Template matching with noisy patches: a contrast-invariant GLR test

Charles Deledalle[1], Loïc Denis[2] and Florence Tupin[3]

[1]IMB, CNRS-Université Bordeaux 1 — [2]Laboratoire Hubert Curien, CNRS-Univ. Saint-Etienne — [3]Institut Mines-Télécom, Télécom ParisTech, CNRS-LTCI

Introduction

Consider

- $\boldsymbol{\theta}$ a noise-free patch of N pixel values,
- $\boldsymbol{x}$ a noisy patch realization of $\boldsymbol{X} \sim p(\cdot|\boldsymbol{\theta})$ where the pdf/pmf p is known,
- $\boldsymbol{a} \in \mathcal{D}$ a template taken from a dictionary $\mathcal{D}$ of noise-free patches (with N pixels),

$\mathcal{D}$ as small as possible $\Rightarrow$ represents classes of patches identical up to a radiometric transformation

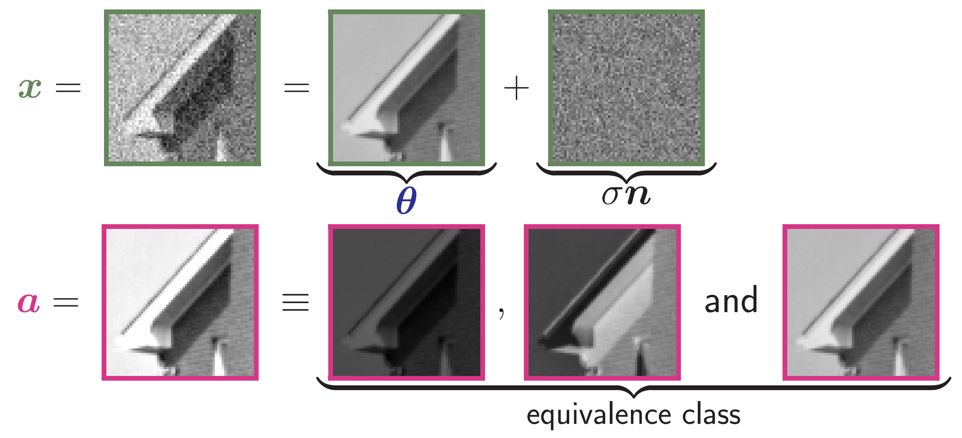

How to match noisy patches $\boldsymbol{x}$ with templates $\boldsymbol{a}$?

Problem statement

Contrast invariance

- Define a template matching criterion $c : (\boldsymbol{x}, \boldsymbol{a}) \mapsto c(\boldsymbol{x}, \boldsymbol{a}) > 0$,
- The larger $c(\boldsymbol{x}, \boldsymbol{a})$, the more relevant the match between $\boldsymbol{x}$ and $\boldsymbol{a}$,
- Consider invariance up to a family of transformation $\mathcal{T}_\rho$ parametrized by ρ

$$\forall \boldsymbol{X}, \boldsymbol{a}, \rho, \quad c(\boldsymbol{X}, \mathcal{T}_\rho(\boldsymbol{a})) = c(\boldsymbol{X}, \boldsymbol{a}) .$$

- Example (affine contrast change): $\mathcal{T}_\rho(\boldsymbol{a}) = \mathcal{T}_{\alpha,\beta}(\boldsymbol{a}) = \alpha \boldsymbol{a} + \beta \mathbb{1}$, where $\mathbb{1}_k = 1$ for all $1 \leq k \leq N$.

Robustness to noise

- Template matching formulation:

$$\exists \rho \quad \boldsymbol{\theta} = \mathcal{T}_\rho(\boldsymbol{a}) \quad \Rightarrow \quad \text{“} \boldsymbol{x} \text{ matches with } \boldsymbol{a} \text{”}$$

- Hypotheses test (parameter test):

$$\mathcal{H}_0 : \exists \rho \quad \boldsymbol{\theta} = \mathcal{T}_\rho(\boldsymbol{a}) \quad \text{(null hypothesis)},$$
$$\mathcal{H}_1 : \forall \rho \quad \boldsymbol{\theta} \neq \mathcal{T}_\rho(\boldsymbol{a}) \quad \text{(alternative hypothesis)}.$$

- Optimal Neyman-Pearson criterion (maximizing P_D for any given P_{FA}) is the likelihood ratio test:

$$\mathcal{L}(\boldsymbol{x}, \boldsymbol{a}) = \frac{p(\boldsymbol{x}|\boldsymbol{\theta} = \mathcal{T}_\rho(\boldsymbol{a}), \mathcal{H}_0)}{p(\boldsymbol{x}|\boldsymbol{\theta}, \mathcal{H}_1)}.$$

$\Rightarrow$ cannot be evaluated since ρ and $\boldsymbol{\theta}$ are unknown.

How to define a suitable template matching criterion?

Typical contrast-invariant template matching

Normalized correlation:

- Most usual way to measure similarity up to an affine change of contrast:

$$\mathcal{C}(\boldsymbol{x}, \boldsymbol{a}) = \left| \frac{\sum_k (x_k - \bar{x})(a_k - \bar{a})}{\sqrt{\sum_k (x_k - \bar{x})^2 \sum_k (a_k - \bar{a})^2}} \right|$$

where $\bar{x} = \frac{1}{N} \sum x_k$ and $\bar{a} = \frac{1}{N} \sum a_k$.

- such that

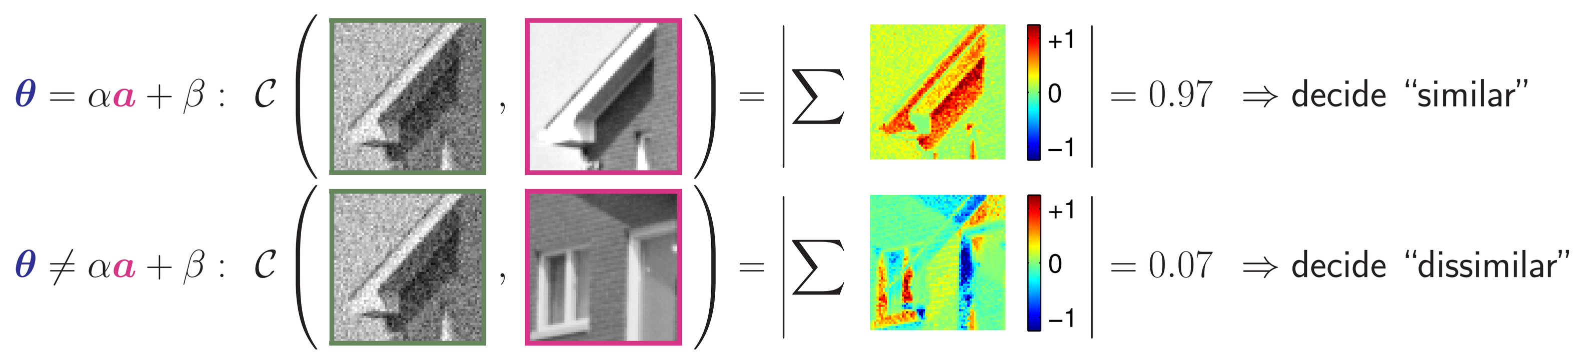

Is correlation a robust template matching criterion wrt different noise statistics?

Contrast-invariant template matching with noisy patches

Using the Generalized Likelihood Ratio (GLR):

- Replaces the unknowns ρ and $\boldsymbol{\theta}$ by their maximum likelihood estimates (MLE) under each hypothesis:

$$\mathcal{G}(\boldsymbol{x}, \boldsymbol{a}) = \frac{\sup_\rho p(\boldsymbol{x}|\boldsymbol{\theta} = \mathcal{T}_\rho(\boldsymbol{a}), \mathcal{H}_0)}{\sup_{\boldsymbol{t}} p(\boldsymbol{x}|\boldsymbol{\theta} = \boldsymbol{t}, \mathcal{H}_1)} = \frac{p(\boldsymbol{x}|\boldsymbol{\theta} = \mathcal{T}_{\hat{\rho}}(\boldsymbol{a}))}{p(\boldsymbol{x}|\boldsymbol{\theta} = \hat{\boldsymbol{t}})}$$

where $\hat{\rho}$ and $\hat{\boldsymbol{t}}$ are the MLE of the unknown ρ and $\boldsymbol{\theta}$,

- Satisfies the contrast invariance property by construction,
- Asymptotically (wrt the SNR) optimal with constant false alarm rate (CFAR).
- Invariant upon changes of variable [Kay and Gabriel, 2003],
- May fail in low SNR conditions, where the MLE is known to behave poorly.

Using variance stabilization:

- Transform patches such that the noise component be approximately Gaussian (with constant variance),
- Example: homomorphic transform for multiplicative noise; Anscombe transform for Poisson noise.
- Given an application s which stabilizes the variance for a specific noise distribution, stabilization-based criteria can be obtained on the output of s as:

$$\mathcal{S}_\mathcal{C}(\boldsymbol{x}, \boldsymbol{a}) = \mathcal{C}(s(\boldsymbol{x}), s(\boldsymbol{a})) \quad \text{and} \quad \mathcal{S}_\mathcal{G}(\boldsymbol{x}, \boldsymbol{a}) = \mathcal{G}(s(\boldsymbol{x}), s(\boldsymbol{a}))$$

where it is assumed that $s(\boldsymbol{X}) \sim \mathcal{N}(s(\boldsymbol{\theta}), \sigma^2 \mathbf{I})$.

- Usually simpler to evaluate in closed-form, and then, leads to faster algorithms.
- Limited to the existence of a stabilization function s.

Proposition (GLR for Gaussian noise)

Consider that $\boldsymbol{X}$ follows an uncorrelated Gaussian distribution:

$$p(x_k|\theta_k) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x_k - \theta_k)^2}{2\sigma^2} \right) ,$$

and consider the class of affine contrast transformations $\mathcal{T}_{\alpha,\beta}(\boldsymbol{x}) = \alpha \boldsymbol{x} + \beta \mathbb{1}$. In this case, we have

$$-\log \mathcal{G}(\boldsymbol{x}, \boldsymbol{a}) = \left(1 - \mathcal{C}(\boldsymbol{x}, \boldsymbol{a})^2\right) \frac{\|\boldsymbol{x} - \bar{x}\mathbb{1}\|_2^2}{2\sigma^2} .$$

Correlation vs GLR under Gaussian noise

- Correlation does not take into account the noise while GLR does, ex.:

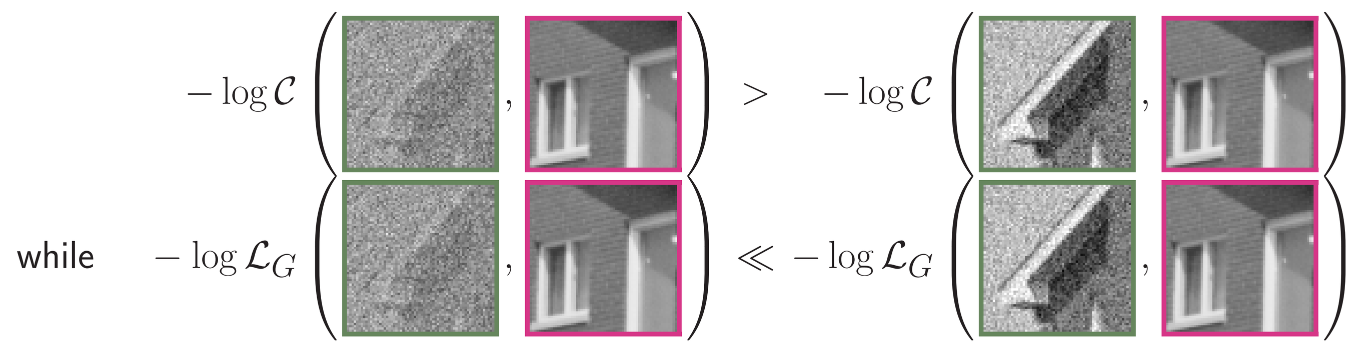

- In fact $0 \times$  $+ \beta$ “explains” better 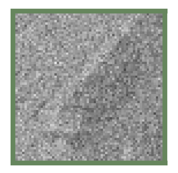 than 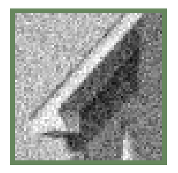.

Proposition (GLR for Gamma noise)

Consider that $\boldsymbol{X}$ follows a gamma distribution such that

$$p(x_k|\theta_k) = \frac{L^L x_k^{L-1}}{\Gamma(L)\theta_k^L} \exp\left(-\frac{L x_k}{\theta_k} \right)$$

and consider the class of log-affine transformations $\mathcal{T}_{\alpha,\beta}(\boldsymbol{x}) = \beta \boldsymbol{x}^\alpha$ where $(.)^\alpha$ is the element-wise power function. In this case, we have

$$-\log \mathcal{G}(\boldsymbol{x}, \boldsymbol{a}) = L \sum_{k=1}^N \log \left(\frac{\hat{\beta} a_k^{\hat{\alpha}}}{x_k} \right)$$

where $\hat{\alpha}$ and $\hat{\beta}$ can be obtained iteratively as

$$\hat{\alpha}_{i+1} = \hat{\alpha}_i - \frac{\sum_k (1 - r_{k,i}) \log a_k}{\sum_k r_{k,i} (\log a_k)^2} \quad \text{and} \quad \hat{\beta}_{i+1} = \frac{1}{N} \sum_k \frac{x_k}{a_k^{\hat{\alpha}_i}}$$

with $r_{k,i} = x_k / (\hat{\beta}_i a_k^{\hat{\alpha}_i})$, whatever the initialization.

Similar result for Poisson noise.

Detection performance

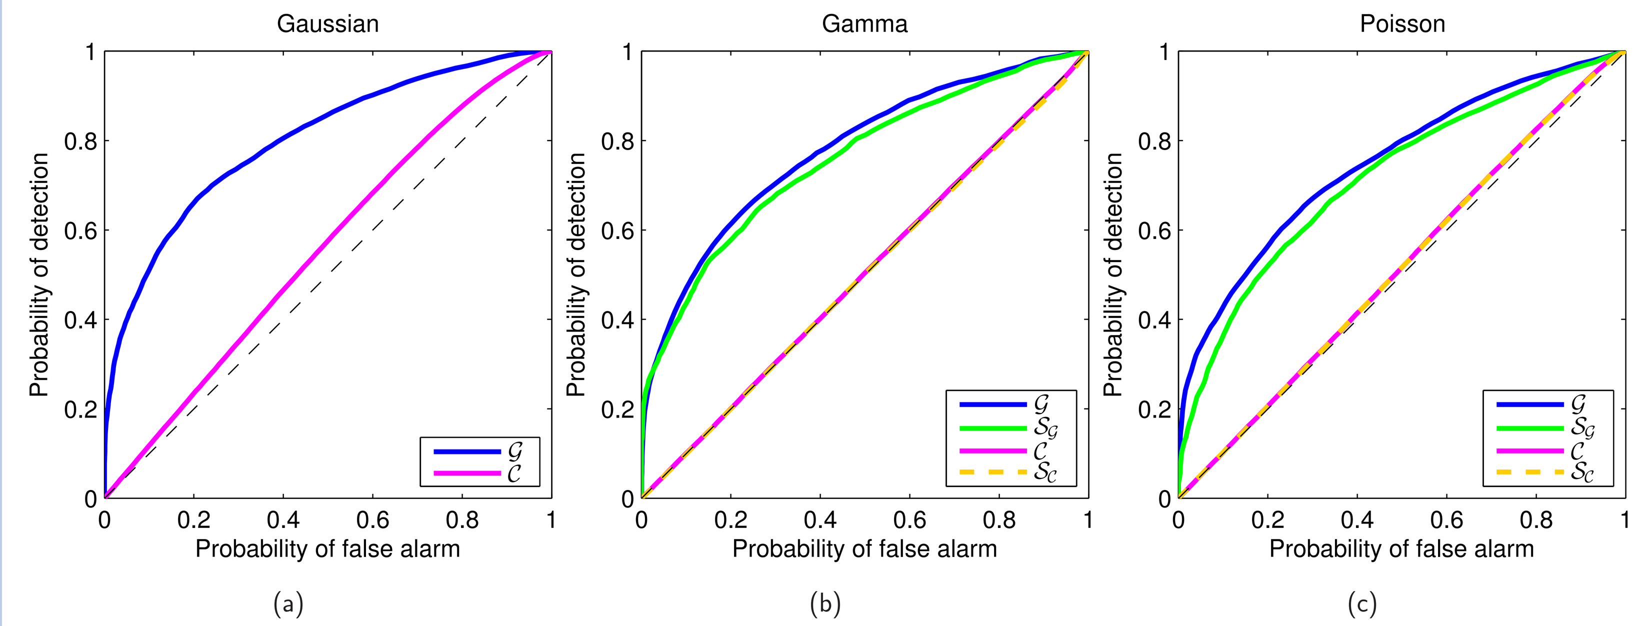

Figure: (a) ROC curve obtained under Gaussian noise, (b) ROC curve obtained under gamma noise and (c) ROC curve obtained under Poisson noise. In all experiments, the SNR is about -3dB.

- Dictionary of 196 templates of size $N = 8 \times 8$ (extracted from the image *“Barbara”* with k-means),
- Several noisy realizations $\boldsymbol{x}$ are generated for several radiometric transformations $\theta = \mathcal{T}_\rho(\boldsymbol{a})$,
- GLR provides the best performance followed by Gaussian GLR after variance stabilization,
- Correlation acts poorly in all situations.

Application to dictionary-based denoising

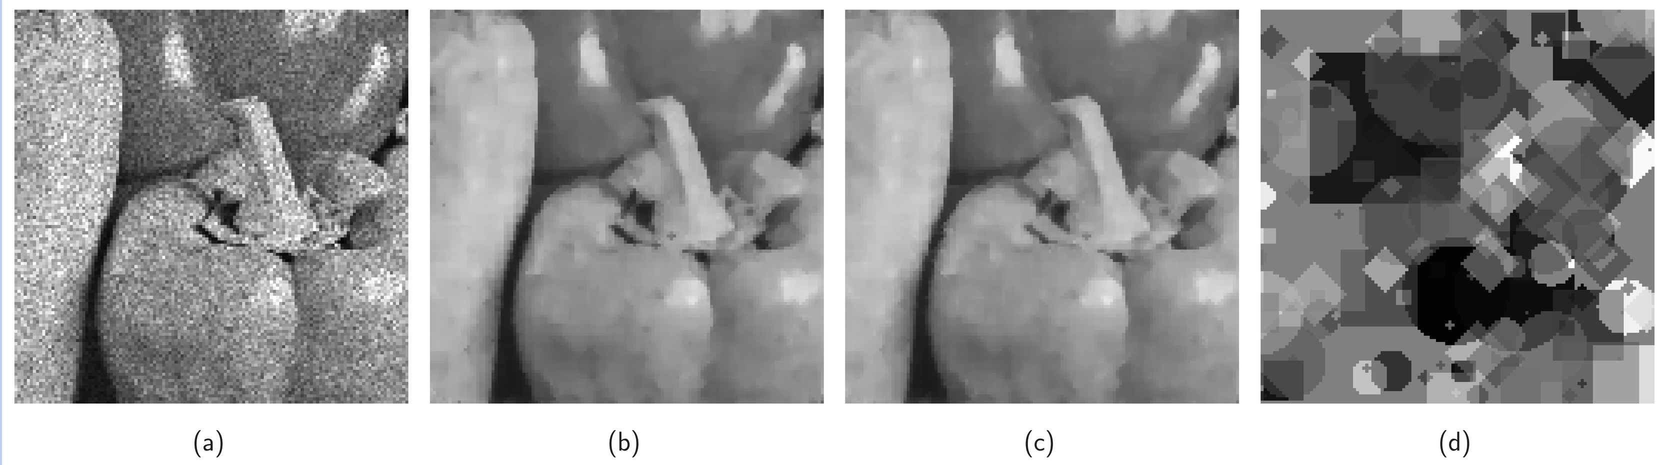

(a) (b) (c) (d)

Figure: (a) Noisy input image damaged by gamma noise (with $L = 10$, $PSNR = 21.14$). (b) Denoised image using the GLR after variance stabilization followed by a debiasing step following [Xie et al., 2002] ($PSNR = 27.42$). (c) Denoised image using the GLR adapted to gamma noise ($PSNR = 27.53$). (d) Image composed of the atoms of the dictionary.

Template-matching based denoising:

- The dictionary $\mathcal{D}$ provides a generative model of the patches $\boldsymbol{x}$ of the noisy image,
- Each patch of the image can then be estimated as:

$$\hat{\boldsymbol{\theta}}(\boldsymbol{x}) = \frac{1}{Z} \sum_{\boldsymbol{a} \in \mathcal{D}} \mathcal{G}(\boldsymbol{x}, \boldsymbol{a}) \boldsymbol{a}^\star \quad \text{with} \quad Z = \sum_{\boldsymbol{a} \in \mathcal{D}} \mathcal{G}(\boldsymbol{x}, \boldsymbol{a}) ,$$

where $\boldsymbol{a}^\star = \mathcal{T}_{\hat{\rho}}(\boldsymbol{a})$ and $\hat{\rho}$ is the MLE of ρ used in the calculation of $\mathcal{G}(\boldsymbol{x}, \boldsymbol{a})$.

“Multi-scale” shift-invariant dictionary:

- $\mathcal{D}$ is composed of the set of all atoms extracted from a 128×128 image (a.k.a., an epitome) built following the transparent dead leaves model of [Galerne and Gousseau, 2012],
- The dictionary is then shift invariant and denoising can be performed in the Fourier domain (see [Jost et al., 2006, Benoît et al., 2011]) while representing information of different scales,
- GLR for the gamma law or for the Gaussian law after stabilizing the variance are both satisfactory visually and in term of PSNR.

References

Benoît, L., Mairal, J., Bach, F., and Ponce, J. (2011).
Sparse image representation with epitomes.
In *Proceedings of CVPR, 2011*. IEEE.

Galerne, B. and Gousseau, Y. (2012).
The transparent dead leaves model.
Advances in Applied Probability, 44(1):1–20.

Jost, P., Vandergheynst, P., Lesage, S., and Gribonval, R. (2006).
Motif: an efficient algorithm for learning translation invariant dictionaries.
In *Proceedings of ICASSP 2006*. IEEE.

Kay, S. and Gabriel, J. (2003).
An invariance property of the generalized likelihood ratio test.
IEEE Signal Processing Letters, 10(12):352–355.

Xie, H., Pierce, L., and Ulaby, F. (2002).
SAR speckle reduction using wavelet denoising and Markov random field modeling.
IEEE Transactions on Geoscience and Remote Sensing, 40(10):2196–2212.

http://www.math.u-bordeaux1.fr/~cdeledal/

charles-alban.deledalle@math.u-bordeaux1.fr