Discriminative sparse representations
[Mairal, Bach, Ponce, Sapiro & Zisserman CVPR '08]

Let us consider 2 sets S_-, S_+ of signals representing 2 different classes.

Idea:

Each set should admit a specific dictionary best adapted to its reconstruction.

Discriminative sparse representations
[Mairal, Bach, Ponce, Sapiro & Zisserman CVPR '08]

Classification procedure for a signal $\mathbf{x} \in \mathbb{R}^n$:

$$\min(\mathbf{R}^\star(\mathbf{x}, \mathbf{D}_-), \mathbf{R}^\star(\mathbf{x}, \mathbf{D}_+))$$

where

$$\mathbf{R}^\star(\mathbf{x}, \mathbf{D}) = \min_{\boldsymbol{\alpha} \in \mathbb{R}^p} ||\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}||_2^2 \text{ s.t. } ||\boldsymbol{\alpha}||_0 \leq L.$$

Discriminative sparse representations
[Mairal, Bach, Ponce, Sapiro & Zisserman CVPR '08]

"Reconstructive" training

$$\begin{cases} \min_{\mathbf{D}_-} \sum_{i \in S_-} \mathbf{R}^\star(\mathbf{x}_i, \mathbf{D}_-) \\ \min_{\mathbf{D}_+} \sum_{i \in S_+} \mathbf{R}^\star(\mathbf{x}_i, \mathbf{D}_+) \end{cases}$$

$$\mathbf{R}^\star(\mathbf{x}, \mathbf{D}) = \min_{\boldsymbol{\alpha} \in \mathbb{R}^p} ||\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}||_2^2 \text{ s.t. } ||\boldsymbol{\alpha}||_0 \leq L.$$

Discriminative sparse representations

[Mairal, Bach, Ponce, Sapiro & Zisserman CVPR '08]

"Discriminative" training

$$\min_{\mathbf{D}_-,\mathbf{D}_+} \sum_i \mathcal{C}\Big(\lambda z_i \big(\mathbf{R}^\star(\mathbf{x}_i, \mathbf{D}_-) - \mathbf{R}^\star(\mathbf{x}_i, \mathbf{D}_+)\big)\Big),$$

where $z_i \in \{-1, +1\}$ is the label of $\mathbf{x}_i$.

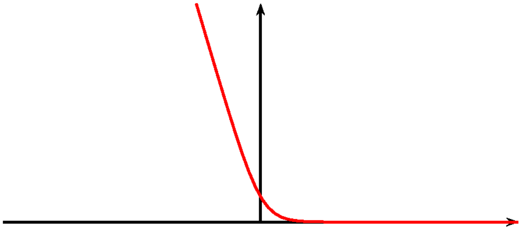

Logistic loss function

Discriminative sparse representations
[Mairal, Bach, Ponce, Sapiro & Zisserman CVPR '08]

Mixed approach

$$\min_{\mathbf{D}_-,\mathbf{D}_+} \sum_i \mathcal{C}\Big(\lambda z_i\big(\mathbf{R}^\star(\mathbf{x}_i,\mathbf{D}_-) - \mathbf{R}^\star(\mathbf{x}_i,\mathbf{D}_+)\big)\Big) + \mu \mathbf{R}^\star(\mathbf{x}_i,\mathbf{D}_{z_i}),$$

where $z_i \in \{-1,+1\}$ is the label of $\mathbf{x}_i$.

Discriminative sparse representations

Keys of the optimization framework

- Alternation of sparse coding and dictionary updates, as in MOD and K-SVD.
- Continuation path with decreasing values of μ.
- Greedy procedure to address the NP-hard sparse coding problem. [Weisbert '80], [Mallat '93].
- or LARS to address a convex relaxation of the sparse coding using the ℓ_1 norm. [Efron '00].
- Use softmax instead of logistic regression for $N > 2$ classes.

Discriminative sparse representations

New feature space: Use one classifier per scale

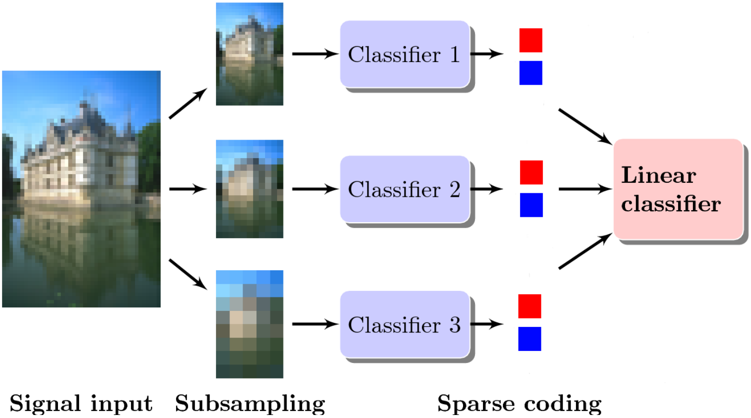

Discriminative sparse representations

New feature space: Use reconstruction error paths

$$\mathbf{R}^{\star}(\mathbf{x}, \mathbf{D}) = \min_{\boldsymbol{\alpha} \in \mathbb{R}^p} ||\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}||_2^2 \text{ s.t. } ||\boldsymbol{\alpha}||_0 \leq L.$$

After the learning of the dictionaries, why not use different values for L?

Discriminative sparse representations

New feature space: Use reconstruction error paths

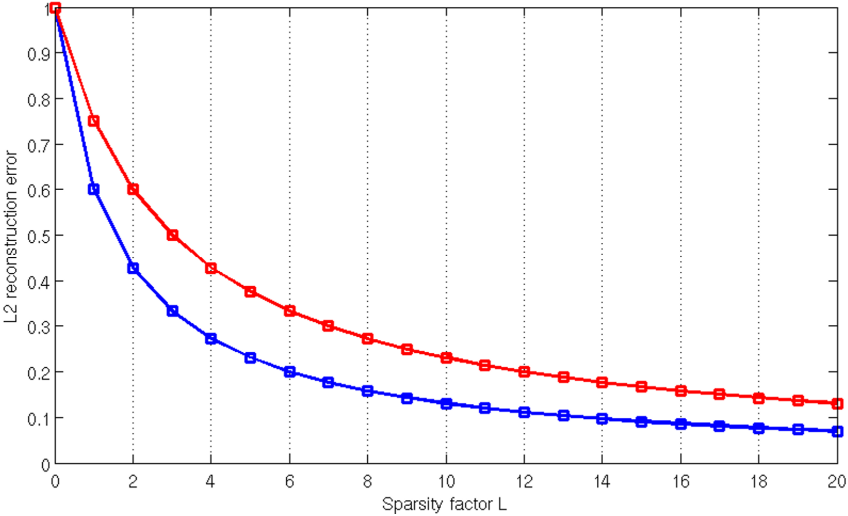

Discriminative sparse representations

New feature space

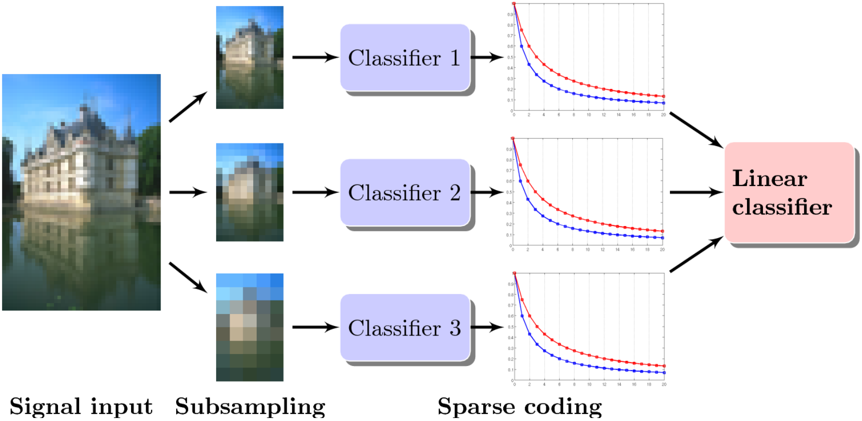

Discriminative sparse representations

Some related works

- Generative models: [Wright et al. '07],[Grosse et al. '07],[Huang & Aviyente '06]
- Another discriminative model: [Rodriguez & Sapiro '08]
- Textons: [Malik et al. '99]
- Discriminative codebooks: [Lazebnik & Raginsky '08], [Winn et al. '05]
- pLSA: [Hoffman '01]
- Neural nets: [Lecun, Hinton ∼90s-today.]

Discriminative sparse representations

[Mairal, Bach, Ponce, Sapiro & Zisserman NIPS '08]

Idea: Using the coefficients as features

Work in progress...

1. Sparse representations for image restoration
2. Discriminative sparse representations for computer vision
3. Applications to recognition and image interpretation

Applications to computer vision

Texture segmentation

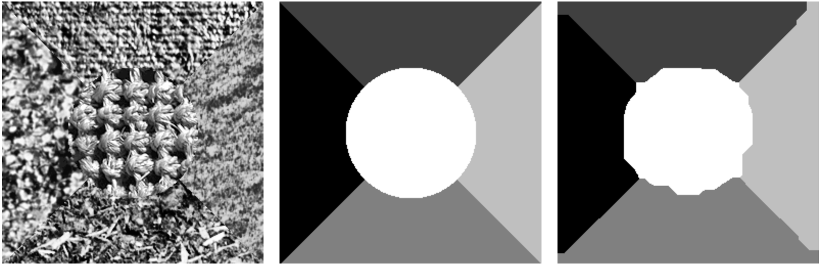

Applications to computer vision

Texture segmentation

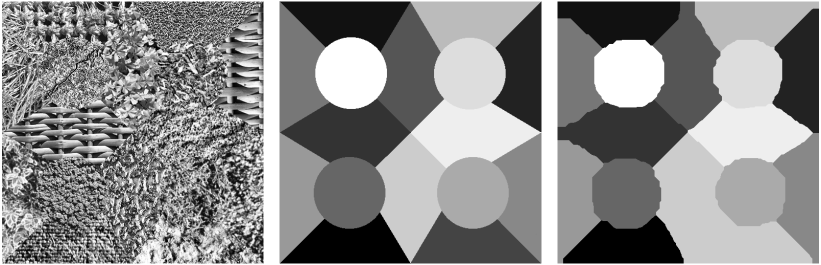

Applications to computer vision

Pixelwise classification

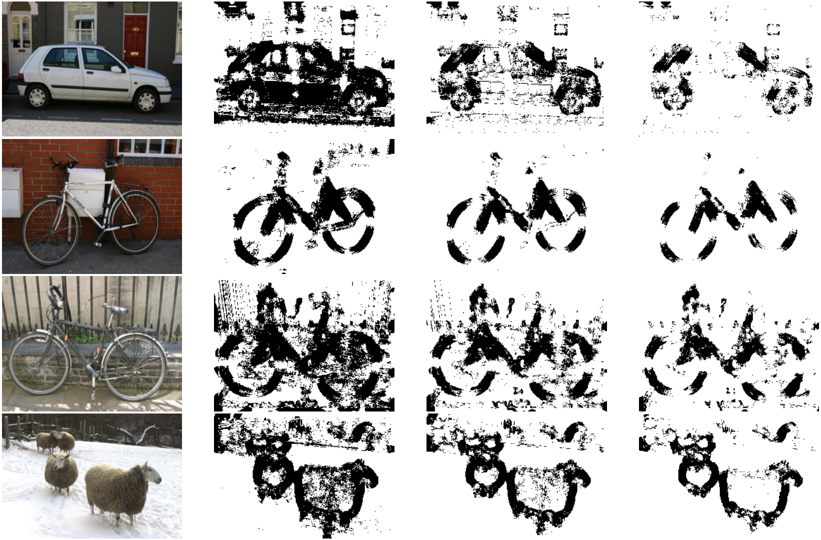

Applications to computer vision

Example of learned dictionaries

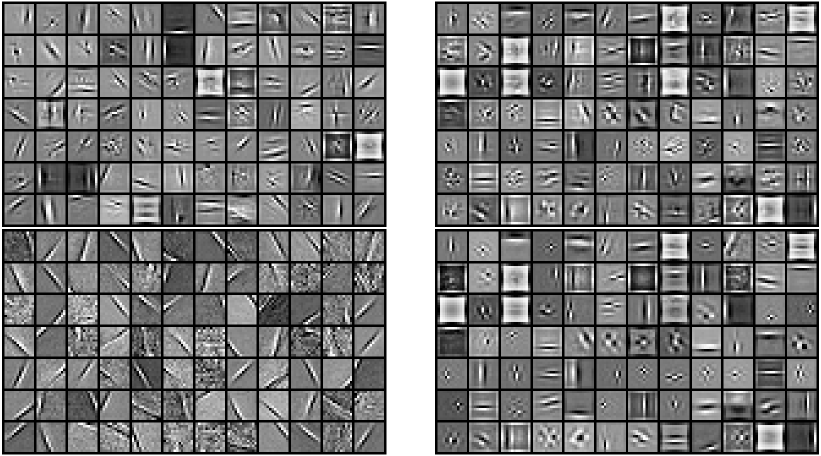

Figure: Top: reconstructive, Bottom: discriminative, Left: Background, Right: Bicycle

Applications to computer vision

Example of object detection, qualitative evaluation

Applications to computer vision

Example of object detection, quantitative evaluation

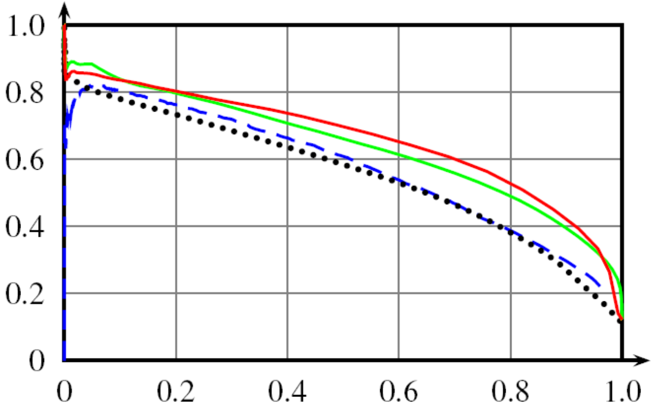

Figure: comparison with Tuytelaars '07 and Pantofaru & Schmidt '06

Applications to computer vision

Discriminative dictionaries for edge detection

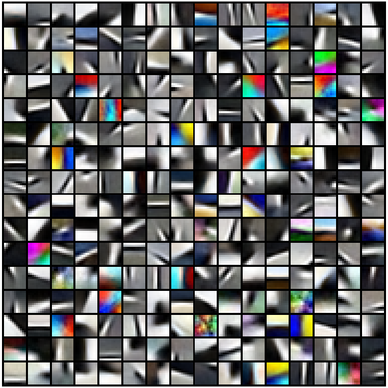

Good edges

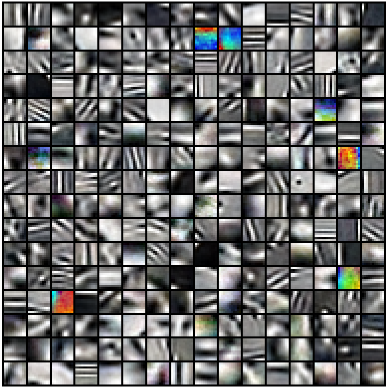

Bad edges

Applications to computer vision

Berkeley segmentation benchmark

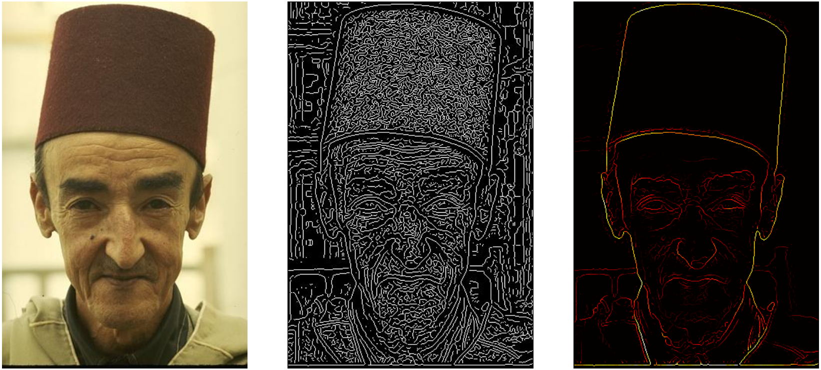

Raw edge detection on the right

Applications to computer vision
Berkeley segmentation benchmark

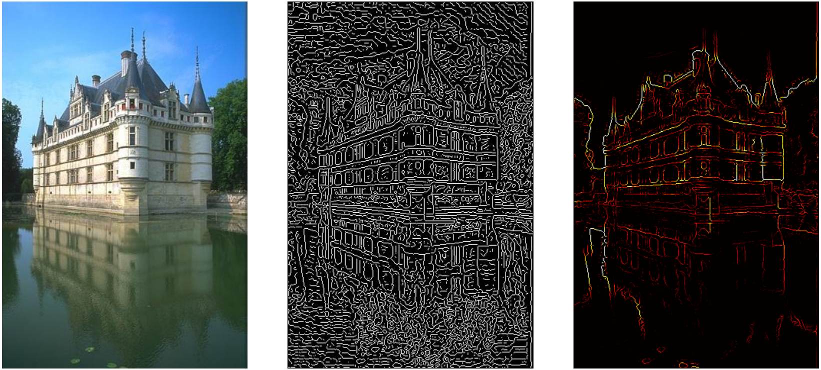

Raw edge detection on the right

Applications to computer vision

Berkeley segmentation benchmark

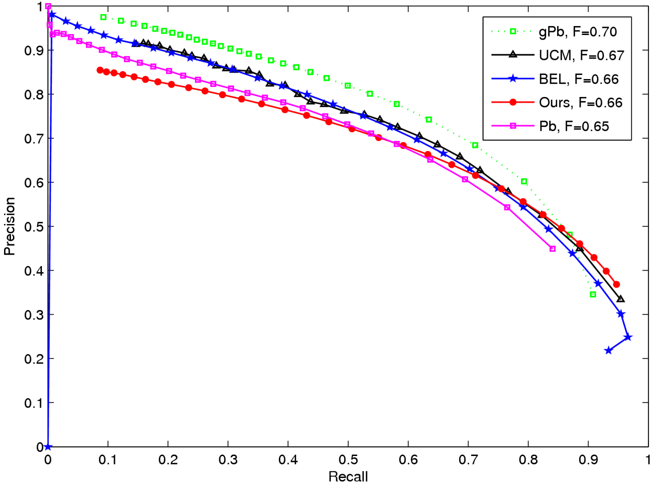

Applications to computer vision

Contour-based classifier: [Leordeanu, Hebert & Sukthankar '07]

Is there a bike, a motorbike, a car or a person on this image?

Application to edge detection and classification

Question:

Can a local analysis of these edges help this classifier?

Answer: Yes !

1. Train class-specific local classifiers of edges.
2. Given an edge map, obtain one class-specific edge map per class.
3. Train the contours-based classifier on these new maps.

Applications to computer vision

Prefiltering of edges map with class-specific edge detection

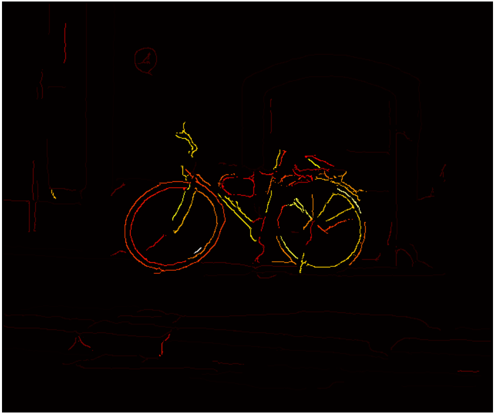

Is there a bike?

Applications to computer vision

Prefiltering of edges map with class-specific edge detection

Is there a car?

Applications to computer vision
Prefiltering of edges map with class-specific edge detection

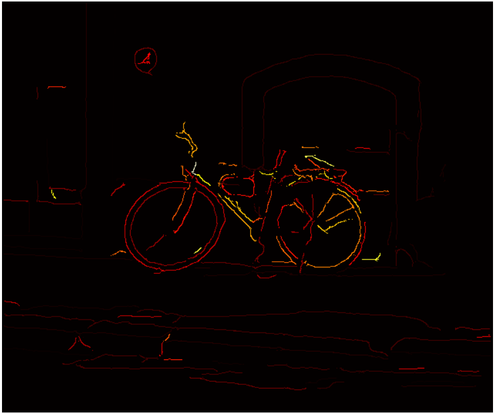

Is there a motobike?

Applications to computer vision

Prefiltering of edges map with class-specific edge detection

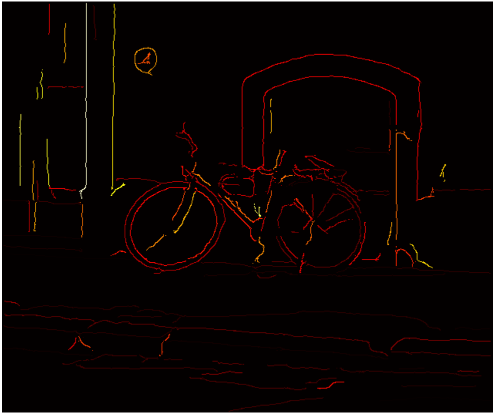

Is there a person?

Applications to computer vision

Prefiltering of edges map with class-specific edge detection

Is there a bike, a motorbike, a car or a person on this image?

Applications to computer vision

Prefiltering of edges map with class-specific edge detection

Is there a bike?

Applications to computer vision

Prefiltering of edges map with class-specific edge detection

Is there a car?

Applications to computer vision

Prefiltering of edges map with class-specific edge detection

Is there a motobike?

Applications to computer vision

Prefiltering of edges map with class-specific edge detection

Is there a person?

Applications to computer vision

performance gain due to the prefiltering

Ours + [Leordeanu '07]	[Leordeanu '07]	[Winn '05]
96.8%	89.4%	76.9%

Recognition rates for the same experiment as [Winn '05] on VOC 2005.

Applications to computer vision

performance gain due to the prefiltering

Category	Ours+[Leordeanu '07]	[Leordeanu '07]
Aeroplane	71.9%	61.9%
Boat	67.1%	56.4%
Cat	82.6%	53.4%
Cow	68.7%	59.2%
Horse	76.0%	67%
Motorbike	80.6%	73.6%
Sheep	72.9%	58.4%
Tvmonitor	87.7%	83.8%
Average	75.9%	64.2 %

Recognition performance at equal error rate for 8 classes on a subset of images from Pascal 07.

Applications to computer vision

Some related works on edges

- Pb: [Martin et al. '04]
- UCM: [Arbelaez '06]
- BEL: [Dollar et al. '06]
- gPb: [Maire et al. '08]
- Class-specific edge detection: [Prasad et al. '06]

A few conclusions

- Sparse representations are a powerful tool for image restoration.
- The learning of sparse representations should be discriminative for recognition tasks.
- Discriminative sparse representations are well adapted to some computer vision tasks such as edge analysis.

Some future directions

- Learning jointly global and local classifiers.
- Learning sparse representations for bags of features.
- Exploiting the coefficients of the sparse decompositions: [Mairal, Bach, Ponce, Sapiro & Zisserman NIPS '08].