How to use mixture models on patches for solving inverse problems in imaging?

Antoine Houdard — MAP5, LTCI

UNIVERSITÉ PARIS DESCARTES

TELECOM ParisTech

Context: patch-based methods with priors

1. patches extraction

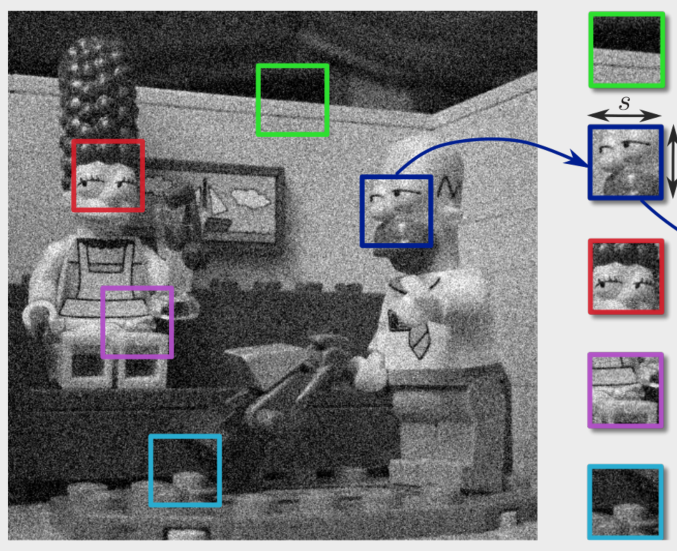

2. modeling in patch-space

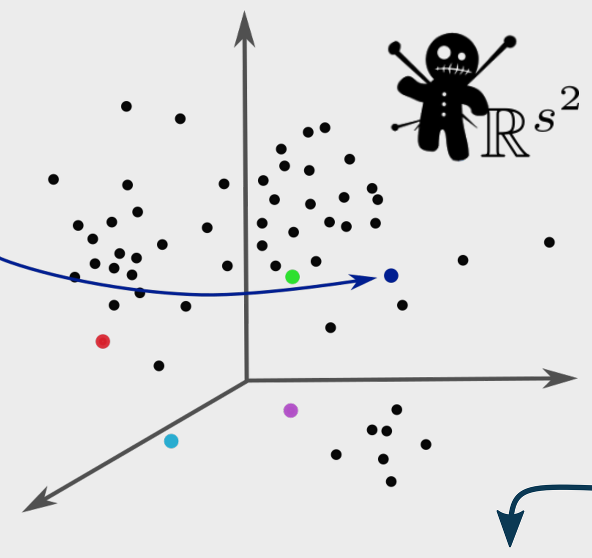

3. prior choice

Bayesian paradigm: put a prior on the underlying clean patches

$$Y_i = D_i X_i + N_i$$

observed data — degradation operator — noise

4. parameters inference

maximum likelihood estimation (MLE)

$$\mathcal{L}(y;\theta) = -\sum_{i=1}^{n} \log\left(\phi(y_i;\theta)\right)$$

issues due to the $\mathbb{R}^{s^2}$ curse of dimensionality

5. clean patches estimation

- MAP
- MMSE (conditional expectation)
- Linear MMSE (Wiener)

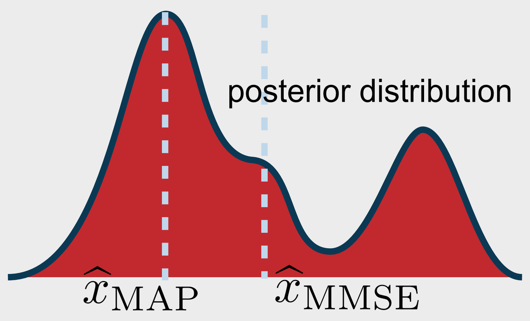

High-Dimensional Mixture Models for Image denoising (HDMI): a Gaussian Mixture Model with dimension regularization

Denoising case: Gaussian white noise model

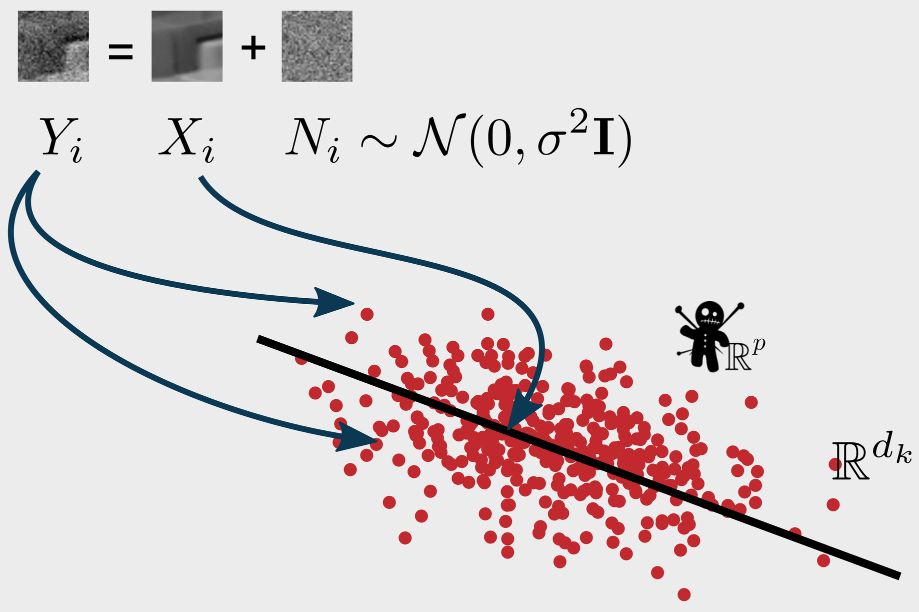

The group memberships are modeled as realizations of a random variable Z and the clean patchs are modeled as

$$\mathrm{X}_{|\mathrm{Z}=k} = U_k \mathrm{T} + \mu_k \text{ avec } \mathrm{T} \in \mathbb{R}^{d_k} \text{ et } \mathrm{T}_{|\mathrm{Z}=k} \sim \mathcal{N}(0, S_k)$$

This implies a full rank GMM on the noisy patches

$$Y \sim \sum_{k=1}^{K} \pi_k \mathcal{N}(y; \mu_k, \Sigma_k) \quad \text{with} \quad Q_k^t \Sigma_k Q_k = \Delta_k = \left(\begin{array}{cc} \begin{array}{ccc} a_{k1} & & 0 \\ & \ddots & \\ 0 & & a_{kd} \end{array} & 0 \\ 0 & \begin{array}{ccc} \sigma^2 & & 0 \\ & \ddots & \\ 0 & & \sigma^2 \end{array} \end{array}\right) \begin{array}{l} \}\, d_k \\ \}\, (p-d_k) \end{array}$$

Inference with EM algorithm

maximize w.r.t. θ the conditional expectation of the complete log-likelihood:

$$\Psi(\theta, \theta^*) \stackrel{\text{def}}{=} \sum_{k=1}^{K} \sum_{i=1}^{n} t_{ik} \log\left(\pi_k g\left(y_i; \theta_k\right)\right),$$

where $t_{ik} = E\left[z = k | y_i, \theta^*\right]$ and θ^* a given set of parameters

E-step: update the t_{ik} knowing current parameters

M-step: update the parameters with MLE

$$\widehat{\pi}_k = \frac{n_k}{n}, \quad \widehat{\mu}_k = \frac{1}{n_k} \sum_i t_{ik} y_i,$$

$$\widehat{S}_k = \frac{1}{n_k} \sum_i t_{ik}(y_i - \mu_k)(y_i - \mu_k)^T,$$

Intrinsic dimensions selection

- when the noise variance σ is known, each group dimension is estimated as

$$\hat{d}_k = \mathrm{argmin}_d \left| \frac{1}{p-d} \sum_{j=d+1}^{p} \lambda_{kj} - \sigma^2 \right|$$

where λ_{kj} are the eigenvalues of the sample covariance $\widehat{S}_k$ from the M-step

- when the noise variance is not known we use the previous approach with various values of σ and we select the best model with BIC criterion

$$\mathrm{BIC}(\mathcal{M}) = \ell(\hat{\theta}) - \frac{\xi(\mathcal{M})}{2} \log(n)$$

where $\xi(\mathcal{M})$ is the complexity of model $\mathcal{M}$ and ℓ the log-likelihood of the parameters

Denoising results

Each patch is denoised by taking the minimum mean square error (MMSE)

Noisy σ = 50 — Denoised with HDMI

HDMI slightly outperform state-of-the-art denoising methods. See preprint : up5.fr/HDMI for PSNR tables

Inpainting

70% missing pixels — recovered image

Since the EM algorithm is designed for data with missing values, the HDMI model can also be used for image interpolation.

Conclusion

HDMI provides a denoising method based on a statistical model with no hack

- demonstrates state-of-the-art performance
- few parameters ton adjust
- can be used in a "blind" way (σ unknown)
- can be adapted for other inverse problems

- high computation cost
- some visual artifacts at high noise levels
- slight low-frequency noise in flat areas

Contact

webpage:
houdard.wp.imt.fr

mail:
antoine.houdard@telecom-paristech.fr

preprint :
up5.fr/HDMI

additional experiments:
houdard.wp.imt.fr/HDMI