A Practical Guide to Multi-image Alignment

Cecilia Aguerrebere Mauricio Delbracio Alberto Bartesaghi Guillermo Sapiro

Problem

Problem: Multi-image alignment: bring different images into one coordinate system

Challenge:

- Several applications with very different SNR conditions
- Great amount of effort invested into developing alignment methods
- We seek to organize the main multi-image alignment methods under a common framework and provide practical answers to fundamental questions

Fundamental Questions:

- Fundamental limits in multi-image alignment performance?
- Best possible accuracy?
- Does having more images help?
- Does shift prior information help?
- Is there any room for improvement?

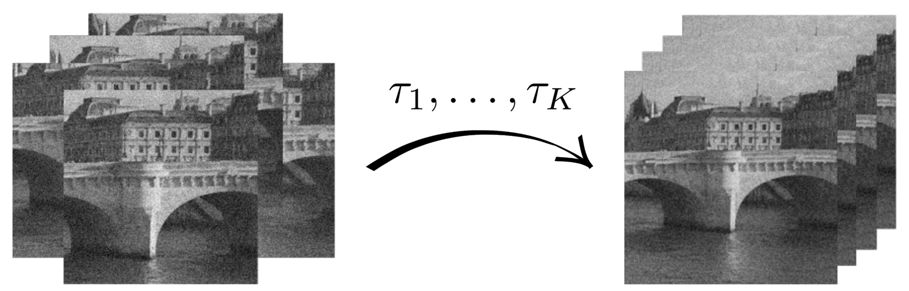

Image Alignment Model

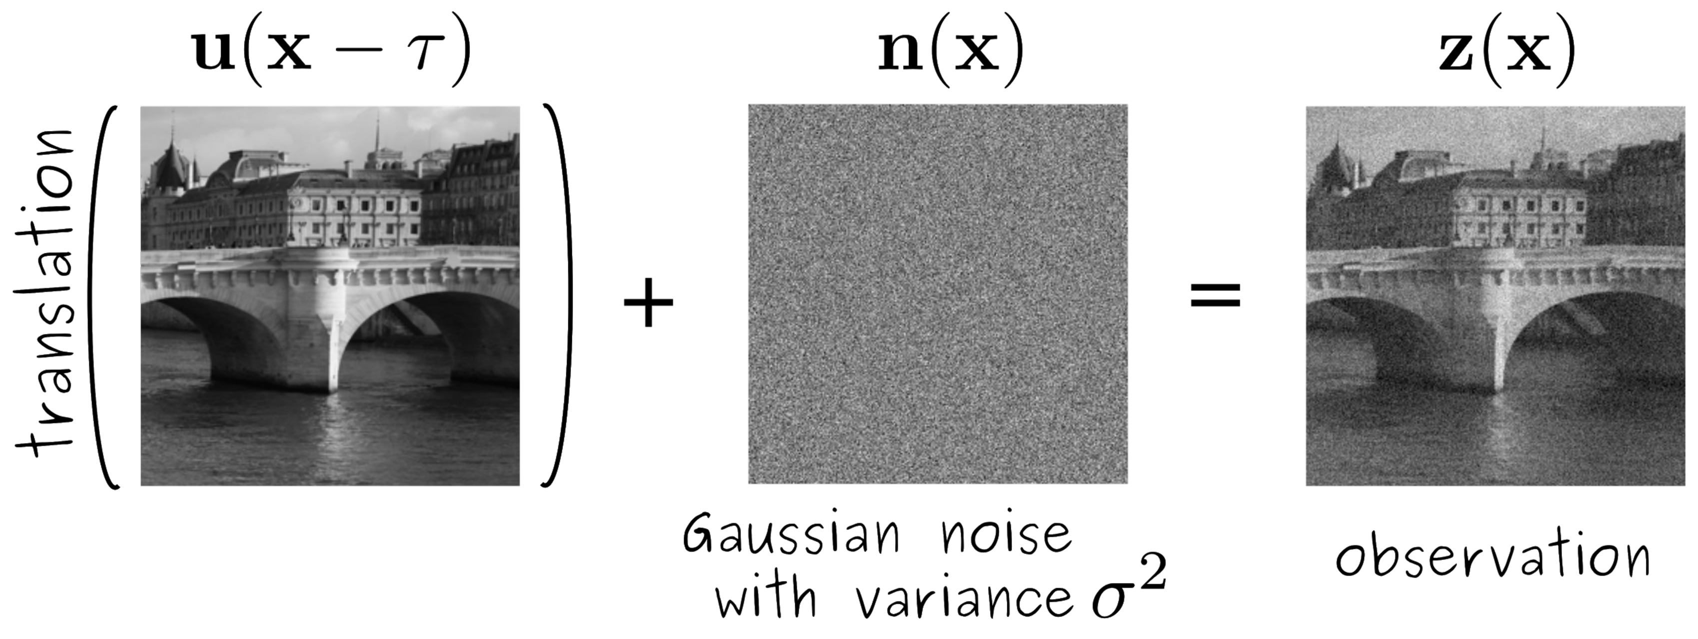

$$z_i(\mathbf{x}) = u(\mathbf{x}-\boldsymbol{\tau}_i) + n_i(\mathbf{x}), \quad i = 0, \ldots, K, \qquad \text{(M)}$$

To consider:

- Number of images **K**
- SNR (image energy / noise level)
- Image model (random, determ.)
- Shift prior

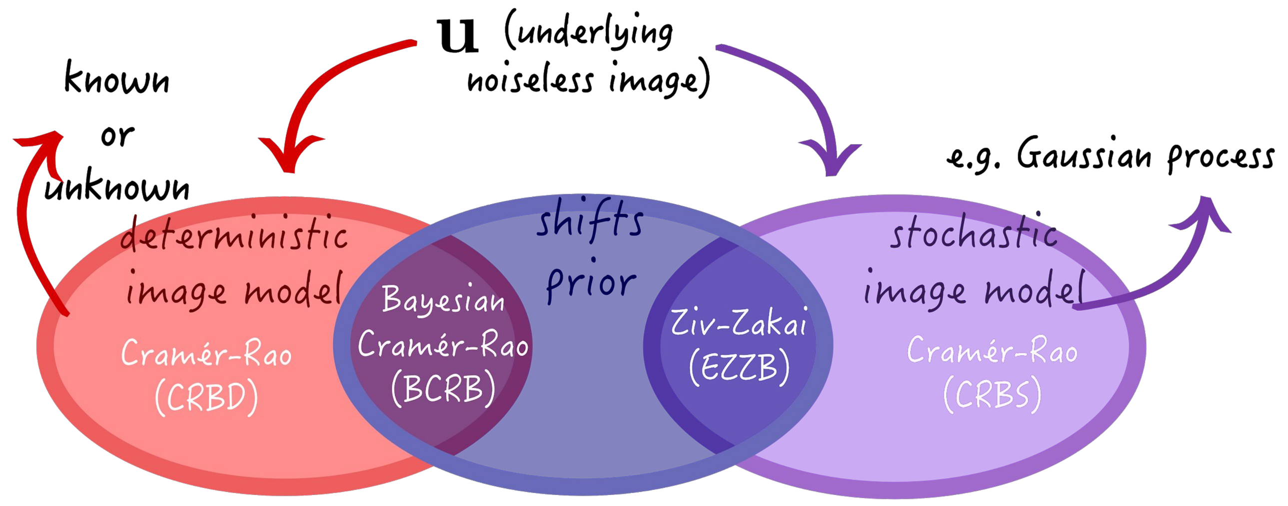

Existing Multi-image Alignment Methods

Common Framework: Many multi-image alignment methods (e.g., MAP, MLE) optimize the "same" cost function, with different strategies.

$$[u,\boldsymbol{\tau}]_{\text{MAP}} = \arg\max_{u,\boldsymbol{\tau}} \left[\log p(\mathbf{z}|u,\boldsymbol{\tau}) + \log p(u)\right]$$
$$= \arg\min_{u,\boldsymbol{\tau}} \frac{1}{2\sigma^2}||\mathbf{z}-\mathbf{B}u||^2 + \frac{1}{2}u^T\boldsymbol{\Sigma}_u^{-1}u$$

Notation:

- $\mathbf{B}(\boldsymbol{\tau}) = [B(\boldsymbol{\tau}_1)^T, \ldots, B(\boldsymbol{\tau}_K)^T]^T$, with $B(\boldsymbol{\tau}_i)$ Shannon shift operator that verifies: $B(\boldsymbol{\tau}_i)u(\mathbf{x}) = u(\mathbf{x}-\boldsymbol{\tau}_i)$.
- $\tilde{\mathbf{B}} = \mathbf{FBF}^T$ is the Fourier equivalent of the shift operator
- u modeled as deterministic unknown (MLE), or as a stationary zero-mean Gaussian process with spectral density $\mathbf{S}_u$

- **Cyclic coordinate descent:** optimize one coordinate at a time: (i) given shifts, estimate u; (ii) estimate new shifts with respect to u. Repeat.
- **Variable Projections**: directly maximize to recover the unknown shifts [2]

$$\boldsymbol{\tau}_{\text{MAP}} = \arg\max_{\boldsymbol{\tau}\in\mathbb{R}^{2K}} \tilde{\mathbf{z}}^T\tilde{\mathbf{B}}\tilde{\mathbf{W}}\tilde{\mathbf{B}}^T\tilde{\mathbf{z}}$$

MAP: $\tilde{\mathbf{W}} = (\mathbf{S}_u(K+1)+\mathbf{I}\sigma^2)^{-1}\mathbf{S}_u$ is the Wiener filter

MLE: $\tilde{\mathbf{W}} = \mathbf{I}$ (no image prior)

- How important is **initialization**? (non-convex)

Theoretical Performance Limits on Multi-alignment Accuracy

Cramer-Rao Bound (CRB)
(stochastic image model)

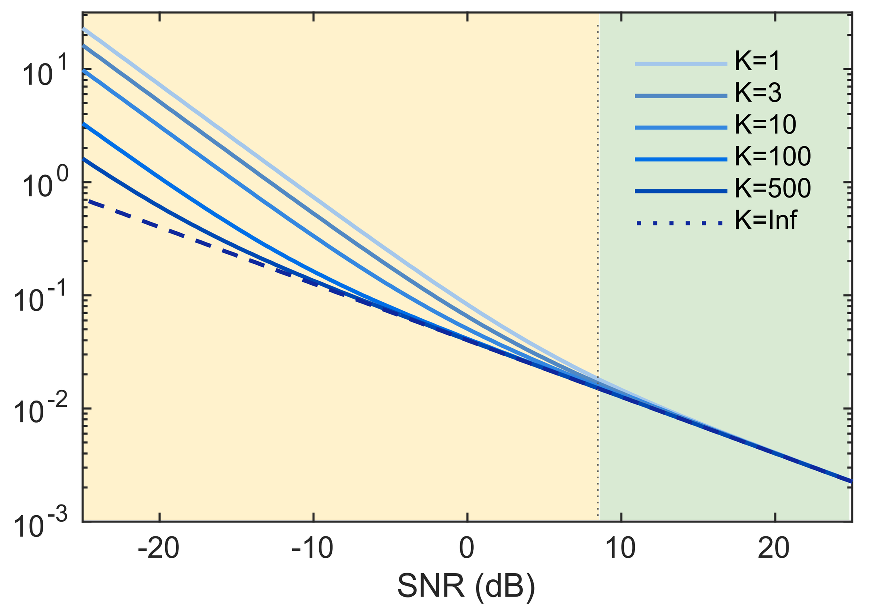

Extended Ziv-Zakai Bound (EZZB)
(white noise image model + uniform shift prior)

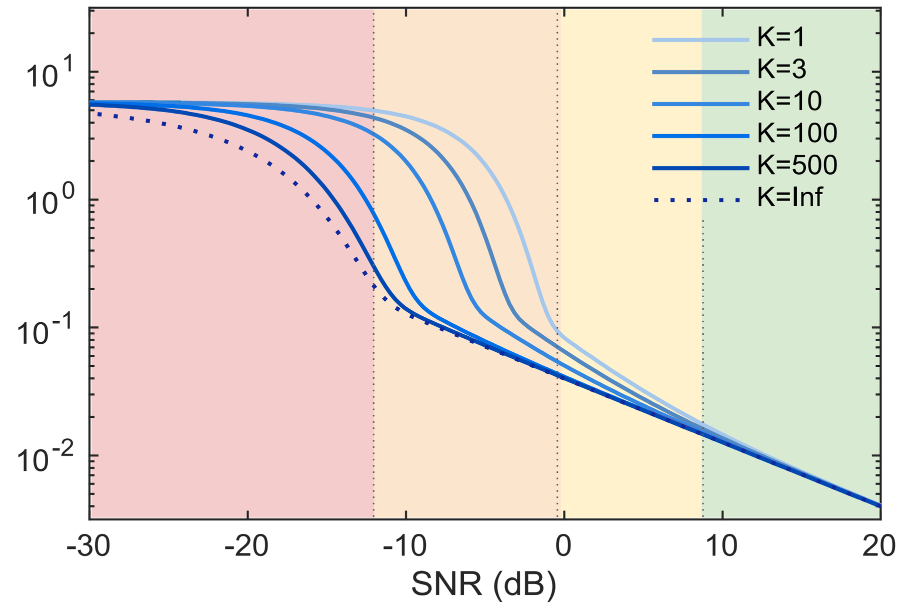

Different image/shift models lead to different performance bounds.
Behaviour depends on image SNR (total energy, noise level) and number of images [3]

Experimental Setup

Simulated Experiments:

- Sets of images generated following (M), different noise levels and number of shifted images
- Motion considered: uniformly distributed independent shifts and drift-driven trajectories
- MAP image prior: zero-mean Gaussian process spectral density decaying with frequency
- Experiments repeated 100 times for each SNR level; RMSE mean and 95% CI reported
- All evaluated methods are almost unbiased (bias orders of magnitude smaller than variance)
- SNR defined as ratio between energy of the derivative and the noise power

Multi-image Denoising: align and average

Evaluate the effect of image prior and patch size on multi-image denoising of real images

Results

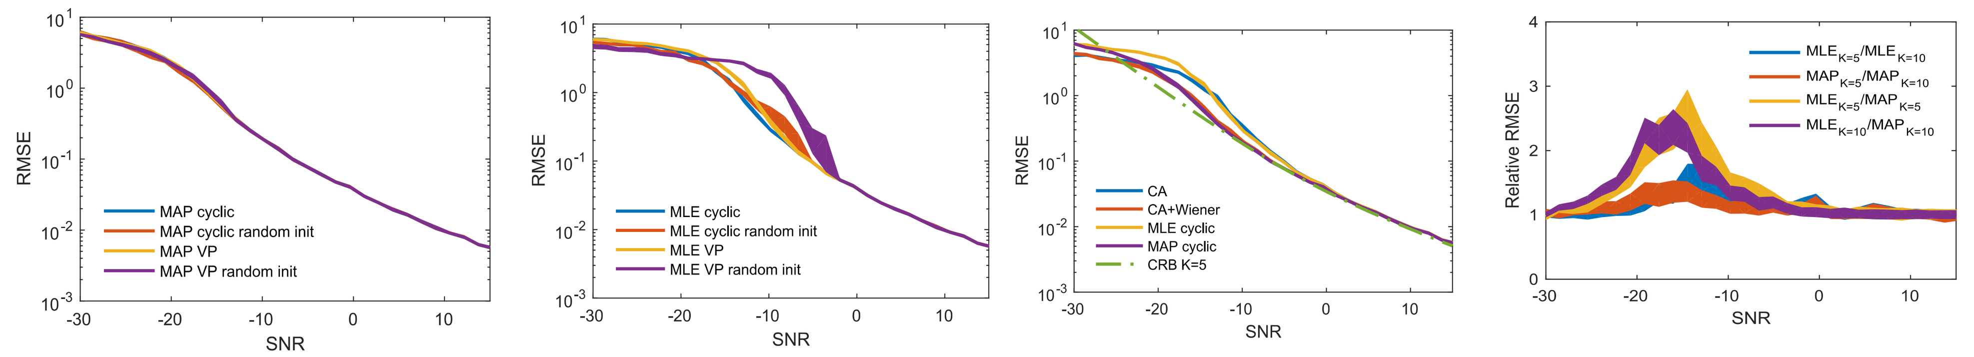

Experiments show three SNR regions high, moderate and very low SNR

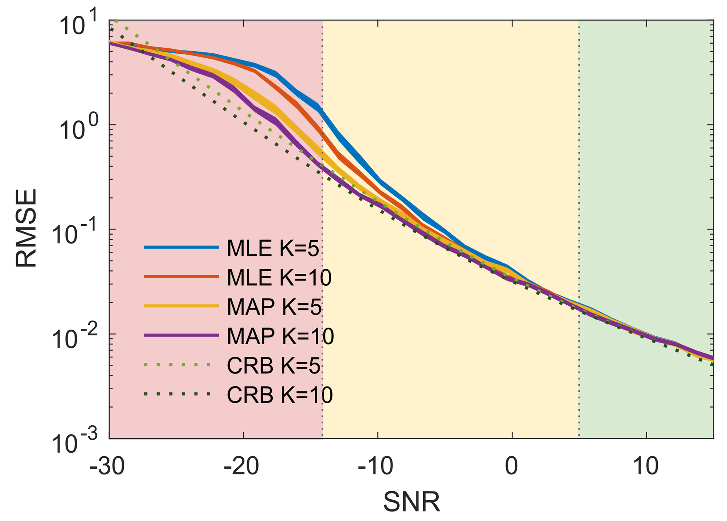

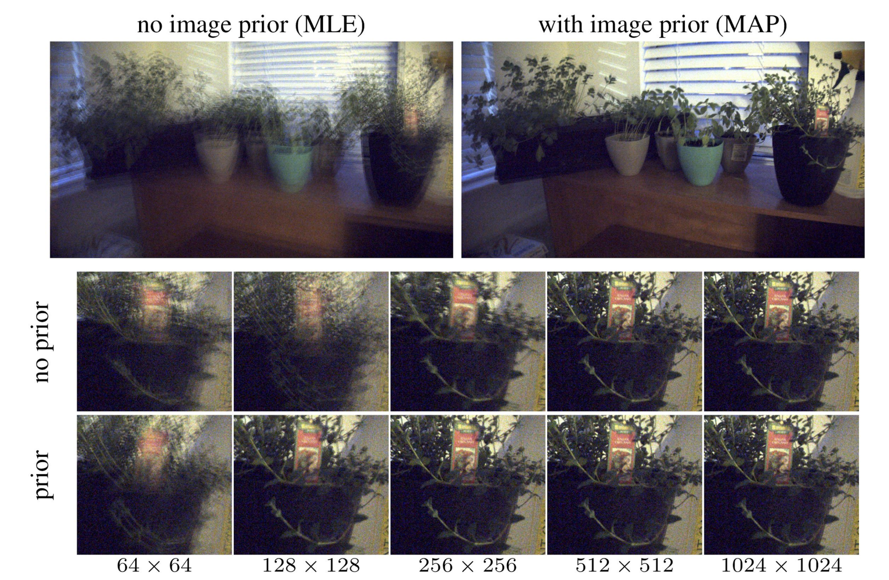

Align and average 5 real images using MLE (no image prior) or MAP (Gaussian image prior) when using different patch-sizes

Concluding Remarks

Per-region behavior depending on the SNR:

High SNR

- All evaluated methods perform very similarly and very close to performance bounds
- Simplest methods already achieve the best possible performance
- Including image prior image or more images does not improve alignment accuracy

Moderate to low SNR

- Clear performance gain when including image prior or using more images. Twofold gain: MSE is reduced; threshold at which performance degrades dramatically is pushed back several dBs
- Performance gain obtained with image prior is larger than that of increasing number of images
- Methods with image prior perform very close to CRB (little room for improvement)
- Optimization/initialization: slight differences observed for MLE in low SNR, not a critical point

Very low SNR

- limit SNR value below alignment is not possible; neither more images nor image priors help (as predicted by [2]). Only way out: increase SNR, e.g., increasing image size (patch size in case of local alignment).

References

[1] C. Aguerrebere et al. A Practical Guide to Multi-image Alignment. In ICASSP 2018.
[2] D. Robinson et al. Optimal registration of aliased images using variable projection with applications to super-resolution. *The Computer Journal,* 52.1 (2007)
[3] C. Aguerrebere et al. Fundamental limits in multi-image alignment. In *IEEE Trans. Signal Proc.*, 64.21 (2016).