# Visual Recognition: Examples of Graphical Models 

Raquel Urtasun

TTI Chicago

March 8, 2012

## Example: Segmentation from Scribles

## ( $\mathrm{n}=$ number of pixels )


$x \in\{0,1\}^{n}$

$$
P(x \mid z)=P(z \mid x) \quad P(x) / P(z) \sim P(z \mid x) P(x)
$$

Posterior
Probability

Likelihood
(data-dependent)

Prior
(data- independent)
(MAP Solution) $x^{*}=\arg \max P(x \mid z)=\arg \min E(x)$ $x$
$x$

## Image Segmentation



$$
\begin{aligned}
& \text { Posterior } \text { Likelihood } \\
& P(x \mid z) P(z \mid x) \\
& \downarrow \\
& \prod_{x_{i}} P\left(z_{i} \mid x_{i}\right)
\end{aligned}
$$

Prior
$P(x)$
[Source: P. Kohli]

## Likelihood $\quad P(x \mid z) \sim P(z \mid x) P(x)$



$$
P(z \mid x)=F_{G M M}(z, x)
$$



Red
[Source: P. Kohli]

## Likelihood $\quad P(x \mid z) \sim P(z \mid x) P(x)$


$\log P\left(z_{i} \mid x_{i}=0\right)$


$$
P\left(z_{i} \mid x_{i}=1\right)
$$

## MAP Solution

$$
\begin{aligned}
x^{*} & =\underset{x}{\operatorname{argmax}} P(z \mid x) \\
& =\underset{x}{\operatorname{argmax}} \prod_{x_{i}} P\left(z_{i} \mid x_{i}\right)
\end{aligned}
$$


[Source: P. Kohli]

## Image Segmentation



$$
\begin{array}{ll}
\text { Posterior } & \text { Likelihood } \\
P(x \mid z)= & P(z \mid x)
\end{array}
$$

## Prior <br> $P(x)$ <br> $\prod f\left(x_{i}, x_{j}\right)$ <br> $x_{i}, x_{j}$

Encourages consistency between labelling of adjacent pixels
[Source: P. Kohli]

## Prior

## $P(x \mid z) \sim P(z \mid x) P(x)$



$$
\begin{aligned}
P(x) & =\prod_{i, j \in N} f_{i j}\left(x_{i}, x_{j}\right) \\
& =\prod_{i, j \in N} \exp \left\{-\left|x_{i}-x_{j}\right|\right\} \quad \text { "MRF Ising prior" }
\end{aligned}
$$

## Posterior and Energy Functions

$$
E(x, z, w)=\sum_{i} \theta_{i}\left(x_{i}, z_{i}\right)+w \sum_{i, j} \theta_{i j}\left(x_{i}, x_{j}, z_{i}, z_{j}\right)
$$

Energy

## Results of the Ising Model


[Source: P. Kohli]

## Conditional Random Fields

$$
P(x \mid z)=\prod_{x_{i}} P\left(z_{i} \mid x_{i}\right) \quad \prod_{x_{i}, x_{j}} P\left(x_{i}, x_{j}, z_{i}, z_{j}\right)
$$

$$
E(x, z, w)=\sum_{i} \theta_{i}\left(x_{i}, z_{i}\right)+w \sum_{i, j} \theta_{i j}\left(x_{i}, x_{j}, z_{i}, z_{j}\right)
$$

[Boykov and Jolly ' 01] [Blake et al. '04] [Rother, Kolmogorov and Blake '04]
[Source: P. Kohli]

## Conditional Random Fields

$$
E(x, z, w)=\sum_{i} \theta_{i}\left(x_{i}, z_{i}\right)+w \sum_{i, j} \theta_{i j}\left(x_{i}, x_{j}, z_{i}, z_{j}\right)
$$




Pairwise Cost

[Boykov and Jolly ' 01] [Blake et al. '04] [Rother, Kolmogorov and Blake '04] [Source: P. Kohli]

Conditional Random Fields

$$
E(x, z, w)=\sum_{i} \theta_{i}\left(x_{i}, z_{i}\right)+w \sum_{i, j} \theta_{i j}\left(x_{i}, x_{j}, z_{i}, z_{j}\right)
$$



Pairwise Cost

[Boykov and Jolly ' 01] [Blake et al. '04] [Rother, Kolmogorov and Blake '04]

## Example: Supervised Semantic Segmentation

- Assign a label to every pixel



## Different Approaches



## Building Unitary Potentials



Image Window (W)
Pixel to be classified (P)

## Image



## Segmentation

## Image Segmentation

$$
\left.\begin{array}{c}
\mathrm{n}= \\
\mathrm{E} \text { number of pixels } \\
\\
\\
\\
0 \rightarrow f(0,1\}^{n} \rightarrow R
\end{array}\right) 1 \rightarrow b g
$$

$$
E(X)=\sum_{i} c_{i} x_{i}+\sum_{i, j} d_{i j}\left|x_{i}-x_{j}\right|
$$



Image



Segmentation
[Boykov and Jolly ' 01] [Blake et al. '04] [Rother, Kolmogorov and Blake '04]

## High order patch potentials



## Patch Dictionary (Tree)


[Source: P. Kohli]

## Image Segmentation

$$
\begin{aligned}
& E:\{0,1\}^{n} \rightarrow R \\
& 0 \rightarrow f g, 1 \rightarrow b g
\end{aligned}
$$

$$
E(X)=\sum_{i} c_{i} x_{i}+\sum_{i, j} d_{i j}\left|x_{i}-x_{j}\right|+\sum_{p} h_{p}\left(X_{p}\right)
$$

$$
h\left(X_{p}\right)= \begin{cases}C_{1} & \text { if } x_{i}=0, i \in p \\ C_{\text {max }} & \text { otherwise }\end{cases}
$$


[Kohli et al. 'o7]
[Source: P. Kohli]

## Image Segmentation

n = number of pixels

$$
\begin{aligned}
& E:\{0,1\}^{n} \rightarrow R \\
& 0 \rightarrow f g, 1 \rightarrow b g
\end{aligned}
$$

$E(X)=\sum_{i} c_{i} x_{i}+\sum_{i, j} d_{i j}\left|x_{i}-x_{j}\right|+\sum_{p} h_{p}\left(X_{p}\right)$


Image


Pairwise Segmentation


Final Segmentation
[Kohli et al. 'o7]
[Source: P. Kohli]

## Minimizing higher order terms

## Higher Order Submodular Functions

## Exact <br> Transformation <br> 

## Pairwise Submodular Function

Billionnet and M. Minoux [DAM 1985]
Kolmogorov \& Zabih [PAMI 2004]
Freedman \& Drineas [CVPR2005]
Kohli Kumar Torr [CVPRz2007, PAMI 2008] Kohli Ladicky Torr [CVPR 2008, IJCV 2009] Ramalingam Kohli Alahari Torr [CVPR 2008] Zivny et al. [CP 2008]

[Source: P. Kohli]

## Qualitative Results

## Image <br> (MSRC-21)

## Pairwise CRF

## Higher order CRF

## Ground Truth


[Source: P. Kohli]

## Example: Holistic Scene Understanding

For an image we would like to reason about:

- Objects: which class, where, how many?
- Segmentation: which semantic label does each pixel take?
- Scene classification: which scene am I looking at?



## Why Holistic?

Let's use a classifier for each task independently. What's in the patch?

- detector: bird
- seg classif.: water
- scene: boat


## Why Holistic?

Let's use a classifier for each task independently. What's in the patch?


- detector: bird
- seg classif.: water
- scene: boat


## Why Holistic?

Let's use a classifier for each task independently. What's in the patch?


- detector: bird
- seg classif.: water
- scene: boat



## Why Holistic?

Let's use a classifier for each task independently. What's in the patch?


- detector: bird
- seg classif.: water
- scene: boat



## Holistic Scene Understanding

We want to reason about the scene as a whole.

- Joint inference of scene type, 2D objects and semantic segmentation
- Efficient learning and inference with structure prediction



## Compact Holistic Model

- Define the problem as hierarchical CRF
- Compatibility potentials + evidence + shape prior



## Compact Holistic Model

We define the problem as a holistic conditional random field

$$
p(\mathbf{a})=p(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{b}, \mathbf{s})=\frac{1}{Z} \prod_{i} \psi_{i}\left(\mathbf{a}_{i}\right) \prod_{\alpha} \psi_{\alpha}\left(\mathbf{a}_{\alpha}\right)
$$

where $\mathbf{a}=(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{b}, \mathbf{s})$ represents the set of all random variables

- $x_{i} \in\{1, \ldots, \mathcal{C}\}$ : class label of the $i$-th super-pixel (first layer of the hierarchy)
- $y_{i} \in\{1, \ldots, \mathcal{C}\}$ : class label of the i -th super-segment (second layer)
- $b_{i} \in\{0,1\}$ : binary variable indicating whether an object detection is on or off
- $z_{i} \in\{0,1\}$ : binary variable indicating the presence of class $i$ in the image
- $s \in\{1, \ldots, \mathcal{S}\}$ : scene type label


## Compact Holistic Model

- Learning the weights $w_{i}$, where $w_{i} \phi_{i}=\log \left(\psi_{i}\right)$, is done with primal-dual approximated learning algorithm
- Joint inference is performed by computing the MAP estimate:

$$
\max _{\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{b}, \mathbf{s}} \frac{1}{Z} \prod_{i} \psi_{i}\left(\mathbf{a}_{i}\right) \prod_{\alpha} \psi_{\alpha}\left(\mathbf{a}_{\alpha}\right)
$$

We use a convergent message-passing algorithm without restriction to submodularity and potential specific moves

## Unitary Potentials

- Super-pixel and super-segment:
$\phi_{i}\left(x_{i}\right)$ and $\phi_{j}\left(y_{j}\right)$ : average of TextonBoost pixel potentials inside each region
- Object detection:

$$
\phi_{l}^{B B o x}\left(b_{i}\right)= \begin{cases}\sigma\left(r_{i}-\lambda_{l}\right) & \text { if } b_{i}=1 \wedge c_{i}=l \\ 0 & \text { otherwise }\end{cases}
$$

Here $r_{i}$ is the score from Felzenswalb et al. detector, $\lambda_{l}$ is the threshold of the detector for that class, $c_{i}$ is the detector class, and $\sigma(x)=1 /(1+\exp (-1.5 x))$ is a logistic function that converts the classifier score into probability.

- Scene:

$$
\phi^{S c e n e}(s=k)=\sigma\left(t_{k}\right)
$$

where $t_{k}$ denotes the classifier score for scene class $k$

## Pairwise potentials

- Super-pixel - Super-segment: we use the $P^{n}$ potentials by Kohli et al.,CVPR'07:

$$
\phi_{i, j}\left(x_{i}, y_{j}\right)= \begin{cases}-\infty & \text { if } x_{i} \neq y_{j} \\ 0 & \text { otherwise }\end{cases}
$$

- Super-segment - Class:

$$
\phi_{i, j}\left(y_{i}, z_{j}\right)= \begin{cases}-\infty & \text { if } y_{i}=j \wedge z_{j}=0 \\ 0 & \text { otherwise }\end{cases}
$$

- Class - Scene:
$\phi^{S C}\left(s, z_{j}\right)= \begin{cases}f_{s, z_{j}} & \text { if } z_{j}=1 \wedge f_{s, z_{j}}>0 \\ -\tau & \text { if } z_{j}=1 \wedge f_{s, z_{j}}=0 \\ 0 & \text { otherwise. }\end{cases}$

where $f_{s, z_{j}}$ represents the probability of occurrence of class $z_{j}$ for scene type $s$


## Pairwise potentials

- Detection - Class:

$$
\phi_{i, j}^{B \text { Class }}\left(\beta_{i}, b_{i}, z_{j}\right)= \begin{cases}-\infty & \text { if } z_{j}=0 \wedge c_{i}=j \wedge b_{i}=1 \\ 0 & \text { otherwise }\end{cases}
$$

- Detection - Super-pixel (shape prior):


aeroplane

chair

car

bird

cow

flower


## Loss function

Structure prediction problems require a specification for the loss. We define it as a weighted sum of task-specific losses, each of order at most 2.

- Super-pixel and super-segment layers: loss is the total number of pixels that were wrongly predicted.
- Class: $0-1$ loss
- Scene: $0-1$ loss
- Detection:

$$
\Delta_{B}\left(b_{i}, \hat{b}_{i}\right)= \begin{cases}1-\frac{\text { intersection }}{\text { Lnion }} & \text { if } b_{i}=1 \\ \frac{\text { intersection }}{\text { union }} & \text { otherwise }\end{cases}
$$

## Inference example

## iteration 0000 , accuracy $=82.36 \%$



## Joint Inference Results



## Segmentation Results MSRC-21

[J. Yao, S. Fidler and R. Urtasun, CVPR12]
Table: MSRC-21 segmentation results

|  | - | $\begin{aligned} & \tilde{\pi} \\ & \frac{\tilde{\omega}}{60} \end{aligned}$ | $\begin{aligned} & \mathbb{\#} \\ & \pm \end{aligned}$ | $3$ | $\begin{aligned} & \stackrel{\circ}{\otimes} \\ & \frac{N}{n} \end{aligned}$ | $\frac{\pi}{n}$ | $\begin{aligned} & \stackrel{\sim}{\Gamma} \\ & \frac{\pi}{0} \\ & \hline \stackrel{0}{0} \\ & \stackrel{0}{0} \end{aligned}$ | $\begin{aligned} & \stackrel{ \pm}{ \pm} \\ & \stackrel{N}{0} \end{aligned}$ | $\underset{\substack{0 \\ \hline \multirow{2}{*}{\hline}\\ \hline}}{ }$ | 牙 | $\frac{0}{\grave{u}}$ | $\begin{aligned} & \bar{\otimes} \\ & \sum_{0}^{2} \\ & \stackrel{0}{4} \end{aligned}$ | $\frac{c}{60}$ | 믈 | $\begin{aligned} & \text { ㄴ } \\ & \text { O } \end{aligned}$ | $\begin{aligned} & \cdot \frac{1}{\pi} \\ & \frac{1}{U} \end{aligned}$ | $\begin{aligned} & \overline{0} \\ & \underline{0} \end{aligned}$ | 范 | $\begin{aligned} & 60 \\ & 0 \\ & 0 \end{aligned}$ | $\begin{aligned} & \text { ते } \\ & \text { oे } \end{aligned}$ | $\begin{aligned} & \stackrel{N}{0} \\ & 0 \end{aligned}$ | ¢ | 7 $\frac{0}{0}$ $\frac{0}{60}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | origMSRC dataset |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
| Shotton et al | 49 | 88 | 79 | 97 | 97 | 78 | 82 | 54 | 87 | 74 | 72 | 74 | 36 | 24 | 93 | 51 | 78 | 75 | 35 | 66 | 18 | 67 | 72 |
| Jiang and Tu | 53 | 97 | 83 | 70 | 71 | 98 | 75 | 64 | 74 | 64 | 88 | 67 | 46 | 32 | 92 | 61 | 89 | 59 | 66 | 64 | 13 | 68 | 78 |
| Pixel-CRF | 73 | 92 | 85 | 75 | 78 | 92 | 75 | 76 | 86 | 79 | 87 | 96 | 95 | 31 | 81 | 34 | 84 | 53 | 61 | 60 | 15 | 72 | 81 |
| Hierarch. CRF | 80 | 96 | 86 | 74 | 87 | 99 | 74 | 87 | 86 | 87 | 82 | 97 | 95 | 30 | 86 | 31 | 95 | 51 | 69 | 66 | 9 | 75 | 86 |
| HCRF+Coocc. | 74 | 98 | 90 | 75 | 86 | 99 | 81 | 84 | 90 | 83 | 91 | 98 | 75 | 49 | 95 | 63 | 91 | 71 | 49 | 72 | 18 | 77.8 | 86.5 |
| Harmony pot. | 60 | 78 | 77 | 91 | 68 | 88 | 87 | 76 | 73 | 77 | 93 | 97 | 73 | 57 | 95 | 81 | 76 | 81 | 46 | 56 | 46 | 75 | 77 |
| Segm.+Class | 72 | 98 | 91 | 77 | 82 | 93 | 86 | 86 | 82 | 82 | 93 | 97 | 71 | 50 | 96 | 59 | 88 | 78 | 51 | 67 | 0 | 76.2 | 85.1 |
| Det 15 class | 69 | 98 | 90 | 78 | 86 | 93 | 88 | 83 | 90 | 83 | 94 | 97 | 73 | 50 | 96 | 71 | 89 | 79 | 54 | 64 | 8 | 77.8 | 85.3 |
| full model | 71 | 98 | 90 | 79 | 86 | 93 | 88 | 86 | 90 | 84 | 94 | 98 | 76 | 53 | 97 | 71 | 89 | 83 | 55 | 68 | 17 | 79.3 | 86.2 |

## Detection and Scene Classification Results

［J．Yao，S．Fidler and R．Urtasun，CVPR12］
Table：MSRC－21 object detection results

|  | Z | $\begin{aligned} & \stackrel{\text { O}}{む} \\ & \frac{ֻ}{\omega} \end{aligned}$ | $\begin{aligned} & \underset{\sim}{\tilde{\sigma}} \\ & \frac{\pi}{0} \\ & \stackrel{O}{0} \\ & \underset{\sim}{\omega} \end{aligned}$ | $\underset{\sim}{\text { U }}$ | ¢ัర | $\begin{aligned} & \frac{0}{U} \\ & \frac{0}{0} \end{aligned}$ | $\begin{aligned} & \text { 亠凶 } \\ & \stackrel{3}{3} \\ & \frac{0}{4} \end{aligned}$ | $\frac{c}{60}$ | 믈 | $\begin{aligned} & \text { 등 } \end{aligned}$ | $\begin{aligned} & \cdot \frac{1}{\pi} \\ & \frac{\pi}{U} \end{aligned}$ | $\stackrel{\sim}{0}$ | $\begin{aligned} & 60 \\ & 0 \\ & 0 \end{aligned}$ | ते | $\begin{aligned} & \stackrel{\rightharpoonup}{\circ} \\ & \hline 0 \end{aligned}$ | 0 0 00 0 0 0 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Recall at equal FPPI |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
| FPPI rate | 0.03 | 0.02 | 0.00 | 0.01 | 0.05 | 0.03 | 0.04 | 0.02 | 0.02 | 0.01 | 0.00 | 0.02 | 0.04 | 0.04 | 0.02 | 0.02 |
| LSVM | 84.6 | 73.9 | 84.6 | 59.4 | 50.0 | 63.6 | 16.9 | 40.0 | 16.2 | 23.7 | 50.0 | 20.0 | 20.0 | 43.2 | 18.8 | 44.3 |
| cont．LSVM | 76.9 | 17.4 | 23.1 | 50.0 | 50.0 | 68.2 | 15.3 | 40.0 | 8.1 | 18.4 | 50.0 | 30.0 | 33.3 | 38.6 | 21.9 | 36.1 |
| Detection | 88.5 | 78.3 | 100.0 | 43.8 | 52.4 | 63.6 | 20.3 | 53.3 | 16.2 | 42.1 | 62.5 | 50.0 | 26.7 | 38.6 | 6.3 | 49.5 |
| full model | 88.5 | 82.6 | 100.0 | 46.9 | 52.4 | 63.6 | 20.3 | 53.3 | 16.2 | 44.7 | 62.5 | 40.0 | 26.7 | 38.6 | 12.5 | 49.9 |
|  |  |  |  |  |  |  | Avera | ge Prec | ision |  |  |  |  |  |  |  |
| LSVM | 78.6 | 76.5 | 96.2 | 56.4 | 54.1 | 61.7 | 19.9 | 45.0 | 18.5 | 30.0 | 59.2 | 31.4 | 28.0 | 45.5 | 22.1 | 48.2 |
| cont．LSVM | 75.8 | 37.0 | 85.1 | 58.2 | 52.1 | 60.8 | 19.1 | 38.5 | 12.3 | 28.6 | 60.5 | 32.1 | 32.1 | 41.7 | 26.2 | 44.0 |
| Detection | 78.1 | 72.7 | 100.0 | 45.5 | 53.1 | 60.9 | 22.9 | 48.9 | 18.2 | 42.9 | 63.6 | 46.0 | 27.3 | 34.3 | 9.1 | 48.2 |
| full model | 78.1 | 81.8 | 100.0 | 45.5 | 53.1 | 60.9 | 22.9 | 48.9 | 18.2 | 44.4 | 63.6 | 45.6 | 27.3 | 34.3 | 16.4 | 49.4 |

Table：MSRC－21 scene classification

|  | classifier | full m. |
| :---: | :---: | :---: |
| accuracy | 79.5 | $\mathbf{8 0 . 6}$ |

## More Results ...



Figure: Segmentation examples: (image, groundtruth, our holistic scene model)


Figure: Examples of failure modes.

## Let's talk about attributes

## Zero-shot learning

- Can I leaned what a mule is without seen a single instance if I know what horses and donkeys are?
- Traditional paradigm is not very appropiate

[Source: D. Parikh]


## Zero-shot learning

- Can I leaned what a mule is without seen a single instance if I know what horses and donkeys are?
- Traditional paradigm is not very appropiate

[Source: D. Parikh]


## Zero-shot learning

- Can I leaned what a mule is without seen a single instance if I know what horses and donkeys are?
- Traditional paradigm is not very appropiate

[Source: D. Parikh]


## Zero-shot learning

- Can I leaned what a mule is without seen a single instance if I know what horses and donkeys are?
- Traditional paradigm is not very appropiate

[Source: D. Parikh]


## Zero-shot learning

- Can I leaned what a mule is without seen a single instance if I know what horses and donkeys are?
- Traditional paradigm is not very appropiate

[Source: D. Parikh]


## Zero-shot learning

- Can I leaned what a mule is without seen a single instance if I know what horses and donkeys are?
- Traditional paradigm is not very appropiate

[Source: D. Parikh]


## Zero-shot learning

- Can I leaned what a mule is without seen a single instance if I know what horses and donkeys are?
- Traditional paradigm is not very appropiate

[Source: D. Parikh]


## Attributes

- Long history of attributes in vision, starting in 2007.
- They are typically simple classifiers
- The score of those classifiers is an alternative representation
- They are binary

| Is furry | Has four-legs |
| :---: | :---: |
| Legs shorter <br> than horses' | Tail longer <br> than donkeys' |

Has tail
[Oliva 2001] [Ferrari 2007] [Lampert 2009] [Farhadi 2009] [Kumar 2009] [Wang 2009] [Wang 2010] [Berg 2010] [Branson 2010] [Parikh 2010] [ICCV 2011] ...
[Source: D. Parikh]

## Attributes

- Long history of attributes in vision, starting in 2007.
- They are typically simple classifiers
- The score of those classifiers is an alternative representation
- They are binary

$$
\begin{array}{cc}
\text { Is furry } & \text { Has four-legs } \\
\text { Legs shorter } & \text { Tail longer } \\
\text { than horses' } & \text { than donkeys' }
\end{array}
$$

Has tail
[Source: D. Parikh]

## Attributes

- Some of them are relative

> Is furry

Has four-legs
Legs shorter than horses'

Tail longer than donkeys'

Has tail

## Image Search

- I want to ask about an image of Chicago
- This might bee too crowded for my taste



## Image Search

- I want to ask about an image of Chicago
- This might bee too crowded for my taste



## How do we think about attributes?



## How do we think about attributes?



## How do we think about attributes?



## How do we think about attributes?



## But it's easy to say...



## Relative Attributes [Parikh et al. 11]

Relative attributes

- Allow relating images and categories to each other
- Learn ranking function for each attribute

Novel applications

- Zero-shot learning from attribute comparisons
- Automatically generating relative image descriptions


## Learning Relative Attributes

## For each attribute $a_{m}$ ，open

Supervision is

$$
\begin{aligned}
& o_{m}:\{(\omega-\text { 签) })\}, \\
& S_{m}:\{\text { 長~目\}, }\}
\end{aligned}
$$

## Learning Relative Attributes

$$
\text { Learn a scoring function } r_{m}\left(\boldsymbol{x}_{\boldsymbol{i}}\right)=\boldsymbol{w}_{\boldsymbol{m}}^{T} \boldsymbol{x}_{\boldsymbol{i}}
$$

that best satisfies constraints:

$$
\begin{aligned}
& \forall(i, j) \in O_{m}: \boldsymbol{w}_{\boldsymbol{m}}^{T} \boldsymbol{x}_{\boldsymbol{i}}>\boldsymbol{w}_{\boldsymbol{m}}^{T} \boldsymbol{x}_{\boldsymbol{j}} \\
& \forall(i, j) \in S_{m}: \boldsymbol{w}_{\boldsymbol{m}}^{T} \boldsymbol{x}_{\boldsymbol{i}}=\boldsymbol{w}_{\boldsymbol{m}}^{T} \boldsymbol{x}_{\boldsymbol{j}}
\end{aligned}
$$

## Learning Relative Attributes

Max-margin learning to rank formulation

$$
\begin{array}{cc}
\min \quad\left(\frac{1}{2}\left\|\boldsymbol{w}_{\boldsymbol{m}}^{T}\right\|_{2}^{2}+C\left(\sum \xi_{i j}^{2}+\sum \gamma_{i j}^{2}\right)\right) \\
\text { s.t } \quad \boldsymbol{w}_{\boldsymbol{m}}^{T}\left(\boldsymbol{x}_{\boldsymbol{i}}-\boldsymbol{x}_{\boldsymbol{j}}\right) \geq 1-\xi_{i j}, \forall(i, j) \in O_{m} \\
\left|\boldsymbol{w}_{\boldsymbol{m}}^{T}\left(\boldsymbol{x}_{\boldsymbol{i}}-\boldsymbol{x}_{\boldsymbol{j}}\right)\right| \leq \gamma_{i j}, \forall(i, j) \in S_{m} \\
\xi_{i j} \geq 0 ; \gamma_{i j} \geq 0 \\
\text { Based on [Joachims 2002] } \\
\text { Image } \rightarrow \text { Relative Attribute Score }
\end{array}
$$

## Zero Shot Learning

Training: Images from $S$ seen categories and Descriptions of $U$ unseen categories


Age: Hugh $\succ$ Clive $\succ$ Scarlett

Smiling:


Jared $\succ$ Miley


Miley $\succ$ Jared

Need not use all attributes, or all seen categories
Testing: Categorize image into one of $\mathbf{S}+\mathbf{U}$ categories

## Automatic Relative Description



Conventional binary description: not dense


Not dense:


## Automatic Relative Description



## Automatic Relative Description



## Results



## Relative (ours):

More natural than insidecity
Less natural than highway
More open than street Less open than coast

Has more perspective than highway Has less perspective than insidecity

## Results



## Relative (ours):

More natural than tallbuilding Less natural than forest

More open than tallbuilding Less open than coast

Has more perspective than tallbuilding

## Results

Binary (existing):
Not Young
BushyEyebrows
RoundFace
More Young than CliveOwen
Less Young than ScarlettJohansson
More BushyEyebrows than ZacEfron
Less BushyEyebrows than AlexRodriguez

## Human Studies: Which Image is described?



## Automatic Relative Image Description

## 18 subjects

Test cases: 10 OSR, 20 PubFig


There is much more... for that you need to do a PhD on vision ;)

