A NEW ADMM ALGORITHM FOR THE EUCLIDEAN MEDIAN AND ITS APPLICATION TO ROBUST PATCH REGRESSION

Kunal N. Chaudhury K. R. Ramakrishnan

Department of Electrical Engineering, Indian Institute of Science, Bangalore, India

ABSTRACT

The Euclidean Median (EM) of a set of points Ω in an Euclidean space is the point x minimizing the (weighted) sum of the Euclidean distances of \boldsymbol{x} to the points in Ω . While there exits no closed-form expression for the EM, it can nevertheless be computed using iterative methods such as the Weiszfeld algorithm. The EM has classically been used as a robust estimator of centrality for multivariate data. It was recently demonstrated that the EM can be used to perform robust patch-based denoising of images by generalizing the popular Non-Local Means algorithm. In this paper, we propose a novel algorithm for computing the EM (and its box-constrained counterpart) using variable splitting and the method of augmented Lagrangian. The attractive feature of this approach is that the subproblems involved in the ADMM-based optimization of the augmented Lagrangian can be resolved using simple closed-form projections. The proposed ADMM solver is used for robust patch-based image denoising and is shown to exhibit faster convergence compared to an existing solver.

Index Terms— Image denoising, patch-based algorithm, robustness, Euclidean median, variable splitting, augmented Lagrangian, alternating direction method of multipliers (ADMM), convergence.

1. INTRODUCTION

Suppose we are given a set of points $a_1, \ldots, a_n \in \mathbf{R}^d$, and some non-negative weights w_1, \ldots, w_n to these points. Consider the problem

$$\min_{\boldsymbol{x}\in\mathcal{C}} \quad \sum_{k=1}^{n} w_k \|\boldsymbol{x} - \boldsymbol{a}_k\|_2 \tag{1}$$

where $\|\cdot\|_2$ is the Euclidean norm and $C \subset \mathbb{R}^d$ is closed and convex. This is a convex optimization problem. The (global) minimizer of (1) when C is the entire space is called the *Euclidean median* (also referred to as the geometric median) [1]. The Euclidean median has classically been used as a robust estimator of centrality for multivariate data [1]. This multivariate generalization of the scalar median also comes up in transport engineering, where (1) is used to model the problem of locating a facility that minimizes the cost of transportation [2]. More recently, it was demonstrated in [3, 4] that the Euclidean median and its non-convex extensions can be used to perform robust patch-based regression for image denoising by generalizing the poplar Non-Local Means algorithm [5, 6].

Note that the minimizer of (1) is simply the projection of the Euclidean median onto C. Unfortunately, there is no simple closed-form expression for the Euclidean median, even when it is unique [7]. Nevertheless, the Euclidean median can be computed using numerical methods such as the iteratively reweighted least squares (IRLS). One such iterative method is the so-called Weiszfeld algorithm [8]. This

particular algorithm, however, is known to be prone to convergence problems [2]. Geometric optimization-based algorithms have also be proposed for approximating the Euclidean median [7]. More recently, motivated by the use of IRLS for sparse signal processing and L_1 minimization [9, 10], an IRLS-based algorithm for approximating the Euclidean median was proposed in [3]. More precisely, here the authors consider a smooth convex surrogate of (1), namely,

$$\sum_{k=1}^{n} w_k \left(\| \boldsymbol{x} - \boldsymbol{a}_k \|_2^2 + \varepsilon \right)^{1/2} \qquad (\varepsilon > 0),$$
(2)

and describe an iterative method for optimizing (2). The convergence properties of this iterative algorithm was later studied in [11].

In the last few years, there has been a lot of renewed interest in the use of variable splitting, the method of multipliers, and the alternating direction method of multipliers (ADMM) [12] for various non-smooth optimization problems in signal processing [13, 14, 15]. Motivated by this line of work, we introduce a novel algorithm for approximating the solution of (1) using variable splitting and the augmented Lagrangian in Section 2. An important technical distinction of the present approach is that, instead of working with the smooth surrogate (2), we directly address the original non-smooth objective in (1). The attractive feature here is that the subproblems involved in the ADMM-based optimization of the augmented Lagrangian can be resolved using simple closed-form projections. In Section 3, the proposed algorithm is used for robust patch-based denoising of images following the proposal in [3]. In particular, we incorporate the information that the pixels of the original (clean) patches take values in a given dynamic range (for example, in [0,255] for gravscale images) using an appropriately defined C. Numerical experiments show that the iterates of the proposed algorithm converge much more rapidly than the IRLS iterates for the denoising method in question. Moreover, the reconstruction quality of the denoised image obtained after few ADMM iterations is often substantially better than the IRLS counterpart.

2. EUCLIDEAN MEDIAN USING ADMM

Since we will primarily work with (1) in this paper, we will overload terminology and refer to the minimizer of (1) as the Euclidean median. The main idea behind the ADMM algorithm is to decompose the objective in (1) into a sum of independent objectives, which can then be optimized separately. This is precisely achieved using variable splitting [13, 14]. In particular, note that we can equivalently formulate (1) as

$$\min_{\boldsymbol{x}_1,\dots,\boldsymbol{x}_n} \quad \iota_{\mathcal{C}}(\boldsymbol{z}) + \sum_{k=1}^n f_k(\boldsymbol{x}_k)$$

ubject to $\boldsymbol{x}_k - \boldsymbol{z} = 0 \quad (k = 1,\dots,n),$
(3)

z

Correspondence: kunal@ee.iisc.ernet.in. The first author was partially supported by the Startup Grant provided by the Indian Institute of Science.

where

$$f_k(\boldsymbol{x}_k) = w_k \| \boldsymbol{x}_k - \boldsymbol{a}_k \|_2, \qquad (4)$$

and $\iota_{\mathcal{C}}$ is the indicator function of \mathcal{C} [13]. In other words, we artificially introduce one local variable for each term of the objective, and a common global variable that forces the local variables to be equal [16]. The advantage of this decomposition will be evident shortly.

Note that in the process of decomposing the objective, we have introduced additional constraints. While one could potentially use any of the existing constrained optimization methods to address (3), we will adopt the method of augmented Lagrangian which (among other things) is particularly tailored to the separable structure of the objective. The augmented Lagrangian for (3) can be written as

$$\mathcal{L} = \iota_{\mathcal{C}}(\boldsymbol{z}) + \sum_{k=1}^{n} \left\{ f_k(\boldsymbol{x}_k) + \boldsymbol{y}_k^T(\boldsymbol{x}_k - \boldsymbol{z}) + \frac{\mu}{2} \| \boldsymbol{x}_k - \boldsymbol{z} \|_2^2 \right\}$$
(5)

where y_1, \ldots, y_n are the Lagrange multipliers corresponding to the equality constraints in (3), and $\mu > 0$ is a penalty parameter [12, 13].

We next use the machinery of generalized ADMM to optimize \mathcal{L} jointly with respect to the variables and the Lagrange multipliers. The generalized ADMM algorithm proceeds as follows [13, 16]: In each iteration, we sequentially minimize \mathcal{L} with respect to the variables x_1, x_2, \ldots, x_n, z in a Gauss-Seidel fashion (keeping the remaining variables fixed), and then we update the Lagrange multipliers y_1, \ldots, y_n using a fixed dual-ascent rule. In particular, note that the minimization over x_k at the (t + 1)-th iteration is given by

$$\boldsymbol{x}_{k}^{(t+1)} = \underset{\boldsymbol{x}}{\operatorname{argmin}} \ f_{k}(\boldsymbol{x}) + (\boldsymbol{x} - \boldsymbol{z}^{(t)})^{T} \boldsymbol{y}_{k}^{(t)} + \frac{\mu}{2} \|\boldsymbol{x} - \boldsymbol{z}^{(t)}\|_{2}^{2} \ (6)$$

where $z^{(t)}$ denotes the variable z at the end of the *t*-th iteration. A similar notation is used to denote the Lagrange multipliers at the end of the *t*-th iteration.

Next, by combining the quadratic and linear terms and discarding terms that do not depend on z, we can express the partial minimization over z as

$$\begin{aligned} \boldsymbol{z}^{(t+1)} &= \operatorname*{argmin}_{\boldsymbol{z}} \ \iota_{\mathcal{C}}(\boldsymbol{z}) + \frac{\mu}{2} \sum_{k=1}^{N} \|\boldsymbol{z} - \boldsymbol{x}_{k}^{(t+1)} - \mu^{-1} \boldsymbol{y}_{k}^{(t)}\|_{2}^{2} \\ &= \operatorname*{argmin}_{\boldsymbol{z} \in \mathcal{C}} \ \|\boldsymbol{z} - \boldsymbol{z}_{0}^{(t+1)}\|_{2}^{2} \end{aligned}$$

where

$$\boldsymbol{z}_{0}^{(t+1)} = \frac{1}{N} \sum_{k=1}^{N} \left(\boldsymbol{x}_{k}^{(t+1)} + \mu^{-1} \boldsymbol{y}_{k}^{(t)} \right).$$

(1 + 1)

Therefore,

$$\boldsymbol{z}^{(t+1)} = \mathscr{P}_{\mathcal{C}}(\boldsymbol{z}_0^{(t+1)}) \tag{7}$$

(4 + 1)

where $\mathscr{P}_{\mathcal{C}}(\boldsymbol{x})$ denotes the projection of \boldsymbol{x} onto the convex set \mathcal{C} . Note that we have used the most recent \boldsymbol{x}_k variables in (7). Finally, we update the Lagrange multipliers using the following standard rule:

$$\boldsymbol{y}_{k}^{(t+1)} = \boldsymbol{y}_{k}^{(t)} + \mu \, (\boldsymbol{x}_{k}^{(t+1)} - \boldsymbol{z}^{(t+1)}). \tag{8}$$

We loop over (6), (7) and (8) until some convergence criteria is met, or up to a fixed number of iterations. For a detailed exposition on the augmented Lagrangian, the method of multipliers, and the ADMM, we refer the interested reader to [12, 13, 16].

We now focus on the update in (6). By combining the linear and quadratic terms (and discarding terms that do not depend on x), we can express (6) in the following general form:

$$\boldsymbol{x}^{\star} = \operatorname*{argmin}_{\boldsymbol{x}} f(\boldsymbol{x}) + \frac{1}{2} \|\boldsymbol{x} - \boldsymbol{v}\|_{2}^{2}$$
(9)

where

$$f(\boldsymbol{x}) = \lambda \|\boldsymbol{x} - \boldsymbol{u}\|_2. \tag{10}$$

Note that the objective in (9) is closed, proper, and strongly convex, and hence x^* exists and is unique. In fact, we immediately recognize x^* to be the *proximal map* of f evaluated at u [13, 17].

Definition 2.1 For any $f : \mathbf{R}^d \mapsto (-\infty, \infty]$ that is closed, proper, and convex, the proximal map $\Psi_f : \mathbf{R}^d \mapsto \mathbf{R}^d$ is defined to be

$$\Psi_f(\boldsymbol{v}) = \operatorname*{argmin}_{\boldsymbol{x}} f(\boldsymbol{x}) + \frac{1}{2} \|\boldsymbol{x} - \boldsymbol{v}\|_2^2 \qquad (\boldsymbol{v} \in \mathbf{R}^d).$$
(11)

As per this definition, $\boldsymbol{x}^{\star} = \Psi_f(\boldsymbol{v})$.

Note that (10) involves the L_2 norm, which unlike the L_1 norm is not separable. A coordinate-wise minimization of (9) is thus not possible (as is the case for the L_1 norm leading to the well-known shrinkage function). However, we have the following result on proximal maps [13, 17].

Theorem 2.2 Let $f : \mathbf{R}^d \mapsto (-\infty, \infty]$ be closed, proper, and convex. Then we can decompose any $x \in \mathbf{R}^d$ as

$$\boldsymbol{x} = \Psi_f(\boldsymbol{x}) + \Psi_{f^\star}(\boldsymbol{x}), \tag{12}$$

where $f^* : \mathbf{R}^d \mapsto (-\infty, \infty]$ is the convex conjugate of f,

$$f^{\star}(\boldsymbol{y}) = \max_{\boldsymbol{x}} \, \boldsymbol{x}^{T} \boldsymbol{y} - f(\boldsymbol{x}) \qquad (\boldsymbol{y} \in \mathbf{R}^{d}). \tag{13}$$

The usefulness of (12) is that both f^* and and its proximal map can be obtained in closed-form. Indeed, by change-of-variables,

$$egin{aligned} f^{\star}(oldsymbol{y}) &= \max_{oldsymbol{x}} oldsymbol{x}^T oldsymbol{y} - \lambda \, \|oldsymbol{x} - oldsymbol{u}\|_2 \ &= oldsymbol{u}^T oldsymbol{y} + \max_{oldsymbol{x}} oldsymbol{x}^T oldsymbol{y} - \lambda \, \|oldsymbol{x}\|_2 \ \end{aligned}$$

Now, if $\|\boldsymbol{y}\|_2 \leq \lambda$, then by Cauchy-Schwarz, $\boldsymbol{x}^T \boldsymbol{y} \leq \lambda \|\boldsymbol{x}\|_2$. In this case, the maximum over \boldsymbol{x} is 0. On the other hand, if $\|\boldsymbol{y}\|_2 > \lambda$, then setting $\boldsymbol{x} = t\boldsymbol{y}$ (t > 0), we have

$$\boldsymbol{x}^{T}\boldsymbol{y} - \lambda \|\boldsymbol{x}\|_{2} = t \|\boldsymbol{y}\|_{2} (\|\boldsymbol{y}\|_{2} - \lambda) > 0,$$

which can be made arbitrarily large by letting $t \to \infty$. Thus,

$$f^{\star}(\boldsymbol{y}) = \boldsymbol{u}^T \boldsymbol{y} + \iota_{\mathcal{B}}(\boldsymbol{y}), \qquad (14)$$

where $\iota_{\mathcal{B}}$ is the indicator function of the ball \mathcal{B} with centre 0 and radius λ ,

$$\iota_{\mathcal{B}}(\boldsymbol{y}) = \begin{cases} 0 & \text{if } \|\boldsymbol{y}\|_2 \leq \lambda, \\ \infty & \text{else.} \end{cases}$$

Having obtained (14), we have from (11),

$$\begin{split} \Psi_{f^{\star}}(\boldsymbol{v}) &= \operatorname*{argmin}_{\boldsymbol{x}} \ \boldsymbol{u}^T \boldsymbol{x} + \iota_{\mathcal{B}}(\boldsymbol{x}) + \frac{1}{2} \|\boldsymbol{x} - \boldsymbol{v}\|_2^2 \\ &= \operatorname*{argmin}_{\boldsymbol{x} \in \mathcal{B}} \ \frac{1}{2} \|\boldsymbol{x} - (\boldsymbol{v} - \boldsymbol{u})\|_2^2, \end{split}$$

which is precisely the projection of v - u onto \mathcal{B} . This is explicitly given by

$$\Psi_{f^{\star}}(\boldsymbol{v}) = \min(\lambda, \|\boldsymbol{v} - \boldsymbol{u}\|_2) \frac{\boldsymbol{v} - \boldsymbol{u}}{\|\boldsymbol{v} - \boldsymbol{u}\|_2}.$$
 (15)

Combining (12) with (15), we have

$$\Psi_f(\boldsymbol{v}) = \Psi_{\lambda}(\boldsymbol{v}, \boldsymbol{u}) = \boldsymbol{v} - \min(\lambda, \|\boldsymbol{v} - \boldsymbol{u}\|_2) \frac{\boldsymbol{v} - \boldsymbol{u}}{\|\boldsymbol{v} - \boldsymbol{u}\|_2}.$$
 (16)

It is reassuring to note that $\Psi_{\lambda}(v; u)$ equals v when $\lambda = 0$, and equals u when $\lambda = \infty$. In the context of the ADMM update (6), note that

$$\lambda = \mu^{-1} w_k, \ \boldsymbol{u} = \boldsymbol{a}_k, \ \text{ and } \boldsymbol{v} = \boldsymbol{z}^{(t)} - \mu^{-1} \boldsymbol{y}^{(t)}_k$$

The overall ADMM algorithm (called EM-ADMM) for computing the Euclidean median is summarized in Algorithm 1.

Data: Dimension d, points
$$a_1, \ldots, a_n \in \mathbb{R}^d$$
, and $\mu > 0$.
Result: z.
Initialize: $z, y_1, \ldots, y_n \in \mathbb{R}^d$;
loop
 $r = 0$;
for $k = 1, \ldots, n$ do
 $\begin{vmatrix} x_k = \Psi_{\mu^{-1}w_k}(z - \mu^{-1}y_k, a_k); \\ r = r + x_k + \mu^{-1}y_k; \\ end$
 $z = \mathscr{P}_{\mathcal{C}}(N^{-1}r);$
for $k = 1, \ldots, n$ do
 $\mid y_k = y_k + \mu(x_k - z);$
end
end

Algorithm 1: Euclidean Median using ADMM (EM-ADMM)

We next demonstrate how EM-ADMM can be used for robust patch-based denoising of images. We also investigate its convergence behavior in relation to the IRLS algorithm in [3].

3. APPLICATION: ROBUST PATCH REGRESSION

In the last ten years or so, some very effective patch-based algorithms for image denoising have been proposed [5, 18, 19, 20, 21]. One of the outstanding proposals in this area is the Non-Local Means (NLM) algorithm [5]. Several improvements and adaptations of NLM have been proposed since its inception, some of which currently offer stateof-the-art denoising results [21]. Recently, it was demonstrated in [3] that the denoising performance of NLM can be further improved by introducing the robust Euclidean median into the NLM framework.

We consider the denoising setup where we are given a noisy image $g = (g_i)_{i \in I}$ (*I* is some linear ordering of the pixels in the image) that is derived from some clean image $f = (f_i)_{i \in I}$ via

$$\boldsymbol{g}_i = \boldsymbol{f}_i + \sigma \, \boldsymbol{\xi}_i \qquad (i \in I),$$

where $\boldsymbol{\xi} = (\boldsymbol{\xi}_i)_{i \in I}$ is iid $\mathcal{N}(0, 1)$. The denoising problem is one of estimating the unknown \boldsymbol{f} from the noisy measurement \boldsymbol{g} .

For any pixel $i \in I$, let patch \mathbf{P}_i denote the restriction of g to a square window centered at i. Letting k be the length of this window, this associates every pixel i with a patch-vector \mathbf{P}_i of length k^2 . For $i, j \in I$, define the weights $w_{ij} = \exp(-||\mathbf{P}_i - \mathbf{P}_j||_2^2/h^2)$, where h > 0 is a smoothing parameter. For every $i \in I$, we consider the following optimization problem

$$\hat{\mathbf{P}}_{i} = \underset{\mathbf{P} \in \mathbf{R}^{k^{2}}}{\operatorname{argmin}} \ \iota_{\mathcal{C}}(\mathbf{P}) + \sum_{j \in S(i)} w_{ij} \|\mathbf{P} - \mathbf{P}_{j}\|_{2}.$$
(17)

Here S(i) denotes the neighborhood pixels of *i*, namely, those pixels that are contained in a window of size $S \times S$ centred at *i*. The convex set C is defined to be the collection of patches whose pixel intensities are in the dynamic range [l, u]; e.g., l = 0 and u = 255 for a



Fig. 1. PSNR evolution with iterations for the *Barbara* image at two different noise levels. For the ADMM, we used $\mu = 10^{-3}$, and for the IRLS, we set $\varepsilon = 10^{-6}$. In either case, we use the noisy patch to seed the iterations (please see main text for further details).

grayscale image. The projection onto C can be computed coordinatewise at negligible cost. In particular, for $1 \le i \le d$,

$$\mathscr{P}_{\mathcal{C}}(oldsymbol{x})[i] = egin{cases} oldsymbol{x}[i] & ext{if } l \leq oldsymbol{x}[i] \leq u, \ l & ext{if } oldsymbol{x}[i] < l, ext{ and } \ u & ext{if } oldsymbol{x}[i] > u. \end{cases}$$

Notice that (17) is exactly the optimization in (1). In other words, we denoise each patch by taking the Euclidean median of its neighbouring noisy patches. The center pixel of the denoised patch is taken to be the denoised pixel. The overall denoising scheme is called the Non-Local Euclidean Medians (in short, NLEM) [3].

Table 1. Denoising performance of (a) NLM, (b) NLEM-IRLS and (c) NLEM-ADMM at $\sigma = 10, 20, 40, 60$, and 80 (PSNRs averaged over 10 noise realizations). Parameters: $S = 21, k = 7, h = 10\sigma$.

Image	Method	PSNR (dB)				
House	(a)	34.22	29.78	25.20	23.37	22.35
	(b)	32.66	29.01	26.68	24.78	23.46
	(c)	34.13	30.77	27.04	24.87	23.47
Lena	(a)	33.24	29.31	26.17	24.54	23.64
	(b)	32.54	29.48	27.31	25.38	24.37
	(c)	33.57	30.38	27.59	25.42	24.38
Peppers	(a)	32.32	27.66	23.11	21.03	19.99
	(b)	30.95	26.95	24.31	22.84	21.24
	(c)	32.14	28.54	25.06	22.98	21.26
Barbara	(a)	32.37	27.39	23.53	22.05	21.34
	(b)	30.77	27.28	25.55	23.77	22.61
	(c)	32.43	28.84	25.67	23.77	22.62

The denoising experiments were performed on several grayscale test images. We used the standard NLM parameters for all the experiments [5]: S = 21, k = 7, and $h = 10\sigma$. The proposed EM-ADMM algorithm is used for computing the solution of (17). Notice that the

main computation here is that of determining the distances in (16). Since we require a total of S^2 such distance evaluations in k^2 dimensions, the complexity is $O(S^2k^2)$ per iteration per pixel. Incidentally, this is also the complexity of the IRLS algorithm in [3]. For all the denoising experiments, we initialized the Lagrange multipliers in EM-ADMM to zero vectors. We considered two possible initializations for z, namely, the noisy patch and the patch obtained using NLM. Exhaustive experiments showed that the best convergence results are obtained for $\mu \sim 10^{-3}$. For example, figure 1 shows the evolution of the peak-signal-to-noise ratio (PSNR) over the first ten iterations obtained using the ADMM algorithm. Also show in this figure is the evolution of the PSNR for the IRLS algorithm from [3]. In either case, we used the noisy patch as the initialization. Notice that the increase in PSNR is much more rapid with ADMM compared to IRLS. In fact, in about 4 iterations, the optimal PSNR is attained with ADMM. Notice that at $\sigma = 10$, the PSNR from IRLS grows much more slowly compared to ADMM.



Fig. 2. Convergence of EM-ADMM and IRLS at a given pixel of the *House* image at $\sigma = 40$. The inset shows the position of the particular pixel where the patch regression is performed. The objective functions shown in the plot correspond to (1) and (2) that are respectively optimized by ADMM and IRLS (note that the objectives are very close since we use $\varepsilon = 10^{-6}$ in (2)). The noisy image patch was used to initialize both algorithms. For EM-ADMM, the objective function converges to the optimal value (up to to three decimal places) in just 2 iterations. The convergence is relatively slow for IRLS (accuracy up to three decimals obtained after 6 iterations).

The above observation is also supported by the convergence result shown in figure 2 for the *House* image. Notice that ADMM converges in just 2 iterations, while IRLS takes about 6 iterations to attain the same accuracy (we recall that the cost per iteration is comparable). In general, we observed that the within 2 iterations the ADMM result converges to an accuracy that is sufficient for the present denoising application. At higher noise levels ($\sigma > 60$), about 4 iterations are required. This applies to all the test images that we have experimented with.

The denoising results for some test images obtained using NLM and NLEM are provided in Table 1. For NLEM, we used both the ADMM and IRLS solvers. Following the previous observations, we used 4 iterations for both solvers. For $\sigma \leq 60$, we used the noisy

patch to initialize the iterations, and for $\sigma > 60$ we used the NLM patch. This initialization scheme was found to give the best PSNR results. At large noise levels, there is not much difference after 4 iterations. However, at low noise levels, the PSNR obtained using ADMM is often substantially larger than that obtained using IRLS after 4 iterations. The reason for this is that IRLS converges really slow in such situations. This is evident from the PSNR plots in figure 1 at $\sigma = 10$. For a visual comparison, a particular denoising result for the *Barbara* image obtained using NLM and NLEM (using the ADMM solver) is shown in figure 3. Roughly speaking, at the cost of just four NLMs, we are able to improve the PSNR by more than 2 dB. Notice that the NLEM result is much more sharp compared to the NLM counterpart.



(c) NLM (PSNR =23.53 dB) (d) NELM (PSNR = 25.67 dB).

Fig. 3. Denoising results using NLM and NLEM at $\sigma = 40$. We used the proposed EM-ADMM algorithm for computing the Euclidean medians at each pixel (please see the main text for other details about the experiment). The EM-ADMM algorithm was initialized using the noisy patch and run for four iterations.

4. CONCLUSION

We proposed a new algorithm for computing the (constrained) Euclidean median using variable splitting and the augmented Lagrangian. In particular, we demonstrated how the ADMM-based optimization of the augmented Lagrangian can be resolved using simple closed-form projections. The proposed algorithm was used for image denoising using the Non-Local Euclidean Medians, and was generally found to exhibit faster convergence compared to the IRLS algorithm. One interesting direction that we did not pursue is to adapt the penalty μ at each iteration, which is known to speed up the convergence [16]. Yet another interesting direction is the use of accelerated ADMM [22] to further speed up the convergence.

5. REFERENCES

- P. H. Lopuhaa and J. P. Rousseeuw, "Breakdown points of affine equivariant estimators of multivariate location and covariance matrices," *Annals of Statistics*, vol. 19, no. 1, pp. 229-248, 1991.
- [2] H. A. Eiselt and V. Marianov, *Foundations of Location Analysis*. International Series in Operations Research and Management Science, New York: Springer, 2011.
- [3] K. N. Chaudhury and A. Singer, "Non-local Euclidean medians," *IEEE Signal Processing Letters*, vol. 19, no. 11, pp. 745 - 748, 2012.
- [4] K. N. Chaudhury and A. Singer, "Non-local patch regression: Robust image denoising in patch space," *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, pp.1345-1349, 2013.
- [5] A. Buades, B. Coll, and J. M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Modeling and Simulation*, vol. 4, pp. 490-530, 2005.
- [6] A. Buades, B. Coll, and J. M. Morel, "Image denoising methods. A new nonlocal principle," *SIAM Review*, vol. 52, pp. 113-147, 2010.
- [7] P. Bose, A. Maheshwari, and P. Morin, "Fast approximations for sums of distances, clustering and the Fermat–Weber problem," *Computational Geometry*, vol. 24, no. 3, pp. 135-146, 2003.
- [8] E. Weiszfeld, "Sur le point par lequel le somme des distances de *n* points donnes est minimum," *Tohoku Mathematical Journal*, vol. 43, pp. 355-386, 1937.
- [9] R. Chartrand and Y. Wotao, "Iteratively reweighted algorithms for compressive sensing," *Proc. IEEE International Conference* on Acoustics, Speech and Signal Processing, pp. 3869-3872, 2008.
- [10] I. Daubechies, R. Devore, M. Fornasier, and C. S. Gunturk "Iteratively reweighted least squares minimization for sparse recovery," *Communications on Pure and Applied Mathematics*, vol. 63, pp. 1-38, 2009.
- [11] K. N. Chaudhury, "On the convergence of the IRLS algorithm in non-local patch regression," *IEEE Signal Processing Letters*, vol. 20, no. 8, 815-818, 2013.
- [12] J. Eckstein and D. P. Bertsekas, "On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Mathematical Programming*, vol. 55, no. 1-3, pp. 293-318, 1992.
- [13] P. L. Combettes and J. C. Pesquet, "Proximal splitting methods in signal processing," *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, pp. 185-212, Springer New York, 2011.
- [14] M. V. Afonso, J. M. Bioucas-Dias, and M. A. Figueiredo, "Fast image recovery using variable splitting and constrained optimization," *IEEE Transactions on Image Processing*, vol. 19, no. 9, pp. 2345-2356, 2010.
- [15] M. V. Afonso, J. M. Bioucas-Dias, and M. A. Figueiredo, "An augmented Lagrangian approach to the constrained optimization formulation of imaging inverse problems," *IEEE Transactions* on *Image Processing*, vol. 20, no. 3, pp. 681-695, 2011.
- [16] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1-122, 2011.

- [17] R. T. Rockafellar, *Convex Analysis*. Princeton University Press, 1970.
- [18] C. Kervrann, J. Boulanger, "Optimal spatial adaptation for patch-based image denoising," *IEEE Transactions on Image Processing*, vol. 15(10), 2866-2878, 2006.
- [19] M. Aharon, M. Elad, A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Transactions on Image Processing*, vol. 54, no. 11, pp. 4311-4322, 2006.
- [20] K. Dabov, A. Foi, V. Katkovnik, K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Transactions on Image Processing*, vol. 16, pp. 2080-2095, 2007.
- [21] P. Milanfar, "A tour of modern image filtering: New insights and methods, both practical and theoretical," *IEEE Signal Processing Magazine*, vol. 30, no. 1, pp. 106-128, 2013.
- [22] T. Goldstein, B. O'Donoghue, S. Setzer, and R. Baraniuk, "Fast alternating direction optimization methods," *CAM report*, pp. 12-35, 2012.