Bayesian Average SParsity

Bayesian Average SParsity

Valentina Masarotto*, Yves Wiaux†,
* Chair of Mathematical statistics, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland.
† Institute of Sensors, Signals, and Systems, Heriot-Watt University, Edinburgh EH14 4AS, UK.

***Abstract*—The Average Sparsity regularisation prior was recently leveraged for compressive imaging in the convex optimisation algorithm SARA, and shown to provide superior reconstruction quality in comparison to simple sparsity [1]. We discuss an approach for introducing Average Sparsity in a Bayesian algorithm for compressive imaging.**

Compressed sensing (CS) enables the accurate recovery of sparse signals from well-designed sub-Nyquist sampling. The main stream signal reconstruction approaches rely on convex optimisation, in particular on the Basis Pursuit (BP) problem searching for the signal with minimum ℓ_1-norm representation in the sparsity basis, simultaneously satisfying the data constraint. Among many variations of BP, the Average Sparsity regularisation prior was recently leveraged for compressive imaging in the convex optimisation algorithm SARA, and shown to provide superior reconstruction quality in comparison to simple sparsity [1]. Average sparsity exploits the hypothesis that a natural image is simultaneously sparse in multiple bases, thus providing a more powerful structured sparsity prior. Bayesian reconstruction algorithms have also been proposed. In this case, prior information will be encapsulated into a prior distribution function. Bayesian methods offer several advantages. Here, posterior evaluation will yield the full distribution of the signal rather than just a point estimate. Importantly, in a parametric Bayesian approach, prior parameters can be attached to the signal parameter space and estimated simultaneously. This yields a completely automatic approach that avoids fine-tuning of the algorithms, as is the case in convex optimisation (and greedy) approaches. To our knowledge, no Bayesian approach has made use of average sparsity so far.

We discuss an approach for introducing Average Sparsity in a Bayesian algorithm for compressive imaging. The key question we aim to answer is the following. Can we promote a sparsity prior that accounts for the signal being simultaneously sparse in several bases? The method we present originates from the Bayesian Compressed sensing with Laplace priors introduced in [2] and [3] and incorporates a sparsity prior in line with Average Sparsity. The basic setting in [2] and [3] to recover a signal x from the data $\boldsymbol{y} = \Phi\boldsymbol{x} + \boldsymbol{n}$, assuming sparsity of the representation $\boldsymbol{w}$ with $\boldsymbol{x} = \Psi\boldsymbol{w}$, relies on assuming Gaussian noise and a Laplace distribution for $\boldsymbol{w}$:

$$\begin{aligned} n_j &\sim \text{i.i.d.}\, \mathcal{N}(0, \beta^{-1}); \\ w_i &\sim \text{i.i.d.}\, \text{Laplace}(\lambda). \end{aligned}$$

As for the noise variance β^{-1}, one can either model it through a hyperprior distribution, or take it as fixed if we assume we know the noise variance. In our simulations below we consider the latter option. In order to exploit conjugacy and have analytical solutions, the Laplace prior is replaced with the following Normal-Gamma hierarchical model:

$$\begin{aligned} w_i|\gamma_i &\sim \mathcal{N}(0, \gamma_i); \\ \gamma_i|\lambda &\sim \Gamma(1, \lambda/2); \\ \lambda|\nu &\sim \Gamma(\nu/2, \nu/2). \end{aligned}$$

The hyperparameter $\lambda/2$ identifies the inverse of the mean of $\Gamma(1, \lambda/2)$. The hyperparameter ν is assumed to be fixed. The sparsity of the wavelet representation $\boldsymbol{w}$ is captured into the variance vector $\boldsymbol{\gamma}$. Small variances will lead to negligible coefficients.

We hypothesise that the image of interest is simultaneously sparse in bases $(\Psi^1, \cdots, \Psi^s)$. Average sparsity here really means giving meaning to averaging the variances. In order to implement it, we propose an empirical Bayes two-stage model. As a first step, we estimate the variance vectors $(\boldsymbol{\gamma}^1, \cdots, \boldsymbol{\gamma}^s)$ separately, by using the maximum likelihood estimation procedure proposed in [3] assuming sparsity in one of the sparsity basis at the time. The variance vectors are then all imported in an arbitrarily chosen reference basis through the adequate transformation on covariance matrices. A this stage, s estimates are available for each γ_i in the reference sparsity basis. Assuming that those γ_i's are samples of the prior distribution Γ on variances, their sample mean over the s available values at each point i provides the maximum likelihood estimation of the mean of the Γ, which we take as our final estimate for γ_i. That is Average Sparsity. As a second step, the final estimate for the variances in the reference basis is then used to estimate the wavelet coefficients as the expected values of a Gaussian posterior as prescribed by [3].

We provide below results of preliminary simulations on a $N = 64 \times 64$ image of Lena. We consider 0.5N measurements defined as projections of $\boldsymbol{x}$ on random i.i.d. Gaussian vectors, affected by a fixed input SNR of 30dB. For BASP, we assume average sparsity in the Db1 to Db8 wavelet bases, with the chosen reference sparsity basis being Db5. BASP gives an output SNR = 25.71 Db. The best reconstruction using the method of [3] with a single basis chosen within Db1 - Db8 is provided with Db5, which leads to an output SNR = 22.69 Db. The figure shows the original image (left panel), the BASP reconstruction (middle panel), and the single basis reconstruction (right panel). These preliminary results suggest that BASP significantly outperforms the original Bayesian CS approach of [3]. Future work should provide detailed comparison between these Bayesian methods as well as with SARA.

References

[1] R. Carrillo, J. McEwen, and Y. Wiaux, "Sparsity averaging reweighted analysis (sara): a novel algorithm for radio-interferometric imaging," *Monthly Notices of the Royal Astronomical Society*, vol. 426, no. 2, pp. 1223–1234, 2012.

[2] S. Ji, Y. Xue, and L. Carin, "Bayesian compressive sensing," *Signal Processing, IEEE Transactions on*, vol. 56, no. 6, pp. 2346–2356, 2008.

[3] S. D. Babacan, R. Molina, and A. K. Katsaggelos, "Bayesian compressive sensing using laplace priors," *Image Processing, IEEE Transactions on*, vol. 19, no. 1, pp. 53–63, 2010.