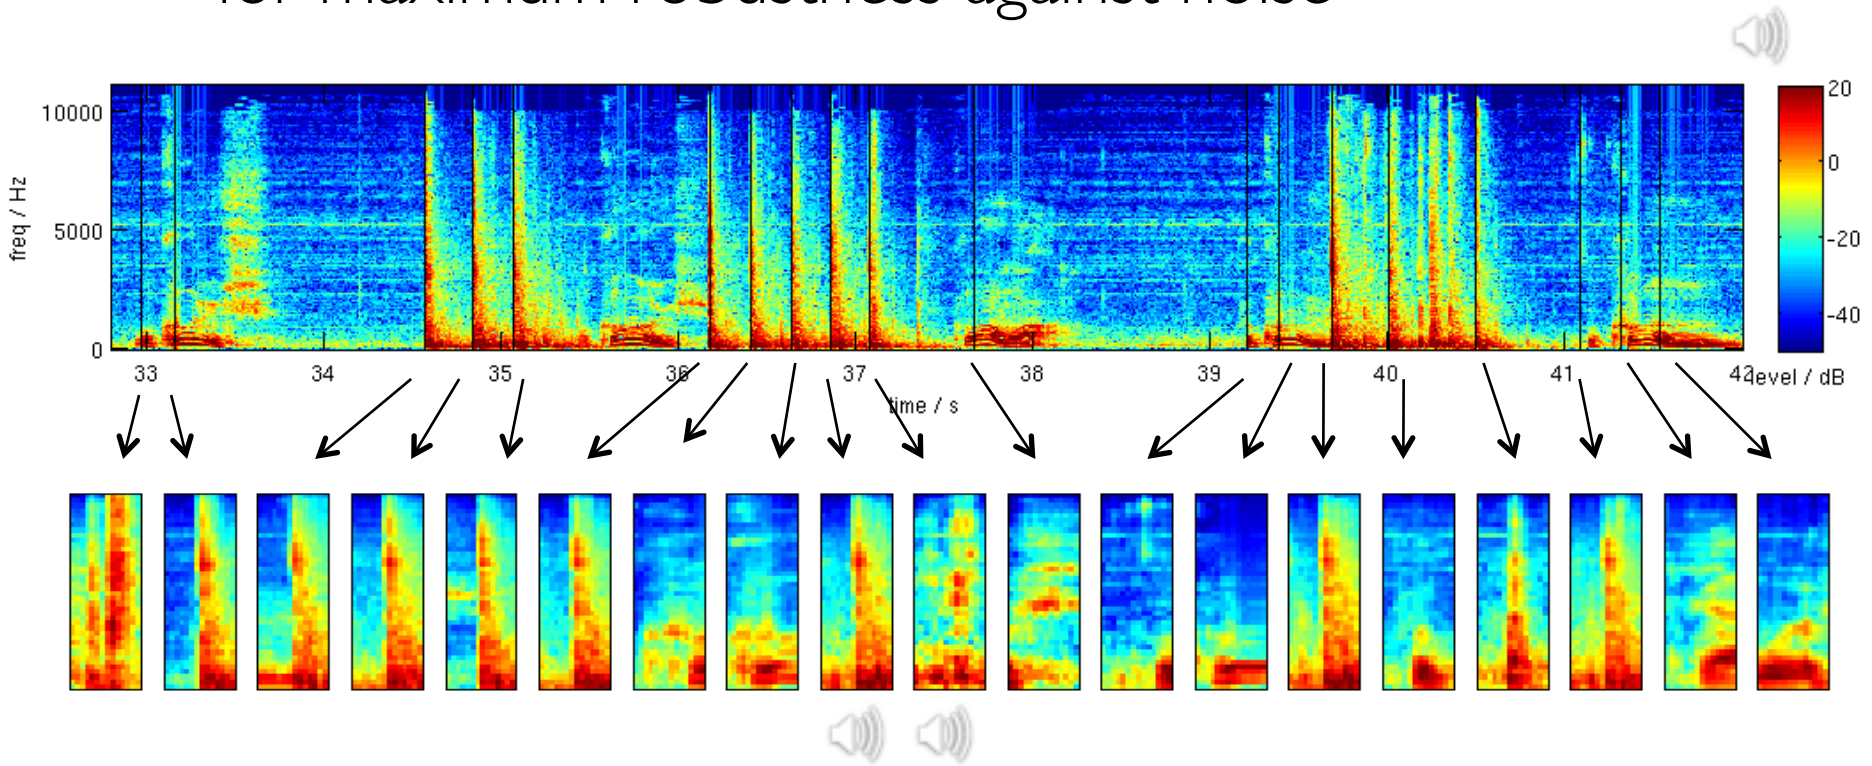

- **Multiscale analysis to find energy "bursts"**
  - extract 250 ms mel-spectrum window
  - describe with PCA

Audio Atom Results

- K-means clustering to form codewords
- Better noise resistance than MFCCs

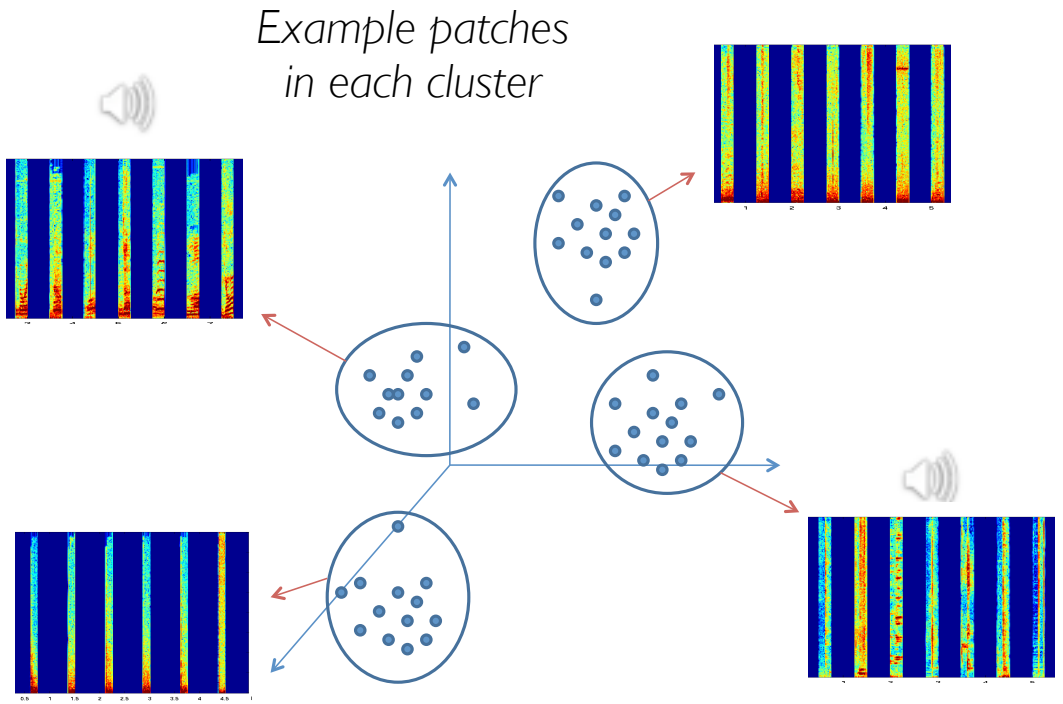

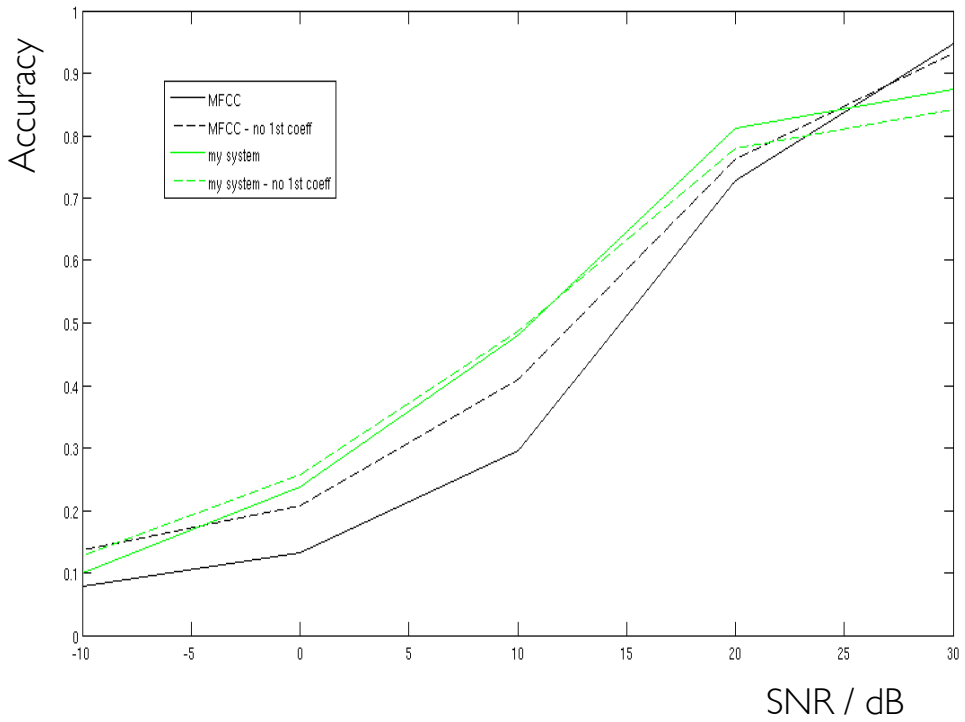

Joint Audio-Visual Atoms

- Consider all possible A-V combinations

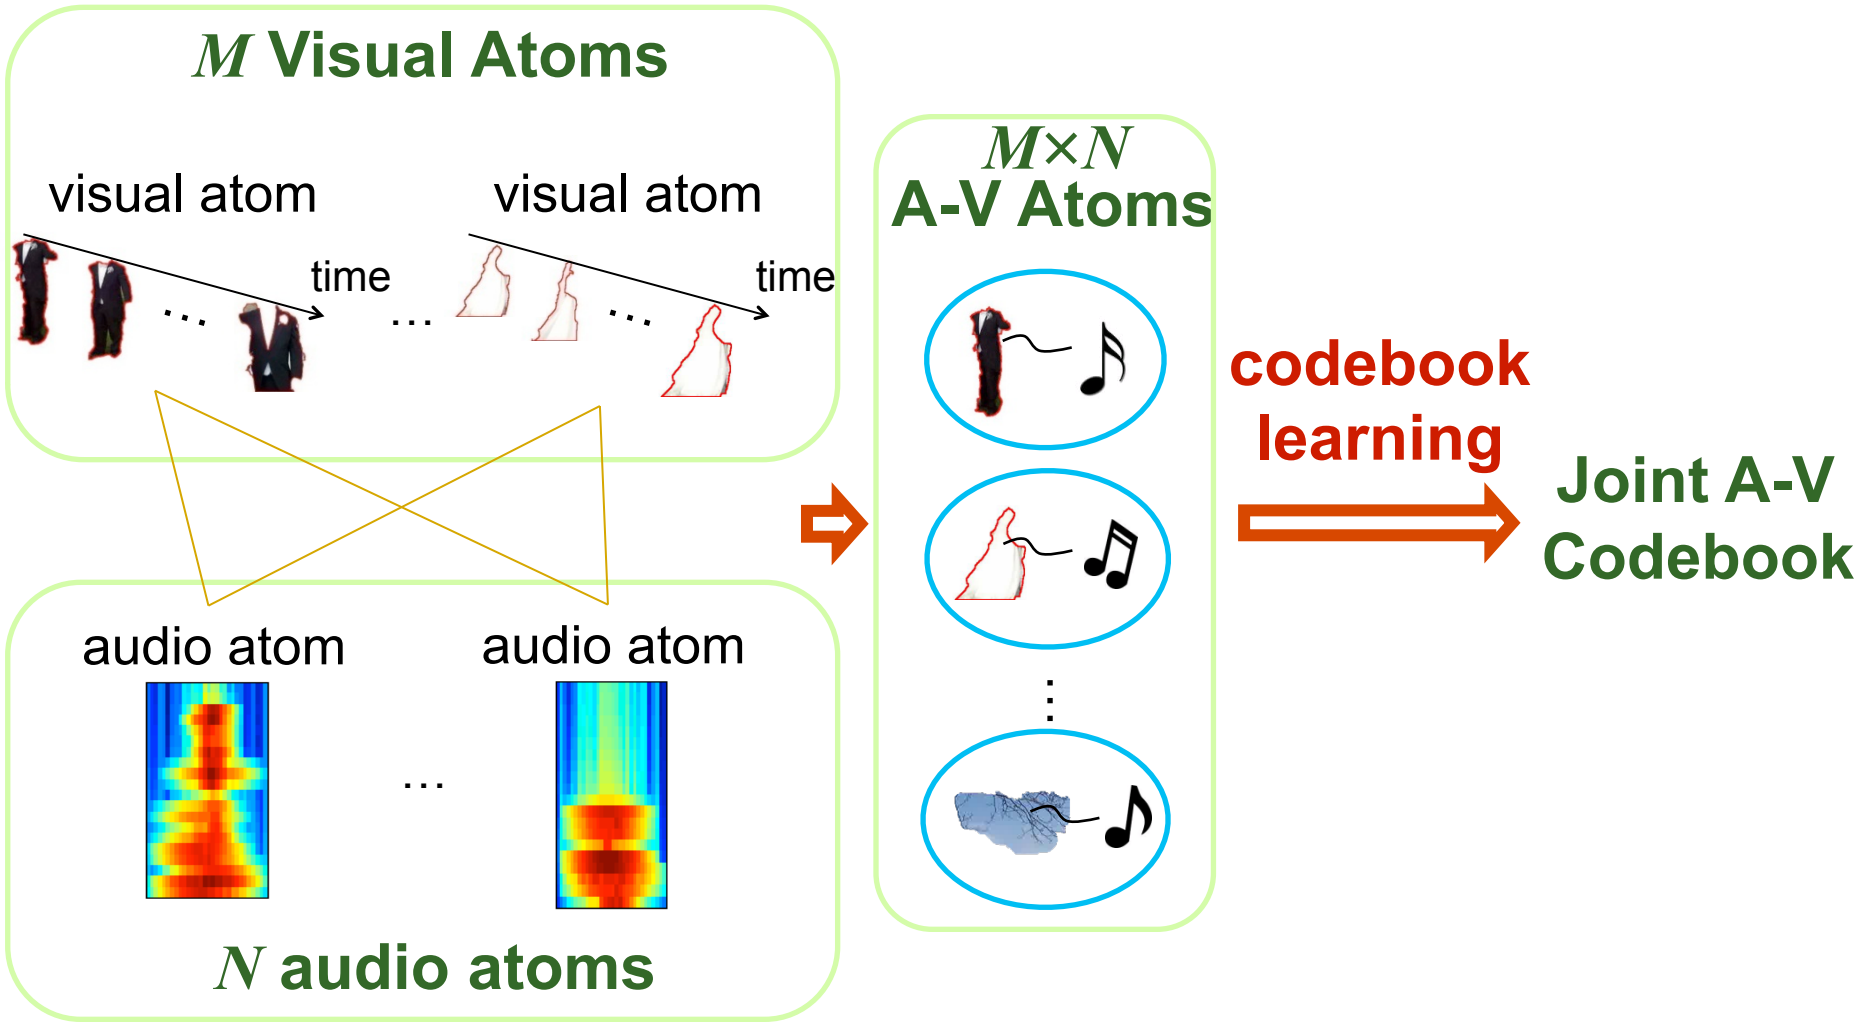

Concept Codebook Learning

- by Multiple-Instance Learning (MIL)
  - only have clip-level labels

[Maron et al., NIPS, 1998]

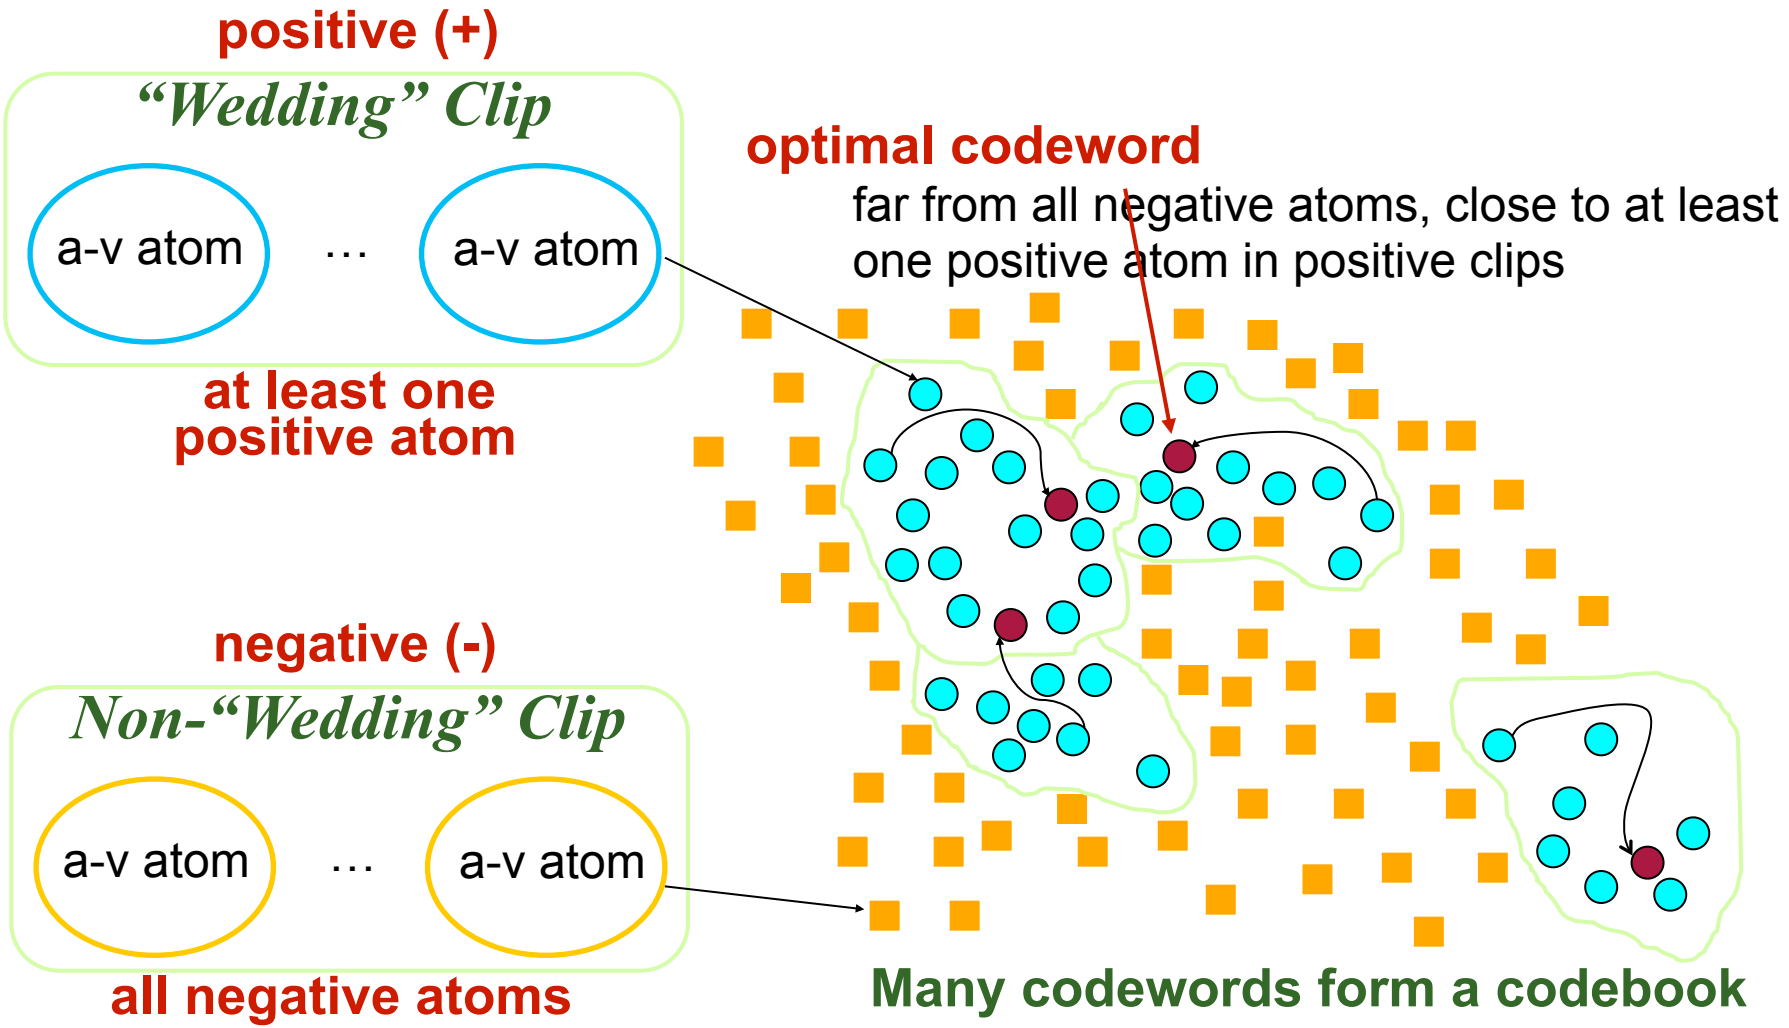

Evaluation on Consumer Video

- Kodak consumer video benchmark set
  - 1358 videos (813 training)
  - 25 labels

[Loui et al. MIR 2007]

Example A-V Codewords

- "Wedding" class

black suit + romantic music

red carpet + romantic music

green grass + romantic music

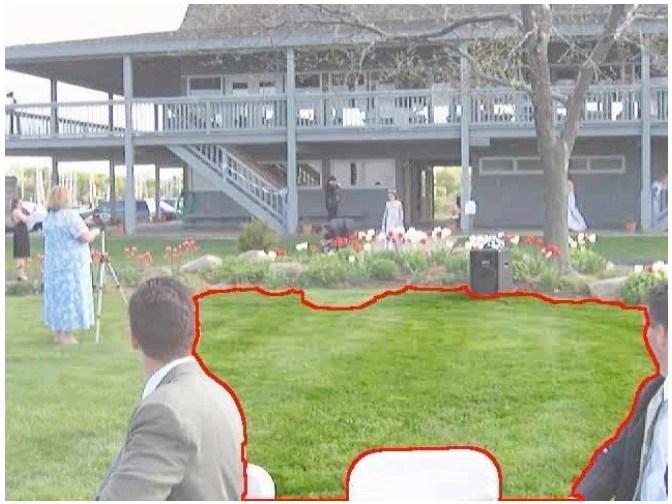

white gown + romantic music

Example A-V Codewords

- "Parade" class

marching people + parade sound

road + parade sound

marching people + parade sound

road + parade sound

Example A-V Codewords

- "Beach" class

sand + beach sound

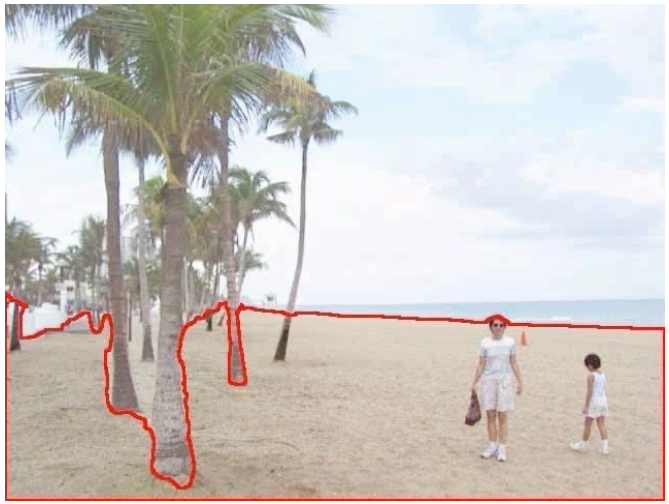

water + beach sound

beach chair + beach sound

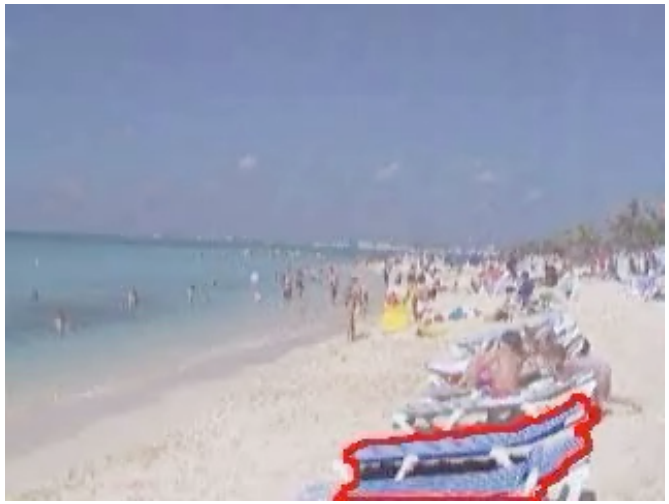

people with swim suit + beach sound

Consumer Video Evaluation

- A-V atoms vs. Static regions
  - Average Precision on test set

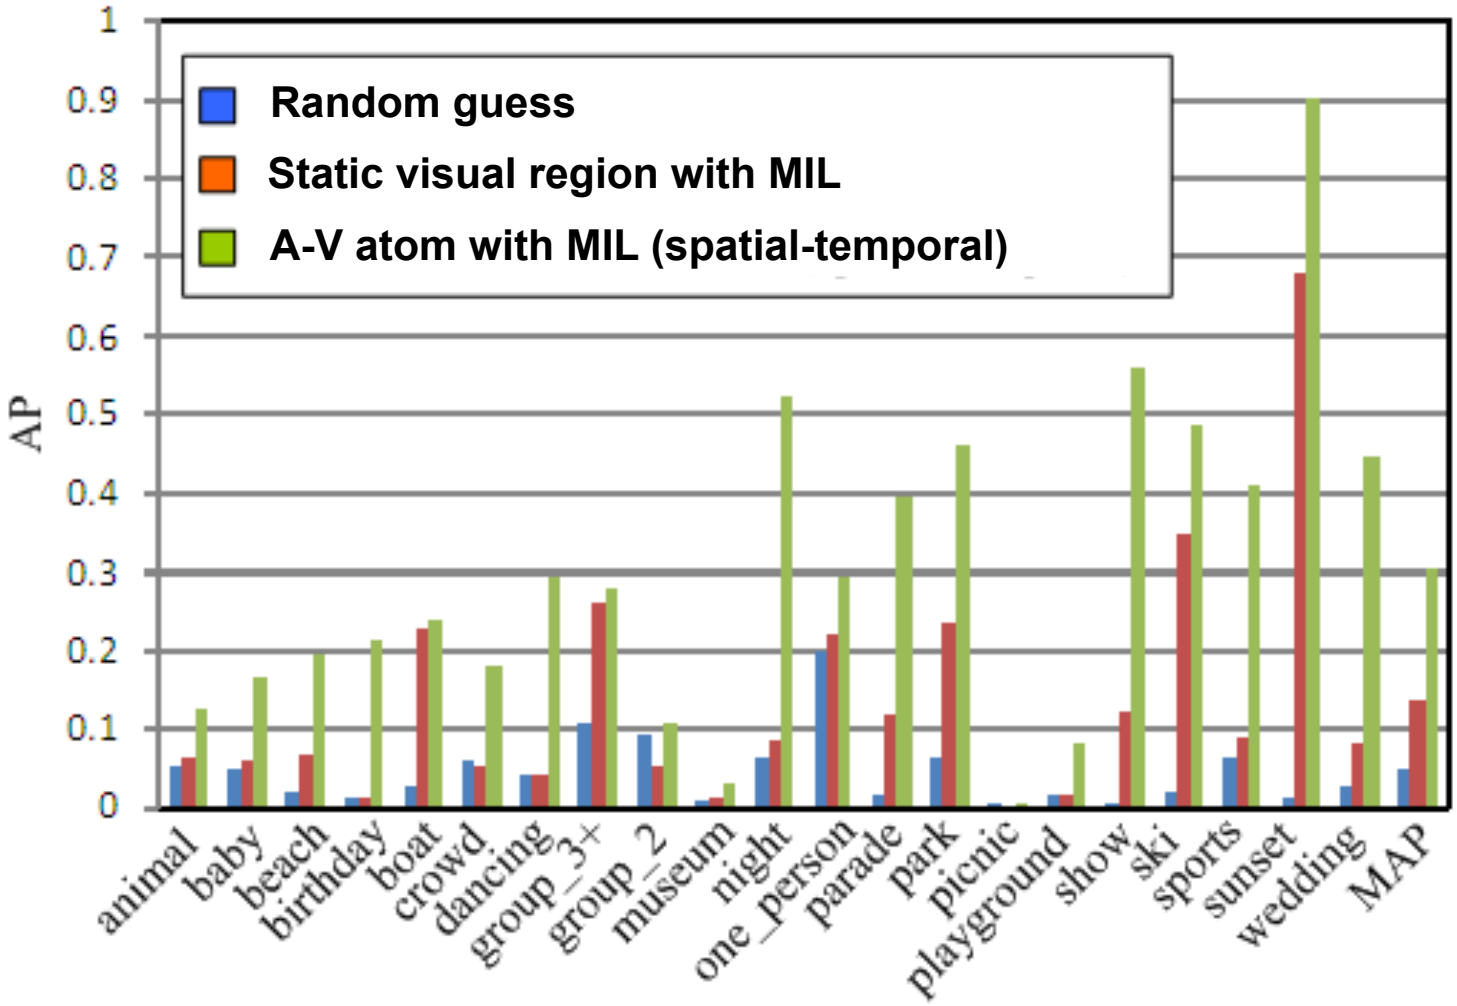

Consumer Video Evaluation

- Audio / Video / both
  - audio useful for many classes

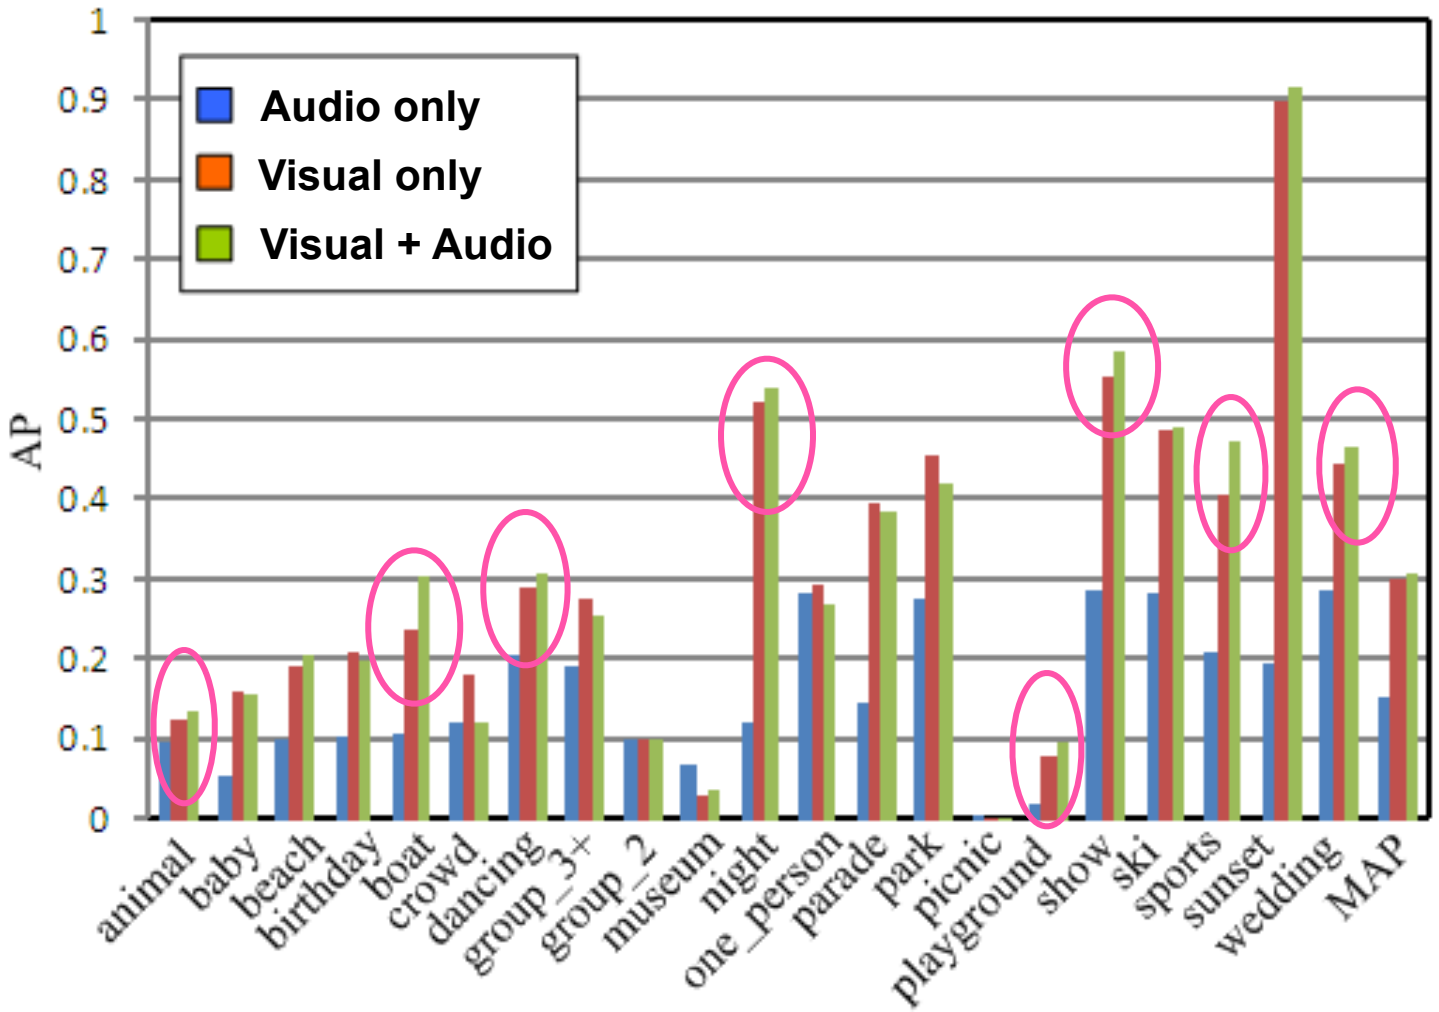

?

2. Labeled Data Gathering

- Learning codebooks requires labeled data
  - novel concept → need new labels
  - more labels → better performance
- Amazon "Mechanical Turk"

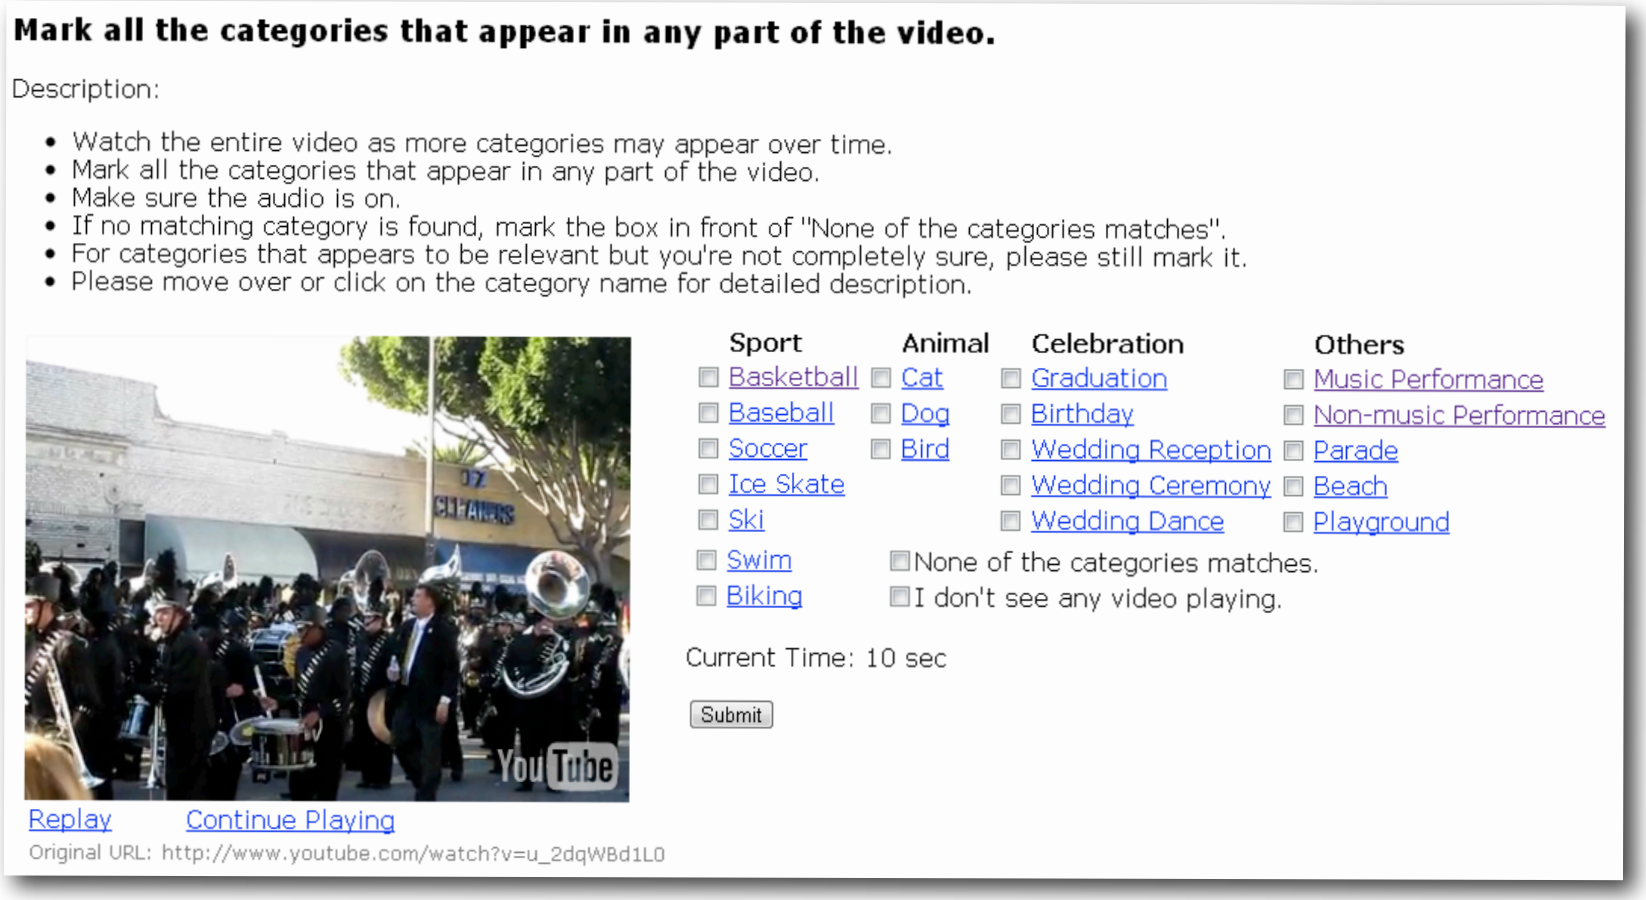

MTurk Results

- **Data: YouTube** raw camera uploads
  - based on keyword search for 20 categories
- **MTurk Human Intelligence Tasks (HITs)**
  - paid $0.02 per 10s clip (~$7/hr)
  - 4 labelers/clip, finished 9,641 videos in 2 weeks

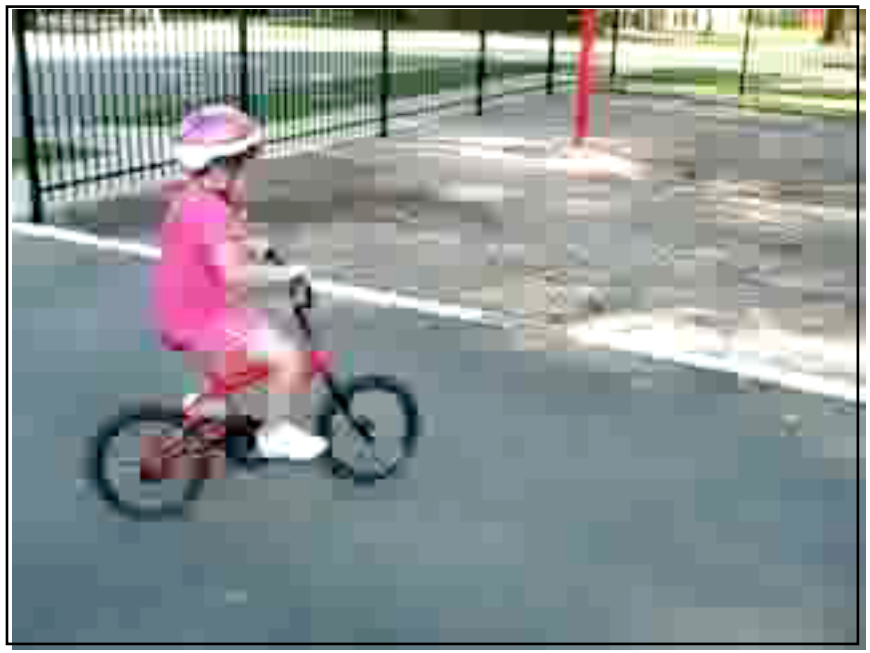

Playground, Biking

Non-music Performance, Ice Skating

3. Scene / Object Context

- How to identify events in video?
  - not objects, not locations
  - e.g. "people kissing"
- Traditional approach:
  - get low level features for large training set
  - statistical classifier
- Our approach
  - use specialized mid-level detectors (faces, cars)
  - learn context, relationships
  - e.g. "kissing" = 2 faces moving together

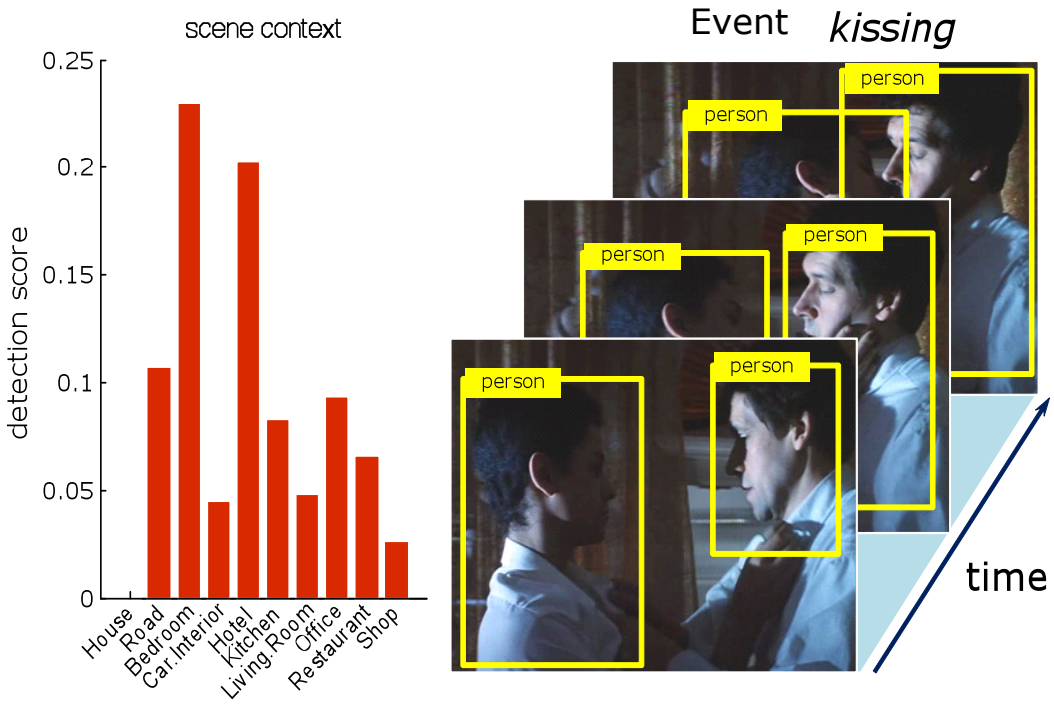

[Jiang, Li, Chang, TSCVT 2011]

Action-Scene-Object

- Identify relevant objects, scenes from a few training examples (~ 10)
- Learn relationships for action
  - accuracy much better than raw classifier

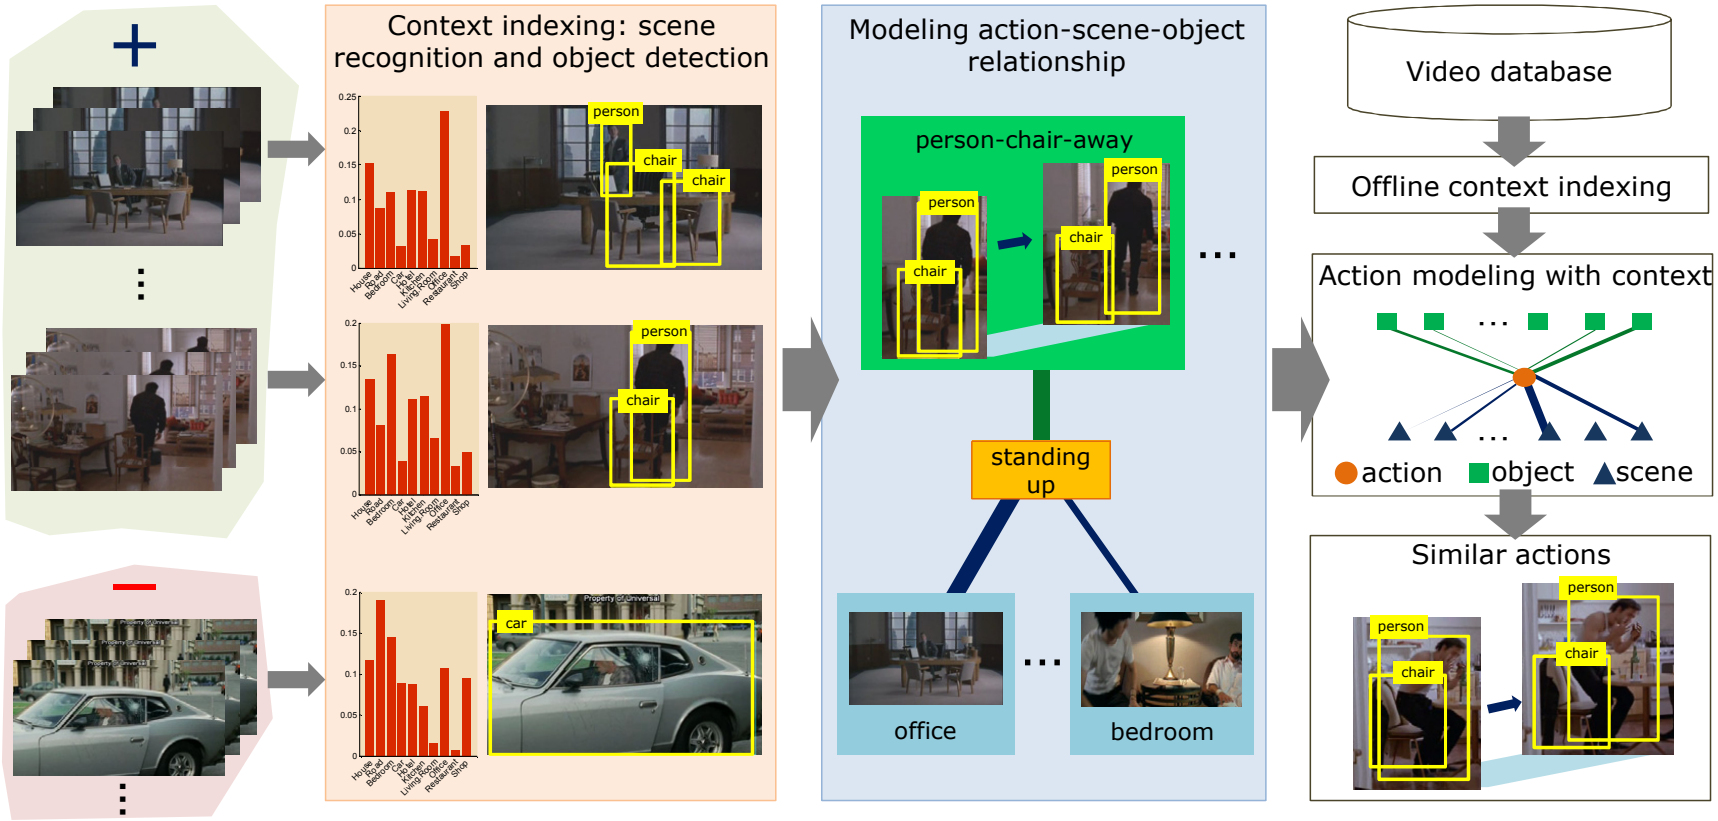

4. Future Work

- Existing joint Audio-Visual atoms are based on simple co-occurrence
  - no temporal structure
- Synchrony?
  - too hard to detect in web video
- Causality?
  - e.g. simple ordering
  - "Causal Audio-Video Atoms' CAVAs

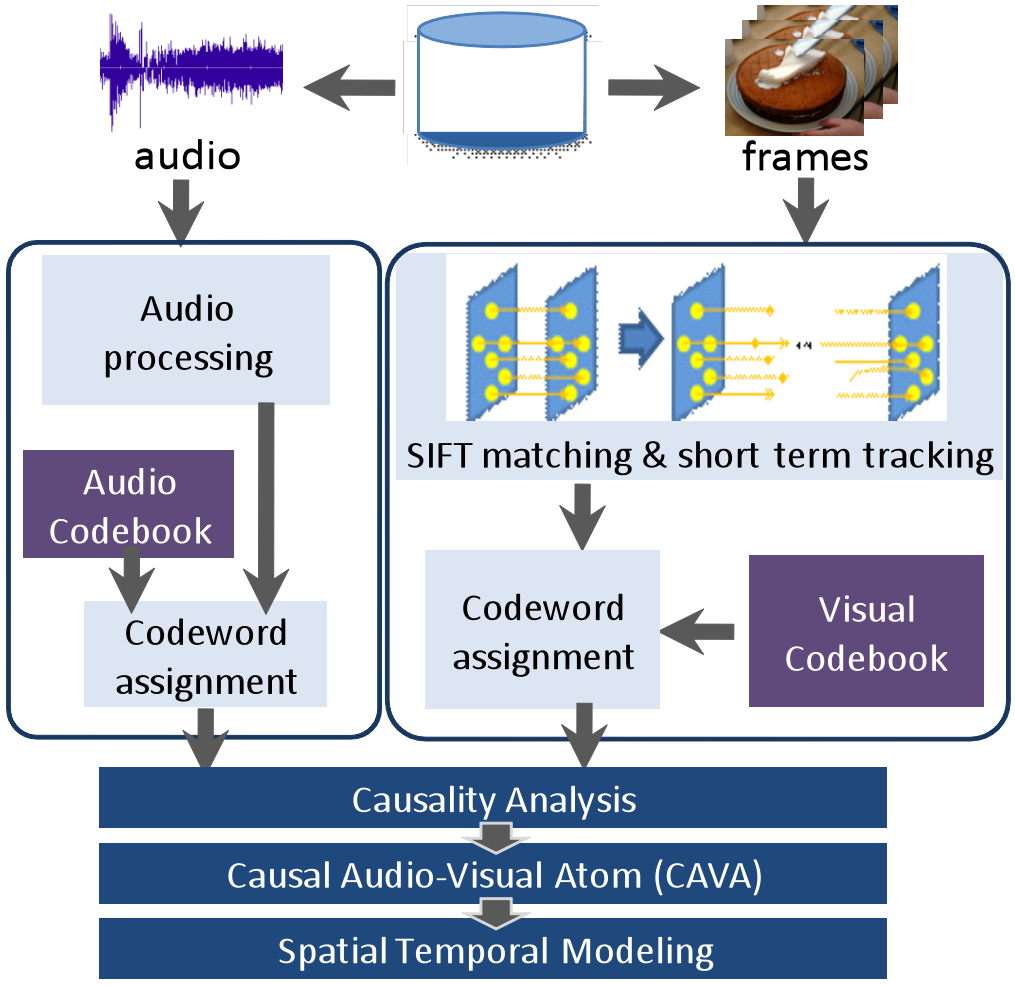

Summary

- Web video analysis
  - desperate need for automatic analysis
  - must be in terms of objects, scenes, actions
- Joint Audio-Visual Atoms
  - object-related codebooks for audio, video
  - MIL of all possible combinations to find cues
- Labeled Data
  - Mechanical Turk quickly labels web video examples
- Context-Based Action Detection
  - uses mature existing object and scene detectors
- Better "Causal Audio-Visual Atoms"