

Complex Nonconvex l_p Norm Minimization for Underdetermined Source Separation

Emmanuel Vincent

METISS Group, IRISA-INRIA
Campus de Beaulieu, 35042 Rennes Cedex, France
`emmanuel.vincent@irisa.fr`

Abstract. Underdetermined source separation methods often rely on the assumption that the time-frequency source coefficients are independent and Laplacian distributed. In this article, we extend these methods by assuming that these coefficients follow a generalized Gaussian prior with shape parameter p . We study mathematical and experimental properties of the resulting complex nonconvex l_p norm optimization problem in a particular case and derive an efficient global optimization algorithm. We show that the best separation performance for three-source stereo convolutive speech mixtures is achieved for small p .

1 Introduction

Underdetermined source separation is the problem of recovering the single-channel source signals $s_j(t)$, $1 \leq j \leq J$, underlying a multichannel mixture signal $x_i(t)$, $1 \leq i \leq I$, with $I < J$. The mixing process can be modeled in the time-frequency domain via the Short-Term Fourier Transform (STFT) as [1]

$$\mathbf{X}(n, f) = \mathbf{A}(f)\mathbf{S}(n, f) \quad (1)$$

where $\mathbf{S}(n, f)$ is the vector of source STFT coefficients in time-frequency bin (n, f) , $\mathbf{X}(n, f)$ is the vector of mixture STFT coefficients in the same bin, and $\mathbf{A}(f)$ is a complex mixing matrix. This problem can be tackled by first estimating the mixing matrices and then deriving the Maximum *A Posteriori* (MAP) source coefficients under the constraint (1), based on some prior distribution.

Existing separation methods rely on the assumption that the source coefficients are independent and sparsely distributed, *i.e.* a large proportion of coefficients are close to zero. Examples of sparse priors include mixtures of Dirac impulses and Gaussians [2], mixtures of Gaussians [3], Student t distributions [4] and the Laplacian distribution [5,6,1,7]. The latter is popular since it leads to a convex optimization problem that can be solved efficiently. In this paper, we extend this approach by assuming that the source coefficients follow a generalized Gaussian prior, of which the Laplacian is a special case. This extension is not straightforward, since the resulting criterion can be nonconvex.

The structure of the rest of this paper is as follows. In Section 2, we argue that generalized Gaussian priors are well suited to the modeling of speech signals.

We study the properties of the resulting optimization problem in Section 3 in the case where $J = I + 1$ and derive an efficient global optimization algorithm. We evaluate its performance on instantaneous and convolutive speech mixtures in Section 4 and conclude in Section 5.

2 Generalized Gaussian priors

Generalized Gaussian priors were first introduced in the context of source separation via Independent Component Analysis (ICA) [8,9,10]. The phases of the source STFT coefficients $S_j(n, f)$ are assumed to be uniformly distributed, while their magnitudes are modeled by

$$P(|S_j(n, f)|) = p \frac{\beta^{1/p}}{\Gamma(1/p)} e^{-\beta |S_j(n, f)|^p} \quad (2)$$

where the parameters $p > 0$ and $\beta > 0$ govern respectively the shape and the variance of the prior. This prior includes the Laplacian ($p = 1$) and the Gaussian ($p = 2$) as special cases and its sparsity increases with decreasing p .

In order to assess the benefit of using this prior, we computed the best shape parameters p for 30 speech signals, considering all frequency bins either separately or together. The signals were sampled at 8 kHz and had a duration of 12 s. The STFT was computed using half-overlapping sine windows of various lengths L and each frequency bin was scaled to unit variance. The Maximum Likelihood (ML) parameters were estimated using Matlab `fminunc` optimizer¹.

The observed parameter range is depicted in Figure 1. On average, p varies between 0.4 and 0.9 depending on the window length L and stays almost constant across frequency, except at very low frequencies where background noise dominates. This shows that generalized Gaussian priors with $p < 1$ better fit speech sources than Laplacian priors. Interestingly, the observed value of p reaches a minimum for $L = 64$ ms, which was also previously determined to be the optimal window length for source separation via binary masking [11,12].

3 Properties of the complex l_p norm criterion

Given these results, we now assume that the mixing matrices $\mathbf{A}(f)$ are known and that the source STFT coefficients follow a generalized Gaussian prior with fixed parameters p and β . The MAP source coefficients are given by

$$\hat{\mathbf{S}}(n, f) = \arg \min_{\mathbf{S} \in \mathbb{C}^J} \|\mathbf{S}\|_p^p \text{ subject to } \mathbf{A}(f)\mathbf{S} = \mathbf{X}(n, f) \quad (3)$$

where $\|\mathbf{S}\|_p$ is the l_p norm of the vector \mathbf{S} defined by $\|\mathbf{S}\|_p^p = \sum_{j=1}^J |S_j|^p$. When $p < 1$, this criterion is nonconvex hence difficult to minimize.

¹ This algorithm is based on a subspace trust region method. For more details, see http://www.mathworks.com/access/helpdesk_r13/help/toolbox/optim/fminunc.html




Fig. 1. Variation of the shape parameter p measured on speech STFT coefficients as a function of the window length L (left) and as a function of frequency f with $L = 64$ ms (right). The black curve and the gray area represent respectively the geometric mean and the geometric standard deviation of the measured values. This illustration is motivated by the fact that the measured values are approximately log-Gaussian.

3.1 Unconstrained expression

We focus in the rest of this article on the simple case where $J = I + 1$ and $\mathbf{A}(f)$ has full row rank. Under this assumption, the constrained optimization problem (3) is equivalent to a one-dimensional unconstrained complex optimization problem [1]. The MAP source coefficients are then expressed as

$$\widehat{\mathbf{S}}(n, f) = \mathbf{O} + \widehat{u}\mathbf{V} \quad (4)$$

where \mathbf{O} is any vector satisfying the constraint, *e.g.* $\mathbf{O} = \mathbf{A}(f)^\dagger \mathbf{X}(n, f)$ with † denoting pseudo-inversion, \mathbf{V} is any vector spanning the null space of $\mathbf{A}(f)$ and

$$\widehat{u} = \arg \min_{u \in \mathbb{C}} \|\mathbf{O} + u\mathbf{V}\|_p^p. \quad (5)$$

This optimization problem has to be solved for each STFT bin (n, f) individually. Using the complex derivative notation [13], the first and second order derivatives of the criterion $\mathcal{L}(u) = \|\mathbf{O} + u\mathbf{V}\|_p^p$ are given by

$$\frac{\partial \mathcal{L}}{\partial \bar{u}} = \frac{\overline{\partial \mathcal{L}}}{\partial u} = \frac{p}{2} \sum_{j=1}^J |O_j + uV_j|^{p-2} \overline{V_j} (O_j + uV_j) \quad (6)$$

$$\frac{\partial^2 \mathcal{L}}{\partial u \partial \bar{u}} = \frac{\partial^2 \mathcal{L}}{\partial \bar{u} \partial u} = \frac{p^2}{4} \sum_{j=1}^J |O_j + uV_j|^{p-2} |V_j|^2 \quad (7)$$

$$\frac{\partial^2 \mathcal{L}}{\partial \bar{u}^2} = \frac{\overline{\partial^2 \mathcal{L}}}{\partial u^2} = \frac{p(p-2)}{4} \sum_{j=1}^J |O_j + uV_j|^{p-4} \overline{V_j}^2 (O_j + uV_j)^2 \quad (8)$$

3.2 Singular and non-singular solutions

It is well known that for real variables the global minimum of \mathcal{L} with $p \leq 1$ results in at least one source coefficient being zero and can be found by combinatorial optimization [6,1]. However, this is not true anymore with complex variables, as shown in [1,7] in the particular case $p = 1$. Nevertheless, the local minima of \mathcal{L} can still be characterized using the two lemmas below.

Lemma 1. *Let $\mathcal{J} = \{j : V_j \neq 0\}$. When $p < 1$, the points $z_j = -\frac{O_j}{V_j}$, $j \in \mathcal{J}$, are singular (i.e. non-differentiable) local minima of \mathcal{L} .*

Proof. Let $\mathcal{Z}_j = \{k : z_k = z_j\}$. The point z_j is characterized by the fact that $S_k(n, f) = 0$ for all $k \in \mathcal{Z}_j$ and $S_k(n, f) \neq 0$ for all $k \notin \mathcal{Z}_j$. By developing the expression of \mathcal{L} around this point when $p < 1$, we get

$$\mathcal{L}(z_j + u) = \mathcal{L}(z_j) + \left(\sum_{k \in \mathcal{Z}_j} |V_k|^p \right) |u|^p + O(u). \quad (9)$$

Thus \mathcal{L} is non-differentiable at z_j and $\mathcal{L}(z_j + u) > \mathcal{L}(z_j)$ for small $u \neq 0$. \square

Lemma 2. *The other local minima of \mathcal{L} are non-singular and within the convex hull of z_j , $j \in \mathcal{J}$.*

Proof. If $u \neq z_j$ for all j and u is a local minimum of \mathcal{L} , then \mathcal{L} is differentiable at u according to (6) and $\frac{\partial \mathcal{L}}{\partial u} = 0$. After rearranging this equality, we get

$$u = \frac{\sum_{j \in \mathcal{J}} |O_j + uV_j|^{p-2} |V_j|^2 z_j}{\sum_{j \in \mathcal{J}} |O_j + uV_j|^{p-2} |V_j|^2}. \quad (10)$$

Thus u can be expressed as a weighted sum of z_j , $j \in \mathcal{J}$, with positive weights summing to one. \square

In the following, we use the term ‘singular’ to characterize by extension the local minima of \mathcal{L} where at least one source coefficient is zero, although \mathcal{L} is differentiable at these minima when $p > 1$.

3.3 Critical value of p for the existence of non-singular solutions

The above distinction between singular and non-singular local minima raises the question whether non-singular minima can exist for all values of p and whether the global minimum can be non-singular. We studied this question experimentally with $I = 2$ and $J = 3$.

We draw 100 independent source coefficient vectors following the generalized Gaussian distribution with shape parameter $p = 0.4$ using the Metropolis-Hastings algorithm [14]. We also draw 100 instantaneous (real) mixing matrices of the form $A_{1j} = \cos(\theta_j)$ and $A_{2j} = \sin(\theta_j)$ with θ_j uniformly distributed in $[-\frac{\pi}{2}, \frac{\pi}{2}]$ and 100 convolutive (complex) mixing matrices of the form $A_{1j} = 1$ and

$A_{2j} = e^{2i\pi\theta_j}$ with θ_j uniformly distributed in $(-\pi, \pi]$. The multiplication of each source coefficient vector by each matrix resulted in a total of 10000 instantaneous mixtures and 10000 convolutive mixtures.

For each mixture, we tested whether non-singular minima of \mathcal{L} existed and whether the global minimum was non-singular as follows. Given Lemma 2, we sampled \mathcal{L} on a discrete grid spanning the convex hull of z_j , $1 \leq j \leq J$, containing points of the form $u = \frac{k_1}{3K}z_1 + \frac{k_2}{3K}z_2 + \frac{3K-k_1-k_2}{3K}z_3$, $1 \leq k_j \leq K$, with $K = 50$, and we selected the global minimum \tilde{u} on this grid. If \tilde{u} was non-singular, then the true global minimum was necessarily non-singular. Otherwise, we decided that the global minimum was singular. In the latter case, we also sampled the gradient and the Hessian of \mathcal{L} on the same grid, selected as a potential local minimum the point with the smallest gradient among all points with positive definite Hessian and refined it using the `fminunc` optimizer. We then observed whether the optimizer converged to a non-singular local minimum or not.

The results were very similar for instantaneous and convolutive mixtures. The average percentage of mixture draws resulting in a non-singular local minimum or a non-singular global minimum is depicted in Figure 2 as a function of p . Both quantities decrease with decreasing p , with a large drop around $p = 1$. For $p \lesssim 0.75$, there remains a few non-singular local minima, but no global minima. This can be illustrated in a more general case using the example below.




Fig. 2. Percentage of mixture draws resulting in a non-singular local minimum (dashed curve) or a non-singular global minimum (plain curve) of the l_p norm criterion with three-source two-channel mixtures.

Example 1. Let $O_j = e^{2i\pi \frac{j}{J}}$ and $V_j = 1$, $1 \leq j \leq J$. Then $u = 0$ is a non-singular local minimum of \mathcal{L} for all $p > 0$. Furthermore, the value of \mathcal{L} at this minimum is smaller than at all singular local minima for $p > p_{\text{crit}}$ where p_{crit} is defined implicitly by $\sum_{j=1}^{J-1} |1 - O_j|^{p_{\text{crit}}} = J$ and equals respectively 0.738, 0.612, 0.534, 0.481 for $J = 3, 4, 5, 6$, and decreases with increasing J .

Proof. Using the fact that $\sum_{j=1}^J O_j = \sum_{j=1}^J O_j^2 = 0$, the coefficients of the complex gradient, the diagonal coefficients of the complex Hessian and the off-diagonal coefficients of the complex Hessian of \mathcal{L} , defined in [13], are given respectively at $u = 0$ by $\frac{\partial \mathcal{L}}{\partial \bar{u}} = \frac{\partial \mathcal{L}}{\partial u} = 0$, $\frac{\partial^2 \mathcal{L}}{\partial u \partial \bar{u}} = \frac{\partial^2 \mathcal{L}}{\partial \bar{u} \partial u} = \frac{Jp^2}{4}$ and $\frac{\partial^2 \mathcal{L}}{\partial \bar{u}^2} = \frac{\partial^2 \mathcal{L}}{\partial u^2} = 0$. Thus the complex gradient is zero and the complex Hessian is positive-definite, which proves that $u = 0$ is a non-singular local minimum of \mathcal{L} [13].

The values of the criterion at this non-singular minimum and at the singular local minima $z_j = -O_j$ are given by $\mathcal{L}(0) = J$ and $\mathcal{L}(z_j) = \sum_{j=1}^{J-1} |1 - O_j|^p$ for all j . The latter is a strictly increasing function of p . Indeed, it can be checked that $\frac{d\mathcal{L}(z_j)}{dp} = \log J > 0$ at $p = 0$ and $\frac{d^2\mathcal{L}(z_j)}{dp^2} > 0$ for all $p > 0$. Thus $\mathcal{L}(0) > \mathcal{L}(z_j)$ if and only if $p > p_{\text{crit}}$ where p_{crit} is the value of p such that $\mathcal{L}(0) = \mathcal{L}(z_j)$. \square

This shows that the global minimum of \mathcal{L} can be non-singular when $p > p_{\text{crit}}$. We conjecture that p_{crit} is the lowest value of p for which this can happen.

Conjecture 1. The global minimum of \mathcal{L} with $p < p_{\text{crit}}$ is always singular.

We have not yet managed to prove this conjecture mathematically. However we verified it experimentally with $J = 3$ (see Figure 2) and with $4 \leq J \leq 6$ using the same number of mixture draws and a similar discrete grid for optimization.

3.4 Efficient optimization algorithm

Lemmas 1 and 2 and Conjecture 1 suggest the following efficient algorithm for the estimation of the MAP source coefficients.

- If $p \geq 1$, run any gradient-based optimizer initialized randomly using (6)–(8).
- If $p \leq p_{\text{crit}}$, sample the criterion at the singular points z_j , $1 \leq j \leq J$, and select the minimum of the criterion among these points.
- If $p_{\text{crit}} < p < 1$, sample the criterion on a discrete grid spanning the convex hull of the singular points z_j , $1 \leq j \leq J$, and containing these points. Select the minimum of the criterion on this grid. If it is non-singular, refine it via any gradient-based optimizer using (6)–(8).

Provided that Conjecture 1 is true, this algorithm is guaranteed to find the global minimum of the criterion when $p \geq 1$ or $p \leq p_{\text{crit}}$, but also when $p_{\text{crit}} < p < 1$ if the discrete grid is tight enough. Moreover, this algorithm is quite fast, particularly for small p . Using Matlab on a 1.2 GHz CPU with the `fminunc` optimizer and the discrete grid defined in Section 3.3, the computation time equals on average 0.15 s, 0.0065 s and 0.00025 s for $p = 1$, $p = 0.9$ and $p = 0.5$ respectively with $I = 2$ and $J = 3$. By contrast, the optimization via Second Order Cone Programming (SOCP) for $p = 1$ takes about 0.36 s, using the same Matlab toolbox as in [1,7].

4 Source separation results

We evaluated the proposed algorithm for the separation of 10 instantaneous and 10 convolutive speech mixtures with $I = 2$ and $J = 3$. The mixture signals were obtained by mixing the source signals of Section 2 either with a matrix of positive coefficients or with a set of simulated room impulse responses corresponding to a reverberation time of 250 ms, as described in [12]. Following [11,12], the STFT window length L was set to 512 (64 ms) for instantaneous mixtures and 2048 (256 ms) for convolutive mixtures. The frequency-domain complex mixing matrices $\mathbf{A}(f)$ were computed by Fourier transform of the mixing filters in the convolutive case. The performance was measured in decibels (dB) for each estimated source by $\text{SDR}_j = 20 \log_{10}(\|s_j\|/\|\hat{s}_j - s_j\|)$ and subsequently averaged. For comparison, we also evaluated the performance when the criterion was optimized over the singular points only, as suggested in [1].

The results are shown in Figure 3. In the instantaneous case, the best SDR is achieved for $p = 1$ and the SDR for smaller values of p is 0.6 dB smaller. In the convolutive case, the best SDR is achieved for $p \rightarrow 0$ and it is 1.2 dB larger than for $p = 1$. Note that the ML value of p determined in Section 2, namely $p = 0.4$, does not result in the best SDR. This suggests that algorithms estimating the value of p from the data might not improve performance. Note also that the SDR remains much smaller than the theoretical upper bound computed in [12].




Fig. 3. Average SDR as a function of p for the separation of instantaneous (left) and convolutive (right) three-source two-channel speech mixtures (plain curve: optimization over the full space, dashed curve: optimization over the singular points only).

5 Conclusion

In this article, we investigated the benefit of modeling the sources via generalized Gaussian priors instead of Laplacian priors for underdetermined source

separation. This generalization is not straightforward, since the resulting l_p norm criterion is nonconvex for $p < 1$. In the simple case where $J = I + 1$, we characterized mathematically the local minima of this criterion, conjectured that the global maximum is always singular below a critical value of p and derived an efficient global optimization algorithm. We evaluated this algorithm on speech mixtures and showed that small values of p resulted in the best separation performance in the convolutive case, but also in the fastest optimization. This work raises further research issues, including the proof of the above conjecture and the extension of the proposed algorithm when $J > I + 1$.

References

1. Winter, S., Sawada, H., Makino, S.: On real and complex valued l_1 -norm minimization for overcomplete blind source separation. In: Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA). (2005) 86–89
2. Vielva, L., Erdoğmuş, D., Príncipe, J.C.: Underdetermined blind source separation using a probabilistic source sparsity model. In: Proc. Int. Conf. on Independent Component Analysis and Blind Source Separation (ICA). (2001) 675–679
3. Davies, M.E., Mitianoudis, N.: Simple mixture model for sparse overcomplete ICA. IEE Proceedings on Vision, Image and Signal Processing **151** (2004) 35–43
4. Févotte, C., Godsill, S.J.: A Bayesian approach for blind separation of sparse sources. IEEE Trans. on Audio, Speech and Language Processing **14** (2006) 2174–2188
5. Lee, T.W., Lewicki, M.S., Girolami, M.A., Sejnowski, T.J.: Blind source separation of more sources than mixtures using overcomplete representations. IEEE Signal Processing Letters **6** (1999) 87–90
6. Zibulevsky, M., Pearlmutter, B.A., Bofill, P., Kisilev, P.: Blind source separation by sparse decomposition in a signal dictionary. In: Independent Component Analysis : Principles and Practice. Cambridge Press (2001) 181–208
7. Bofill, P., Monte, E.: Underdetermined convoluted source reconstruction using LP and SOCP, and a neural approximator of the optimizer. In: Proc. Int. Conf. on Independent Component Analysis and Blind Source Separation (ICA). (2006) 569–576
8. Choi, S., Cichocki, A., Amari, S.: Flexible independent component analysis. In: Neural Networks for Signal Processing (NNSP 8). (1998) 83–92
9. Everson, R., Roberts, S.: Independent component analysis: a flexible nonlinearity and decorrelating manifold approach. Neural Computation **11** (1999) 1957–1983
10. Wu, H.C., Príncipe, J.C.: Generalized anti-Hebbian learning for source separation. In: Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP). (1999) II–1073–1076
11. Yılmaz, O., Rickard, S.T.: Blind separation of speech mixtures via time-frequency masking. IEEE Trans. on Signal Processing **52** (2004) 1830–1847
12. Vincent, E., Gribonval, R., Plumbley, M.D.: Oracle estimators for the benchmarking of source separation algorithms. Signal Processing **87** (2007) 1933–1950
13. van den Bos, A.: Complex gradient and Hessian. IEE Proceedings on Vision, Image and Signal Processing **141** (1994) 380–382
14. Casella, G., Robert, C.P.: Monte Carlo Statistical Methods, 2nd Edition. Springer, New York, NY (2005)