Digital Signal Processing 27 (2014) 107–118

Contents lists available at ScienceDirect

Digital Signal Processing

www.elsevier.com/locate/dsp

Fast solver for some computational imaging problems: A regularized weighted least-squares approach

B. Zhang [*], S. Makram-Ebeid, R. Prevost, G. Pizaine

Medisys, Philips Research, Suresnes, France

ARTICLE INFO

Article history:
Available online 31 January 2014

Keywords:
Regularized weighted least-squares
Preconditioned conjugate gradient
Preconditioning
Condition number

ABSTRACT

In this paper we propose to solve a range of computational imaging problems under a unified perspective of a regularized weighted least-squares (RWLS) framework. These problems include data smoothing and completion, edge-preserving filtering, gradient-vector flow estimation, and image registration. Although originally very different, they are special cases of the RWLS model using different data weightings and regularization penalties. Numerically, we propose a preconditioned conjugate gradient scheme which is particularly efficient in solving RWLS problems. We provide a detailed analysis of the system conditioning justifying our choice of the preconditioner that improves the convergence. This numerical solver, which is simple, scalable and parallelizable, is found to outperform most of the existing schemes for these imaging problems in terms of convergence rate.

1. Introduction

In this paper we propose to solve some classical imaging problems with a unified quadratic optimization perspective. These topics include data smoothing and completion, edge-preserving filtering, gradient-vector flow estimation, and image registration. We are particularly interested in high-performance numerical solvers for these problems, as they are widely used as building blocks of numerous applications in many domains such as computer vision and medical imaging [1].

Concretely, we look at the framework of the regularized weighted least-squares (RWLS):

$$\arg\min_{u:\mathbb{R}^D\to\mathbb{R}} J(u) := \int_{\mathbb{R}^D} w(\mathbf{x})\big(u(\mathbf{x}) - u_0(\mathbf{x})\big)^2\, d\mathbf{x} + \gamma^{2\alpha} \int_{\mathbb{R}^D} \big|\mathbf{L}_\alpha u(\mathbf{x})\big|^2\, d\mathbf{x} \tag{1}$$

Here, $u_0(\mathbf{x}) \in \mathbb{R}$ represents the observed measurement at point $\mathbf{x} \in \mathbb{R}^D$, and $u(\mathbf{x})$ the data to estimate. $w(\mathbf{x}) \geqslant 0$ are non-negative weights, which can be considered as confidence levels of the measurements. $\mathbf{L}_\alpha u : \mathbb{R}^D \mapsto \mathbb{R}^n$ represents some regularization operator $\mathbf{L}_\alpha$ applying a penalty on u. We will restrict $\mathbf{L}_\alpha$ to be a linear fractional differential operator of order $\alpha > 0$, such as the gradient and the Laplacian. Further, $\gamma > 0$ is a trade-off parameter between the weighted data-fidelity term and the regularity-penalty term.

In this work we will focus on the discrete RWLS problem, or the discrete counterpart of Eq. (1):

$$\arg\min_{\mathbf{u}\in\mathbb{R}^N} J(\mathbf{u}) := \big\|\mathbf{W}^{\frac{1}{2}}(\mathbf{u} - \mathbf{u}_0)\big\|^2 + \gamma^{2\alpha}\|\mathbf{L}_\alpha \mathbf{u}\|^2 \tag{2}$$

Here, $\mathbf{u}_0 \in \mathbb{R}^N$ is the vector of the measurements of length N. For multi-dimensional measurements, the data are assumed to be vectorized in the lexicographical order. $\mathbf{W} \in \mathbb{R}^{N\times N}$ stands for a diagonal weighting matrix with the weights on its diagonal $\mathbf{W}_{i,i} = w_i \geqslant 0$ for $i = 0, \ldots, N-1$. $\mathbf{L}_\alpha$ in the discrete setting will be a matrix representing the differential operator and we keep the same notation. It will be clear (see Section 4) that each of our aforementioned imaging problems fits Eq. (2) by choosing a particular set of weights $\mathbf{W}$ and a particular regularization operator $\mathbf{L}_\alpha$.

Our main contribution here is proposing an efficient preconditioned conjugate gradient (PCG) scheme which solves RWLS, and hence the above imaging problems. We provide a detailed analysis of the system conditioning justifying our choice of the preconditioner that improves the convergence. Surprisingly, this simple solver is found to outperform most of the state-of-the-art numerical schemes proposed for those problems. In particular, the convergence rate of PCG is spectacular, with a gain up to an order of magnitude observed in some of our experiments. Additionally, the PCG has the advantages of being easily implementable, scalable and parallelizable.

This paper is organized as follows. Section 2 describes in detail the RWLS framework, and the proposed PCG solver. Section 3 analyzes the choice of the preconditioner by showing its potential in reducing the condition number of the problem and hence improving the convergence rate. Then, Section 4 presents the different imaging problems revisited and solved by the RWLS approach. We

* Corresponding author.
E-mail address: bo.wang-zhang@philips.com (B. Zhang).

1051-2004/$ – see front matter
http://dx.doi.org/10.1016/j.dsp.2014.01.007

show the superior performance of our method compared to various existing schemes. We also discuss an extension of the RWLS model in Section 5. Our conclusions are drawn in Section 6. Finally, mathematical details are deferred to the appendices.

2. Regularized weighted least-squares and a PCG solver

2.1. Notations in the 1D case

Let us use an example in the 1-dimensional (1D) RWLS to introduce our notations and present our main results. The setting can be easily extended to multi-dimensional cases (Section 2.5).

Consider the following RWLS in the continuous setting where the regularization operator is the first derivative (i.e., $\mathbf{L}_\alpha = d/dx$ with $\alpha = 1$):

$$\arg\min_u J(\mathbf{u}) = \int_{\mathbb{R}} w(x)\big(u(x) - u_0(x)\big)^2\,dx + \gamma^2 \int_{\mathbb{R}} u'(x)^2\,dx \tag{3}$$

This choice makes Eq. (3) a Dirichlet regularized regression problem. Its solution is the stationary point to the associated Euler–Lagrange equation:

$$w(x)u(x) - \gamma^2 u''(x) = w(x)u_0(x), \quad x \in \mathbb{R} \tag{4}$$

In the discrete version, the operator $\mathbf{L}_1$ will be represented by a first-order finite-difference matrix. For example, let $\mathbf{L}_1$ be the following circulant matrix (Eq. (5)), which corresponds to a filter $g_1 = [1, -1]/h_1$ with a periodic boundary condition.

$$\mathbf{L}_1 := \frac{1}{h_1}\begin{bmatrix} -1 & 1 & 0 & \cdots & 0 \\ 0 & -1 & 1 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & -1 & 1 \\ 1 & 0 & \cdots & 0 & -1 \end{bmatrix} \tag{5}$$

Here $h_1 > 0$ represents the finite-difference spacing.

To solve Eq. (2), one sets the gradient of $J(\mathbf{u})$ to zero, and obtain a linear system which is no more than the discrete counterpart of Eq. (4):

$$\mathbf{Au} = \mathbf{b}, \quad \text{where } \mathbf{A} := \mathbf{W} + \gamma^2 \mathbf{L}_1^* \mathbf{L}_1 \text{ and } \mathbf{b} := \mathbf{W}\mathbf{u}_0 \tag{6}$$

We used $\mathbf{L}_1^*$ to denote the conjugate transpose of $\mathbf{L}_1$. It follows that $(-\mathbf{L}_1^*\mathbf{L}_1)$ is a Hermitian matrix which represents a second-order differential filter [2] $g_2 = [1, -2, 1]/h_1^2$. In addition, $\mathbf{A}$ is Hermitian and semi-positive definite.

$\mathbf{L}_1$ is diagonalizable by the fast Fourier transform (FFT) matrix $\mathbf{F}$ and its k-th eigenvalue is $\lambda_k = (e^{-j\omega_k} - 1)/h_1$ with $\omega_k := 2\pi k/N$:

$$\mathbf{L}_1 = \mathbf{F}^* \Lambda \mathbf{F}, \quad \Lambda := \operatorname{diag}[\lambda_0, \lambda_1, \ldots, \lambda_{N-1}]$$

Due to the orthonormality of FFT, one has $\mathbf{F}^*\mathbf{F} = \mathbf{I}$ where $\mathbf{I}$ is the identity matrix. Therefore $\mathbf{A}$ can be rewritten as:

$$\mathbf{A} = \mathbf{W} + \gamma^2 \mathbf{F}^* \tilde{\Lambda} \mathbf{F}, \quad \tilde{\Lambda} := |\Lambda|^2 := \Lambda^* \Lambda$$

where $\tilde{\Lambda}$ is the diagonal matrix of the eigenvalues $\tilde{\lambda}_k$ of $\mathbf{L}_1^*\mathbf{L}_1$ which are given by $\tilde{\lambda}_k := |\lambda_k|^2 = [\frac{2}{h_1}\sin(\omega_k/2)]^2$.

The FFT choice above is clearly due to the assumed periodic boundary condition. More generally, the Hermitian matrix $\mathbf{L}_1^*\mathbf{L}_1$ always possesses an orthonormal eigen-decomposition:

$$\mathbf{L}_1^*\mathbf{L}_1 = \mathbf{B}^* |\Lambda|^2 \mathbf{B}, \quad |\Lambda|^2 := \operatorname{diag}\big[|\lambda_0|^2, |\lambda_1|^2, \ldots, |\lambda_{N-1}|^2\big]$$

where $\mathbf{B}$ is some orthonormal matrix, and $(|\lambda_k|^2)_{k=0,\ldots,N-1}$ are the eigenvalues of $\mathbf{L}_1^*\mathbf{L}_1$ written in the modulus form to emphasize their non-negative nature. In practice, the basis $\mathbf{B}$ will represent trigonometric transforms, i.e. FFT, DCT (discrete cosine transform), and DST (discrete sine transform), with a periodic, an even symmetric, and an odd-symmetric boundary conditions [2] respectively imposed on the matrix $\mathbf{L}_1^*\mathbf{L}_1$.

Consequently, for any order $\alpha > 0$ one can define $\mathbf{L}_\alpha^*\mathbf{L}_\alpha$ to be a fractional differential operator such that:

$$\mathbf{L}_\alpha^*\mathbf{L}_\alpha := \mathbf{B}^* |\Lambda|^{2\alpha} \mathbf{B}, \quad |\Lambda|^{2\alpha} := \operatorname{diag}\big[|\lambda_0|^{2\alpha}, |\lambda_1|^{2\alpha}, \ldots, |\lambda_{N-1}|^{2\alpha}\big] \tag{7}$$

We will keep noting the spectrum of $\mathbf{L}_\alpha^*\mathbf{L}_\alpha$ by

$$\tilde{\Lambda} := |\Lambda|^{2\alpha}, \qquad \tilde{\lambda}_k := |\lambda_k|^{2\alpha} \tag{8}$$

In the subsequent presentation, we will concentrate on the periodic boundary condition (i.e., $\mathbf{B} = \mathbf{F}$).

Similar to Eq. (6), for an arbitrary $\alpha > 0$, the minimizer of Eq. (2) is the solution to the following linear system:

$$\mathbf{Au} = \mathbf{b}, \quad \text{where } \mathbf{A} := \mathbf{W} + \gamma^{2\alpha} \mathbf{L}_\alpha^* \mathbf{L}_\alpha \text{ and } \mathbf{b} := \mathbf{W}\mathbf{u}_0 \tag{9}$$

where $\mathbf{A}$ is Hermitian and semi-positive definite, and can be written as:

$$\mathbf{A} = \mathbf{W} + \gamma^{2\alpha} \mathbf{F}^* \tilde{\Lambda} \mathbf{F} \tag{10}$$

The k-th eigenvalue of $\mathbf{L}_\alpha^*\mathbf{L}_\alpha$ is given by $\tilde{\lambda}_k = |\lambda_k|^{2\alpha} = [\frac{2}{h_1} \times \sin(\omega_k/2)]^{2\alpha}$. These definitions will be extended to multi-dimensional case in Section 2.5.

2.2. Case of constant weights: a linear filtering

We keep considering the 1D RWLS problem. If the weights are everywhere constant (say $\mathbf{W} = \bar{w}\mathbf{I}$ for some constant $\bar{w}$), $\mathbf{A}$ has an explicit inverse. The solution is given by a linear filtering:

$$\mathbf{u} = \mathbf{A}^{-1}\mathbf{b} = \mathbf{F}^* \bar{w}\big(\bar{w}\mathbf{I} + \gamma^{2\alpha}\tilde{\Lambda}\big)^{-1} \mathbf{F}\mathbf{u}_0 \tag{11}$$

In plain words, Eq. (11) signifies:

(i) take the Fourier transform of $\mathbf{u}_0$;
(ii) weight the spectrum by $S_k := \bar{w}/(\bar{w} + \gamma^{2\alpha}\tilde{\lambda}_k)$ in a pointwise manner;
(iii) take the inverse Fourier transform.

The weights S_k correspond to the spectrum of a low-pass filter: S_k attains its maximum at the zero frequency ($k = 0$) and starts to drop down as k increases. It attains the half of the maximum at the frequency k such that $\tilde{\lambda}_k = \bar{w}/\gamma^{2\alpha}$. Examples of the spectrum for different α are shown in Fig. 1.

2.3. Case of non-constant weights: interpretation of a controlled diffusion

Regarding the case of non-constant weights, $\mathbf{A}$ no longer has an explicit inverse in general. However some asymptotic analysis sheds light on the expected behavior of the solution.

Suppose that the weights are nowhere zero, $\alpha = 1$ and a sufficiently small γ such that one can consider the first-order approximation of the solution:

$$\begin{aligned} \mathbf{u} = \mathbf{A}^{-1}\mathbf{b} &= \big(\mathbf{W} + \gamma^2 \mathbf{L}_1^*\mathbf{L}_1\big)^{-1}\mathbf{W}\mathbf{u}_0 \\ &= \big(\mathbf{I} + \gamma^2 \mathbf{W}^{-1}\mathbf{L}_1^*\mathbf{L}_1\big)^{-1}\mathbf{u}_0 \\ &\approx \big(\mathbf{I} - \gamma^2 \mathbf{W}^{-1}\mathbf{L}_1^*\mathbf{L}_1\big)\mathbf{u}_0 = \mathbf{u}_0 - \gamma^2 \mathbf{W}^{-1}\mathbf{L}_1^*\mathbf{L}_1\mathbf{u}_0 \end{aligned} \tag{12}$$

It can be seen that Eq. (12) represents a step of diffusion on $\mathbf{u}_0$ where the step length is controlled by $\gamma^2\mathbf{W}^{-1}$. Clearly, a data point associated with a large weight has a small step length and will undergo little change, while a point with a small weight (or large step) will tend to be blurred out by the diffusion process.