

SPARSITY-BASED RESTORATION OF SMOS IMAGES IN THE PRESENCE OF OUTLIERS

J. Preciozzi, P. Musé

IIE - Fac. Ingeniería
Univ. de la República
Uruguay

A. Almansa[†], S. Durand[‡]

CNRS - France
[†] LTCI - Telecom ParisTech
[‡] MAP5 - Univ. René Descartes

F. Cabot, Y. Kerr
A. Khazaal, B. Rougé

CESBIO
31400 Toulouse
France

ABSTRACT

Estimates of soil moisture and surface salinity are of significant importance to improve meteorological and climate prediction. The SMOS mission monitor these quantities, by measuring the brightness temperature by means of L-band aperture synthesis interferometry. Despite the L-band being reserved for Earth and space exploration, SMOS images reveal large number of strong outliers, produced by illegal antennas emitting in this band. In this work we propose a variational approach to recover a super-resolved, denoised brightness temperature map. The measurements are modeled as the superposition of three super-resolved components in the spatial domain: the target brightness temperature map u , an image o modeling the outliers, and Gaussian noise n . This decomposition allows to isolate each of its constituent parts, thanks to a sparsity operator that acts on o , and a bounded variation prior on u that extrapolates its spectrum promoting a non-oscillating behavior. The proposed model is interesting in itself, as it is general enough to be applied to other restoration problems. Experiments on real and synthetic data confirm the suitability of the proposed approach.

1. INTRODUCTION

Surface soil moisture (SSM) and sea surface salinity (SSS) measures provide extremely valuable information for meteorological and climate predictions. A direct way to monitor SSM and SSS is through the use of L-band microwave radiometers systems. These sensors present an important practical limitation: to achieve a suitable ground resolution would require an antenna of prohibitive size. Hopefully, this limitation can be overcome using interferometry instruments such as MIRAS (Microwave Imaging Radiometer by Aperture Synthesis). MIRAS payload is carried on the SMOS satellite, launched on November 2009.

Interferometric radiometers measure the complex cross-correlation between two signals collected by two different antennas. Each correlation is a sample of the so called *visibility function*, given by [3]:

$$V_{k,l} = \frac{1}{\Omega_k \Omega_l} \iint_{\|\xi\| \leq 1} U_k(\xi) U_l^*(\xi) (T_b(\xi) - T_r) \tilde{r}_{kl}(t) \frac{e^{-i2\pi \mathbf{u}_{kl}^T \xi}}{\sqrt{1 - \|\xi\|^2}} d\xi, \quad (1)$$

where \mathbf{u}_{kl} are the frequency baselines pair associated to antennas A_k and A_l ; U_k, U_l are the normalized voltage patterns and Ω_k, Ω_l are the solid angles of the two antennas A_k and A_l respectively. The Cartesian coordinates $\xi = (\xi_1, \xi_2)$ are the spatial domain coordinates, restricted to the unit circle. T_r is the physical temperature of the receivers (assumed the same for all receivers); \tilde{r}_{kl} is the

Fringe-Wash function, a function of the spatial delay $t = \frac{\mathbf{u}_{kl}^T \xi}{f_0}$, where $f_0 = \frac{c}{\lambda_0}$ is the central frequency of observation. Note that the brightness temperature T_b is a 2D function restricted to the unit circle ($\|\xi\| \leq 1$).

In MIRAS, the instrument is composed of three arms on a Y-shaped configuration, where each arm consists of an array of antennas. This configuration leads to a hexagonal grid where the *visibility function* is sampled. Figure 1 shows the star shaped domain Ω obtained from the MIRAS array configuration.

If we note $T = T_b - T_r$, samples of T in a hexagonal grid could be obtained from the visibility samples by solving the linear system $\mathbf{G}T = V$, where matrix \mathbf{G} corresponds to the discrete linear operator derived from Eq 1. However, this inverse problem is ill-posed since \mathbf{G} is not invertible (due to the lack of information beyond Ω). As usual, additional constraints must be added to the model in order to obtain a well posed problem.

In [4] the authors propose to regularize the inverse problem by imposing that T have no frequency components beyond Ω . Hence, if P_Ω denotes the spectral projection onto Ω :

$$\min_T \|V - \mathbf{G}T\|_2^2 \quad \text{s.t. } (I - P_\Omega)T = 0. \quad (2)$$

Note that the projection operator can be written as $P_\Omega = F^{-1}Z_\Omega Z_\Omega^* F$ in terms of the Discrete Fourier Transform operator F^1 , the zero padding operator Z_Ω and its adjoint, the restriction to Ω . Thus the non-zero Fourier coefficients of the solution T to the previous problem can be written as $D = \hat{T} = Z_\Omega^* F T$, and T can be recovered from \hat{T} by $T = F^{-1}Z_\Omega D$.

The previous description is a simplified model of how the SMOS L1B data product $D = \hat{T}$ is obtained from the raw L1A visibility measurements V . As we shall see later the simple Fourier inversion $T = F^{-1}Z_\Omega \hat{T}$ leads to potentially very strong Gibbs effects which are partially alleviated (as proposed by [4]) by the use of a Blackman window B : $T = F^{-1}BZ_\Omega \hat{T}$, which can be seen as a kind of linear Tikhonov regularization:

$$T = \arg \min_T \|B(F T - D)\|_2^2. \quad (3)$$

Such linear approaches are however not well suited for the restoration of SMOS images, because usually the measurements are polluted by outliers. Indeed, despite the fact that according to international radio regulations, the L-band in which MIRAS radiometer operates is exclusively allocated to the Earth Exploration Satellite Service, space research and radio astronomy, soon after SMOS was launched the data revealed there were many signals being transmit-

¹F is the Fourier Transform defined on an hexagonal grid

ted within this protected passive band, rendering a great amount of data unusable for scientific purposes. Because these outliers have frequencies beyond Ω , very strong Gibbs effects can be seen on the final brightness temperature images (see Figure 2).

The recovery of the brightness temperature original images is the main motivation of the present work, where we seek to detect and remove signal effects generated from illegal emissions (outliers), while at the same time extrapolating the image spectrum in order to minimize Gibbs effects. As a byproduct we shall also partially reduce measurement noise.




Fig. 1. Spectral domain Ω associated to MIRAS Y-shaped instrument.




Fig. 2. Brightness temperature obtained by the inverse Fourier transform of the acquired data.

2. PROPOSED METHOD

From the previous paragraph we can summarize the image formation model as follows:

$$D = Z_{\Omega}^*(F(u + o)) + \hat{n}, \quad (4)$$

where u is the non-polluted brightness temperature image, o is the outliers' image generated by illegal signals being transmitted, and \hat{n} is an additive Gaussian noise with covariance matrix $\sigma^2 W^{-2}$. This covariance can be derived from the matrix G or approximated (as we do here) by a diagonal matrix from the knowledge of the redundancy of each Fourier coefficient in D with respect to the measured visibilities V .

Hence, $W\hat{n} = D - Z_{\Omega}^*(F(u + o))$ is a white Gaussian noise with zero mean and variance σ^2 . This motivates the following method.

2.1. Variational formulation

We propose to recover u in Eq. (4) by solving the following constrained optimization problem:

$$\begin{aligned} \min_{u,o} \{ \text{TV}(u) + \mu S(o) \} \\ \text{s.t. } \|W(F(o + u) - D)\|_2^2 \leq |\Omega|\sigma^2. \end{aligned} \quad (5)$$

The constraint is the data fidelity term. Data D are SMOS L1B data product, and correspond to the Fourier coefficients of the brightness temperature. The redundancy matrix W weighs repeated measurements, and masks frequencies outside the spectral support Ω . Parameter σ^2 is the measurement noise variance, whose value is known.

The objective function acts as a regularizer, and consists of two terms that were designed in order to separate the outliers o from the original brightness temperature map u . The first term is the total variation semi-norm of u , intended to super-resolve u beyond spectral support Ω while avoiding Gibbs oscillations. The second term

results from applying some sparsity operator to o (in practice ℓ_0 or ℓ_1 norms) in order to recover a very sparse solution for o . The parameter μ controls the trade-off between both terms and its choice is discussed in Section 2.2.2.

Instead of solving the constrained minimization problem (5), we will solve the following unconstrained minimization:

$$\min_{u,o} \{ \|W(F(o + u) - D)\|_2^2 + \lambda(\text{TV}(u) + \mu S(o)) \}. \quad (6)$$

It can be shown that, for a given σ^2 , there exist a unique Lagrange multiplier λ such that this problem has a unique solution, given by the equivalent constrained problem (5). Here the value of λ is derived from that of σ^2 using a version of Uzawa's algorithm proposed in [5].

In order to reduce the "staircasing" effect inherent to many TV minimization methods, we follow an approach based on the *Spectral TV* introduced by Moisan in [6]. Staircasing reduction is achieved by: (i) Computing image derivatives not by finite differences but analytically on Fourier series expansion; (ii) Approximating the continuous TV as a Riemann integral over a grid at least two times finer than the critical sampling rate.

In our case, instead of doubling the sampling rate of ∇u , we chose to reduce the spectral domain of u to an intermediate cell \mathcal{H} , in-between the star domain Ω where measures \hat{u} are taken, and the cell \mathcal{C} corresponding to the (largely overcritical) spatial sampling rate of u .

The final method can be stated as follows:

$$\begin{aligned} \min_{u,o} \{ \|W(F(o + u) - D)\|_2^2 + \lambda(\text{TV}(u) + \mu S(o)) \} \\ \text{s.t. } \text{supp } \hat{u} \subseteq \mathcal{H} \subset \mathcal{C}. \end{aligned} \quad (7)$$

2.2. Numerical implementation

To solve the minimization problem (7) while achieving the sparsest possible solution, we proceed in two steps. First, we solve (7) choosing as sparsity term $S(o) = \|o\|_1$. Under this conditions, the problem is convex and can be solved using a Forward-Backward algorithm which is known to converge to a global minimum. Then, to reinforce the sparsity of the solution o obtained in the first step, we consider the problem with $S(o) = \|o\|_0$, which is also solved iteratively using the Forward-Backward method, starting from the previous solution. While this problem is non-convex due to the ℓ_0 norm, Blumensath and Davies [8] have shown that for this functional the Forward-Backward algorithm converges to a local minimum.

2.2.1. Combettes-Wajs Forward-Backward implementation

Combettes and Wajs Forward-Backward algorithm [7] is designed to minimize a functional $E(x) = E_1(x) + E_2(x)$, where E_1 and E_2 are convex functions such that: E_1 is differentiable with Lipschitz gradient; E_2 is a *simple* function, in the sense that its associated proximal operator can be easily computed as:

$$\text{prox}_{\gamma E_2}(x) = \arg \inf_y E_2(y) + \frac{1}{2\gamma} \|x - y\|^2. \quad (8)$$

If these hypotheses hold, the following generic algorithm can be derived: starting with an arbitrary x^0 , set for each $k \in \mathbb{N}$,

$$\begin{cases} x^{k+1/2} &= x^k - \gamma \nabla E_1(x^k) \\ x^{k+1} &= \text{prox}_{\gamma E_2}(x^{k+1/2}). \end{cases}$$

In order to ensure convergence to the minimizer, γ must be smaller than $2/L$, where L is the Lipschitz constant of ∇E_1 .

For our problem, we define

$$E_1(u, o) = \frac{1}{2} \|W(F(o+u) - D)\|_2^2$$

$$E_2(u, o) = \lambda \text{TV}(u) + \lambda \mu \|o\|_1.$$

Direct differentiation of $E_1(u, o)$ yields

$$\nabla E_1(u, o) = (F^* W^2 (F(u+o) - D), F^* W^2 (F(u+o) - D)),$$

where F^* is the adjoint operator of F , that is, the inverse Fourier transform. As for the proximal operator of $\gamma E_2(u, o)$, we have

$$\text{prox}_{\gamma E_2}(u, o) = \left(\text{prox}_{\gamma \lambda \text{TV}}(u), \text{prox}_{\gamma \lambda \mu \|\cdot\|_1}(o) \right).$$

For $\text{prox}_{\gamma \lambda \text{TV}}$, we implemented Chambolle's algorithm [5], where in each iteration we included a spectral projection into \mathcal{H} in order to force the constraint in (7). As for $\text{prox}_{\gamma \lambda \mu \|\cdot\|_1}(o)$, it is straightforward to show from (8) that it corresponds to the *soft thresholding* operator:

$$s_\Delta(t) = \begin{cases} \text{sign}(t)(|t| - \Delta) & \text{if } |t| \geq \Delta, \\ 0 & \text{if } |t| < \Delta. \end{cases}$$

To summarize, each iteration for the first step ($S(o) = \|o\|_1$) can be expressed as follows:

$$\begin{cases} u^{k+1/2} &= u^k - \gamma(F^* W^2 F(o+u) - F^* W^2 D) \\ o^{k+1/2} &= o^k - \gamma(F^* W^2 F(o+u) - F^* W^2 D) \\ u^{k+1} &= \text{prox}_{\gamma \lambda \text{TV}}(u^{k+1/2}) \\ o^{k+1} &= s_{\gamma \lambda \mu}(o^{k+1/2}). \end{cases}$$

This iteration converges to a global minimum, that corresponds to the solution of problem (7) with sparsity operator $S(o) = \|o\|_1$.

We now proceed to the second step, where the sparsity operator is chosen to be $S(o) = \|o\|_0$. For this problem, the same Forward-Backward method can be considered and is guaranteed to converge to a local minimizer. Now, instead of the soft thresholding, the proximal operator for $S(o)$ becomes the hard thresholding $h_{\sqrt{2\gamma\lambda\mu}}$. Note that the changes in the previous iteration, associated to replacing $\|o\|_1$ by $\|o\|_0$, are easily implemented by transforming the soft thresholding into a hard thresholding.

2.2.2. Selection of μ

We propose to balance the trade-off between the sparsity term $S(o)$ and the regularity term $\text{TV}(u)$ based on modeling the outliers as follows. For a cylinder c of radius r and height h , the involved norms or semi-norms are:

$$\text{TV}(c) = 2\pi r h$$

$$\|c\|_0 = \pi r^2 \mathbb{1}_{[h>0]}$$

$$\|c\|_1 = \pi r^2 h.$$

The selection of μ determines whether the cylinder is considered to be an outlier or part of the image to be recovered (that is, as part of o or u , respectively).

When the sparsity operator is the ℓ_1 norm, c is considered as an outlier if $\text{TV}(c) \geq \mu \|c\|_1$, leading to a $\mu \leq \frac{2}{r}$. In the examples, we have selected a value $\mu \sim \frac{2}{10} = 0.2$, which amounts to consider that the radii of the outliers are at most 10 pixels wide.

In the case of the ℓ_0 norm, a cylinder is considered as an outlier as soon as $\text{TV}(c) \geq \mu \|c\|_0$, i.e. $\mu \leq \frac{2h}{r}$. Hence, μ can then be interpreted as the minimal h/r outlier ratio. Here, contrarily to the ℓ_1 case, the height plays an important role in distinguishing an outlier from the data, and larger values for μ may be chosen ($\mu \sim 100 \frac{2}{10} = 20$).

3. EXPERIMENTS

We present two kind of experiments. In the first one, we compare results from our approach to those obtained by previous works. Experiments are run on several snapshots from SMOS L1B dataset from march 2010. We have set σ equal to $5K$, which is the measurement error reported by the SMOS mission.

Figure 3 shows the results obtained for snapshot 996, which corresponds to a region of central Europe. Note that the acquired image is corrupted with several outliers that considerably degrade the data. It is clear that our method outperforms both the direct inverse Fourier transform and Blackman apodization.

In the second experiment, we perform a quantitative evaluation of our method. To do this, we simulated a new snapshot based on real data in the following way. First, we generated an image o_{gt} composed of synthetic outliers, with spots with diameter ranging between 1 and 3 pixels and with radiometric values between [-30.000, 30.000]. We selected the snapshot 700 corresponding to Madagascar, denoted by u_{gt} , that contains no outliers but is still contaminated by acquisition noise. Hence, if we call u_0 the ideal, noise-free image, and n_0 the acquisition noise, we have: $u_{gt} = u_0 + n_0$. Finally we added both images and filtering out frequency components outside Ω , as the data u_{gt} comes from the instrument and hence lies on this spectral region. In short, the synthetic data is given by:

$$D = P_\Omega(u_{gt} + o_{gt})$$

Let u be the image brightness temperature, o the image outliers and n the image noise obtained from our method. Figure 4 shows the results obtained on this data. The table shows the estimation error on u and o , for different error norms. These errors were computed only over Ω in the Fourier domain, since the spectral support of our ground-truth is Ω . In all cases, the standard deviation of the residual $\|WF(o+u) - D\|/\sqrt{|\Omega|}$ residual $\sqrt{\mathbb{E}(\hat{\sigma}^2)}$ was 5 ± 0.1 , which is in agreement with the restriction imposed to Uzawa's algorithm.

4. CONCLUSIONS

We presented a method to restore images from the L1B SMOS data product. This product consists of measurements of the ground brightness temperature, obtained by interferometry in the spectral domain. The MIRAS instrument is such that the acquisition is obtained in an anisotropic spectral domain, which is moreover smaller than that of the target resolution. In addition, the measurements are severely corrupted by measurement noise and punctual sources (antennas) emitting in forbidden frequency bands. The proposed variational approach simultaneously tackles all these difficulties by modeling the observations as the superposition of three super-resolved components in the spatial domain: the target brightness temperature map u , an image o modeling the outliers due to the illegal emissions, and colored Gaussian noise n . This decomposition allows to isolate each of its constituent parts, thanks to a sparsity enforcing term that acts on o , and a total variation semi-norm that extrapolates the spectral domain of u while promoting a non-oscillating behavior. It is worth noting that the method is based

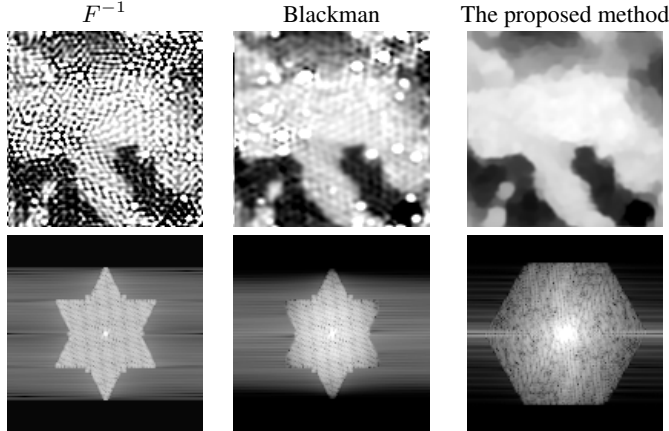


Fig. 3. Comparison between previous work and our method using snapshot 996 from March 2010. The first row shows the temperature map scaled to (0,300) Kelvin, and the second row its corresponding Fourier spectrum. The first column is the result of applying inverse Fourier transform directly. The second column is obtained after applying a Blackman apodization scheme. The third column shows the results obtained with the method proposed here.

on general principles and hence can be applied to other restoration problems.


We performed experimental evaluation on real SMOS measurements, including outliers within a range of 1 to 100 times the dynamic range of the brightness temperature (about 300K). The results obtained from this evaluation proved to be consistent, in the sense that the energy of the noise component is no more than 10% away from the postulated noise level (3.75K), and that the brightness temperature and outliers components exhibit reasonable structure according to the experts knowledge of the ground truth. Experiments over snapshots that appeared not to contain outliers, yielded no detections in o . The previous conclusions on u and n also held.

The suitability of the method was also evaluated on synthetic data. The ground truth was built by superposing sets of synthetic outliers to an outlier-free, real noisy SMOS snapshot. Results were consistent with the ground truth: the estimated noise n is also within 10% of 3.75K, and all the outliers were detected and well localized. On uniform regions of the ground truth, the estimation error in the brightness temperature map u was in the order of magnitude of the measurement noise. In non-uniform regions, the error grows as a result of the (unavoidable) effect of regularization.

4.1. Future work

More accurate results could be obtained if the model was modified to include directly the raw L1A visibilities in the data fidelity constraint as $\|C^{-\frac{1}{2}}(G(o+u) - V)\|_2^2 \leq |\Omega|\sigma^2$. The expected benefit of such a modification would be twofold: (i) a more detailed noise modeling can be incorporated through the visibilities covariance matrix C ; and (ii) no intermediate linear approximation needs to be done (L1B product): the final result is directly obtained from a non-linear optimization on the original raw data V .

However, this approach could not be used in conjunction with the iterative non-differentiable, non-convex optimization methods required by this application. Indeed, the matrix G being very dense and each iteration would require a prohibitive number of operations.



	L_1	L_2	L_∞
$P_\Omega(u - u_{gt})$	11.038	16.557	149.215
$P_\Omega(o - o_{gt})$	9.026	14.482	147.905

Fig. 4. Results obtained with simulated data. Note that since no information outside Ω is available, all errors have been measured after projecting with P_Ω .

By considering the L1B data product D , as was done in this first approach, the bulk of the work of evaluating and inverting G is done only once, during the initialization step that computes D from V .

5. REFERENCES

- [1] Y. H. Kerr, P. Waldteufel, J-P. Wigneron, J-M. Martinuzzi, J. Font, and M. Berger, "Soil moisture retrieval from space: The soil moisture and ocean salinity (smos) mission," *IEEE Trans. Geosci. Remote Sensing*, vol. 39(8), pp. 1729–1735, 2002.
- [2] M. Martin-Neira, Y. Menard, J.M. Goutoule, and U. Kraft, "Miras, a two-dimensional aperture synthesis radiometer," in *Geoscience and Remote Sensing Symposium, 1994. IGARSS '94. Surface and Atmospheric Remote Sensing: Technologies, Data Analysis and Interpretation., International*, aug 1994, vol. 3, pp. 1323–1325 vol.3.
- [3] I. Corbella, N. Duffo, M. Vall-llossera, A. Camps, and F. Torres, "The visibility function in interferometric aperture synthesis radiometry," *IEEE Trans. Geosci. Remote Sensing*, vol. 42(8), pp. 1677–1682, 2004.
- [4] E. Anterrieu, "A resolving matrix approach for synthetic aperture imaging radiometers," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 42, no. 8, pp. 1649 – 1656, aug. 2004.
- [5] A. Chambolle, "An algorithm for total variation minimization and application," *JMIV*, vol. 20(1), 2004.
- [6] L. Moisan, "How to discretize the total variation of an image?," *ICIAM07 Minisymposia*, 2007.
- [7] P. L. Combettes and V. R. Wajs, "Signal recovery by proximal forward-backward splitting," *SIAM Journal on Multiscale Modeling and Simulation*, vol. 4(4), pp. 1168–1200, 2005.
- [8] Thomas Blumensath and Mike E Davies, "Iterative thresholding for sparse approximations," *The Journal of Fourier Analysis and Applications*, vol. 62, pp. 291–294, January 2008.