Electronic Journal of Linear Algebra, ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 32, pp. 163-171, May 2017.

SHRINKAGE FUNCTION AND ITS APPLICATIONS IN MATRIX APPROXIMATION*

TOBY BOAS†, ARITRA DUTTA‡, XIN LI‡, KATHRYN P. MERCIER§, AND ERIC NIDERMAN‡

Abstract. The shrinkage function is widely used in matrix low-rank approximation, compressive sensing, and statistical estimation. In this article, an elementary derivation of the shrinkage function is given. In addition, applications of the shrinkage function are demonstrated in solving several well-known problems, together with a new result in matrix approximation.

Key words. Shrinkage function, Singular value decomposition, Low-rank approximation, Sparse approximation.

AMS subject classifications. 65F15, 65F30, 65F35, 65F50, 65K10.

1. Introduction. Historically some important mathematical functions have been introduced for convenience. For example, the Heaviside step function $H(\cdot)$, a piecewise constant function given by:

$$H(x) = \begin{cases} 1, & x > 0 \\ \frac{1}{2}, & x = 0 \\ 0, & x < 0 \end{cases},$$

and the Dirac delta function $\delta(\cdot)$ (more precisely, a distribution (see, e.g., [1])), a generalized function whose discrete analog is referred to as the Kronecker delta function:

$$\delta_{ij} = \begin{cases} 1, & i = j \\ 0, & i \neq j \end{cases}.$$

This article concerns a newcomer, the shrinkage function $S_\lambda(\cdot)$, first introduced by Donoho and Johnstone in their landmark paper ([7], see also [18]) on function estimation using wavelets in the early 1990's. Recently, the shrinkage function has been heavily used in the solutions of several optimization and approximation problems of matrices (see, e.g., [3, 11, 17, 21]). We give an elementary treatment that is accessible to a vast group of researchers, as it only requires basic knowledge in calculus and linear algebra and show how naturally the shrinkage function can be used in solving more advanced problems.

2. A calculus problem. We start with a simple calculus problem. Let $\lambda > 0$ and $a \in \mathbb{R}$ be given. Consider the following problem:

$$\min_{x \in \mathbb{R}} \left[\lambda |x| + \frac{1}{2}(x-a)^2 \right].$$

*Received by the editors on January 27, 2016. Accepted for publication on March 20, 2017. Handling Editor: Bryan L. Shader.

†Department of Mathematics, University of Florida, Gainesville, FL 32611-8105 (tboas@ufl.edu).

‡Department of Mathematics, University of Central Florida, 4000 Central Florida Blvd, Orlando, Florida 32816 (d.aritra2010@knights.ucf.edu, xli@ucf.edu, eaniederman@knights.ucf.edu). This work is partially supported by a CSUMS program of the National Science Foundation, grant no. DMS-0803059.

§Department of Biology, University of Central Florida, 4000 Central Florida Blvd, Orlando, Florida 32816(katie.mercier@knights.ucf.edu).

Electronic Journal of Linear Algebra, ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 32, pp. 163-171, May 2017.

We adopt the notation $a = \arg\min_{x\in A} f(x)$ to mean that $a \in A$ is a solution of the minimization problem $\min_{x\in A} f(x)$ and define:

$$S_\lambda(a) := \arg\min_{x\in\mathbb{R}} \left[\lambda|x| + \frac{1}{2}(x-a)^2\right]. \tag{2.1}$$

THEOREM 2.1. *Let $\lambda > 0$ be fixed. For each $a \in \mathbb{R}$, there is one and only one solution $S_\lambda(a)$, to the minimization problem (2.1). Furthermore,*

$$S_\lambda(a) = \begin{cases} a-\lambda, & a > \lambda \\ 0, & |a| \le \lambda \\ a+\lambda, & a < -\lambda \end{cases}.$$

REMARK 2.2. The function $S_\lambda(\cdot)$ defined above is called the shrinkage function (also referred to as shrinkage or soft threshold, [7, 18]). One may imagine that $S_\lambda(a)$ "shrinks" a to zero when $|a| \le \lambda$. A plot of $S_\lambda(\cdot)$ for $\lambda = 1$ is given in Fig. 1.

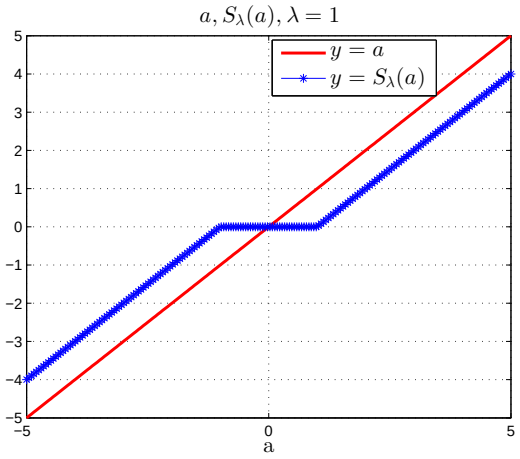

FIGURE 1. *A plot of S_λ for $\lambda = 1$.*

Proof of Theorem 2.1. Let $f(x) = \lambda|x| + \frac{1}{2}(x-a)^2$. Note that $f(x) \to \infty$ when $|x| \to \infty$ and f is continuous on $\mathbb{R}$ and differentiable everywhere except a single point $x = 0$. So, f achieves its minimum value on $\mathbb{R}$ at one of its critical points. A plot of f for different values of a and $\lambda = 1$ is given in Fig. 2. Let $x^* = \arg\min_{x\in\mathbb{R}} f(x)$.

We consider three cases.

Case 1: $x^* > 0$. Since f is differentiable at $x = x^*$ and achieves its minimum, we must have $f'(x^*) = 0$. Note that, for $x > 0$, we have

$$f'(x) = \frac{d}{dx}\left[\lambda x + \frac{1}{2}(x-a)^2\right] = \lambda + (x-a).$$

So,

$$\lambda + (x^* - a) = 0,$$

Electronic Journal of Linear Algebra, ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 32, pp. 163-171, May 2017.

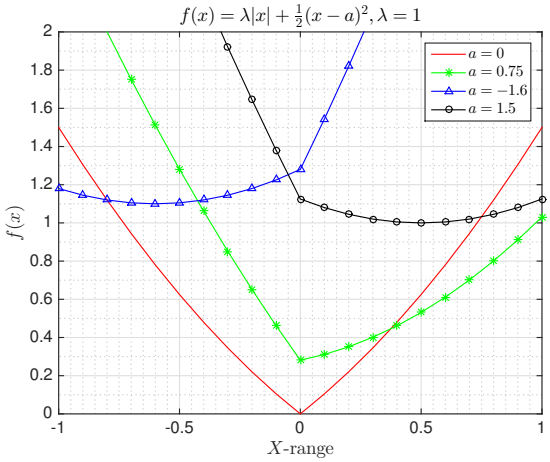

FIGURE 2. *Plots of $f(x)$ for different values of a with $\lambda = 1$.*

which implies

$$x^* = a - \lambda.$$

To be consistent with $x^* > 0$, it is necessary that $a - \lambda > 0$ or, equivalently, $a > \lambda$.

Case 2: $x^* < 0$. By proceeding similarly as in Case 1 above, we can arrive at

$$x^* = a + \lambda \ \text{ with } \ a < -\lambda.$$

Case 3: $x^* = 0$. Note that $f(x)$ is no longer differentiable at $x^* = 0$ (so we could not use the condition $f'(x^*) = 0$ as before). But since f has a minimum at $x^* = 0$ and since f is differentiable on each side of $x^* = 0$, it is necessary that

$$f'(x) > 0 \ \text{ for } \ x > 0 \ \text{ and } \ f'(x) < 0 \ \text{for } \ x < 0.$$

So,

$$\lambda + x - a > 0 \ \text{ for } x > 0 \ \text{ and } \ -\lambda + x - a < 0 \ \text{ for } x < 0.$$

Thus,

$$\lambda - a > 0 \ \text{ and } \ -\lambda - a < 0,$$

or, equivalently,

$$|a| \leq \lambda.$$

To summarize, we have

$$x^* = \begin{cases} a - \lambda & \text{with } a > \lambda \\ a + \lambda & \text{with } a < -\lambda \\ 0 & \text{with } |a| \leq \lambda \end{cases}.$$

Since one and only one of the three cases 1) $a > \lambda$, 2) $a < -\lambda$, and 3) $|a| \leq \lambda$ holds, we obtain the uniqueness. With the uniqueness, it is straightforward to verify that each of the three cases would imply the corresponding formula for x^*. ∎

Electronic Journal of Linear Algebra, ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 32, pp. 163-171, May 2017.

3. A sparse recovery problem. Recently, research in compressive sensing leads to the recognition of the fact that for many optimization problems, the ℓ_1-norm of a vector is a good substitute for the count of the number of non-zero entries of the vector in many minimization problems. In this section, we solve some simple minimization problems using the count of non-zero entries or the ℓ_1-norm. Given a vector $\mathbf{v} \in \mathbb{R}^n$, we want to solve

$$\min_{\mathbf{u}\in\mathbb{R}^n} \left[\|\mathbf{u}\|_{\ell_0} + \frac{\beta}{2}\|\mathbf{u}-\mathbf{v}\|_{\ell_2}^2 \right], \tag{3.2}$$

where $\|\cdot\|_{\ell_0}$ denotes the number of non-zero entries of $\mathbf{u}$, such that, $\|\mathbf{u}\|_{\ell_0} = \#\{i|u_i \neq 0\}$, the cardinality of $\mathbf{u}$, $\|\cdot\|_{\ell_2}$ denotes the Euclidean norm in $\mathbb{R}^n$, and $\beta > 0$ is a given balancing parameter. We can solve problem (3.2) component-wise (in each u_i) as follows. Notice that, given $\mathbf{u} \in \mathbb{R}^n$, each entry u_i of $\mathbf{u}$ contributes 1 to $\|\mathbf{u}\|_{\ell_0}$ if u_i is non-zero, and contributes 0 if u_i is zero. Since we are minimizing $g(\mathbf{u}) := \|\mathbf{u}\|_{\ell_0} + \frac{\beta}{2}\|\mathbf{u}-\mathbf{v}\|_{\ell_2}^2$, if u_i is zero then the contribution to $g(\mathbf{u})$ depending on this u_i is $\frac{\beta}{2}v_i^2$; otherwise, if u_i is non-zero, then we should minimize $\frac{\beta}{2}(u_i - v_i)^2$ for $u_i \in \mathbb{R}\setminus\{0\}$, which forces that $u_i = v_i$ and contributes 1 to $g(\mathbf{u})$ as the minimum value. Therefore, the solution $\mathbf{u}$ to problem (3.2) is given component-wise by

$$u_i = \begin{cases} 0, & \text{if } \frac{\beta}{2}(v_i)^2 \leq 1 \\ v_i, & \text{otherwise} \end{cases}.$$

Next, we replace $\|\mathbf{u}\|_{\ell_0}$ by $\|\mathbf{u}\|_{\ell_1}$ in (3.2) and solve:

$$\min_{\mathbf{u}\in\mathbb{R}^n} \left[\|\mathbf{u}\|_{\ell_1} + \frac{\beta}{2}\|\mathbf{u}-\mathbf{v}\|_{\ell_2}^2 \right], \tag{3.3}$$

where $\|\cdot\|_{\ell_1}$ denotes the ℓ_1 norm in $\mathbb{R}^n$.

Using Theorem 2.1, we can solve (3.3) component-wise as follows.

THEOREM 3.1. [10] *Let $\beta > 0$ and $\mathbf{v} \in \mathbb{R}^n$ be given and let*

$$\mathbf{u}^* = \arg\min_{\mathbf{u}\in\mathbb{R}^n} \left[\|\mathbf{u}\|_{\ell_1} + \frac{\beta}{2}\|\mathbf{u}-\mathbf{v}\|_{\ell_2}^2 \right].$$

Then

$$\mathbf{u}^* = S_{1/\beta}(\mathbf{v}),$$

where, $S_{1/\beta}(\mathbf{v})$ denotes the vector whose entries are obtained by applying the shrinkage function $S_{1/\beta}(\cdot)$ to the corresponding entries of $\mathbf{v}$.

Proof. If u_i and v_i denote the ith entry of the vectors $\mathbf{u}$ and $\mathbf{v}$, respectively, $i = 1, 2, \ldots, n$, then we have,

$$\begin{aligned} \mathbf{u}^* &= \arg\min_{\mathbf{u}\in\mathbb{R}^n} \left[\|\mathbf{u}\|_{\ell_1} + \frac{\beta}{2}\|\mathbf{u}-\mathbf{v}\|_{\ell_2}^2 \right] \\ &= \arg\min_{\mathbf{u}\in\mathbb{R}^n} \left[\sum_{i=1}^{n} \left(\frac{1}{\beta}|u_i| + \frac{1}{2}(u_i - v_i)^2 \right) \right]. \end{aligned}$$

Since the i-th term in the summation depends only on u_i, the vector $\mathbf{u}^*$ must have components u_i^* satisfying

$$u_i^* = \arg\min_{u_i^*\in\mathbb{R}} \left[\frac{1}{\beta}|u_i| + \frac{1}{2}(u_i - v_i)^2 \right],$$

for $i = 1, 2, \ldots, n$. But by Theorem 2.1, the solution to each of these problems is given precisely by $S_{1/\beta}(v_i)$. This yields the result. ∎

Electronic Journal of Linear Algebra, ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 32, pp. 163-171, May 2017.

REMARK 3.2. The previous proof still works if we replace the vectors by matrices and use the extension of the ℓ_1 and ℓ_2-norms to matrices *by treating them as vectors*. Thus by using the same argument we can easily show the following matrix version of the previous theorem.

THEOREM 3.3. [10] *Let $\beta > 0$ and $\mathbf{V} \in \mathbb{R}^{m\times n}$ be given. Then*

$$S_{1/\beta}(\mathbf{V}) = \arg\min_{\mathbf{U}\in\mathbb{R}^{m\times n}} \left[\|\mathbf{U}\|_{\ell_1} + \frac{\beta}{2}\|\mathbf{U}-\mathbf{V}\|_{\ell_2}^2 \right],$$

where $S_{1/\beta}(\mathbf{V})$ is defined component-wise.

Theorem 3.3 solves the problem of approximating a given matrix by a sparse matrix by using the shrinkage function.

4. Approximation by low rank matrices. The sparse approximation as given by Theorem 3.3 has many applications such as data compression and dimension reduction [2, 9]. In these areas, one may also be interested in finding matrices of low rank (see, e.g., [8, 11, 17, 21]). For example, given a matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$, we want to solve the following approximation problem:

$$\min_{\mathbf{X}\in\mathbb{R}^{m\times n}} \left[\operatorname{rank}(\mathbf{X}) + \frac{\beta}{2}\|\mathbf{X}-\mathbf{A}\|_F^2 \right],$$

where $\|\cdot\|_F$ denotes the Frobenius norm of matrices (which turns out to be equivalent to the vector l_2 norm if we treat a matrix as a vector - see more discussion on the matrix norms in Subsection 4.1).

This is a harder problem since $\operatorname{rank}(X)$ is not a convex function. A convex relaxation (see, e.g., [9]) of the problem is provided by replacing the term $\operatorname{rank}(\mathbf{X})$ by the nuclear norm of $\mathbf{X}$, $\|\mathbf{X}\|_*$, (again, see subsection 4.1 for a discussion on the nuclear norm and its properties). The problem then becomes:

$$\min_{\mathbf{X}\in\mathbb{R}^{m\times n}} \left[\|\mathbf{X}\|_* + \frac{\beta}{2}\|\mathbf{X}-\mathbf{A}\|_F^2 \right]. \tag{4.4}$$

This problem again yields an explicit solution ([3, 12]), but in this literature, the formula is derived by using advanced tools from convex analysis ("subdifferentials" to be more specific) [20]. Here, we will show how we can obtain the solution by using simple ideas from the previous section.

4.1. Singular value decomposition and matrix norms. It will be beneficial to recall the various matrix norms. Many useful matrix norms can be defined in terms of the singular values of the matrices. We will deal with two of them: the nuclear norm $\|\cdot\|_*$ and the Frobenius norm $\|\cdot\|_F$.

Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ and $\mathbf{A} = \mathbf{U}\tilde{\mathbf{A}}\mathbf{V}^T$ be a singular value decomposition (SVD) of $\mathbf{A}$ with $\mathbf{U} \in \mathbb{R}^{m\times m}$ and $\mathbf{V} \in \mathbb{R}^{n\times n}$ being two orthogonal matrices (that is, $\mathbf{U}^{-1} = \mathbf{U}^T$ and $\mathbf{V}^{-1} = \mathbf{V}^T$) and $\tilde{\mathbf{A}} = \operatorname{diag}(\sigma_1(\mathbf{A}),\ \sigma_2(\mathbf{A}), \ldots, \sigma_{\min\{m,n\}}(\mathbf{A}))$ being a $m \times n$ non-square diagonal matrix having $\sigma_1(\mathbf{A}) \geq \sigma_2(\mathbf{A}) \geq \cdots \geq \sigma_{\min\{m,n\}}(\mathbf{A}) \geq 0$ on its diagonal and 0's elsewhere. The $\sigma_i(\mathbf{A})$'s are called the singular values of $\mathbf{A}$. It is known that every matrix in $\mathbb{R}^{m\times n}$ has a SVD and that SVD of a matrix is not necessarily unique [16]. Then the nuclear norm of $\mathbf{A}$ is defined as

$$\|\mathbf{A}\|_* = \sum_{i=1}^{\min\{m,n\}} \sigma_i(\mathbf{A}),$$

Electronic Journal of Linear Algebra, ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 32, pp. 163-171, May 2017.

and the Frobenius norm of $\mathbf{A}$ as

$$\|\mathbf{A}\|_F = \left(\sum_{i=1}^{\min\{m,n\}} (\sigma_i(\mathbf{A}))^2 \right)^{1/2}.$$

This norm turns out to be the same as the ℓ_2 norm of $\mathbf{A}$, treated as a vector in $\mathbb{R}^{mn\times 1}$. This is because the nonzero singular values $\sigma_i(\mathbf{A})$'s are exactly the square root of the nonzero eigenvalues of $\mathbf{A}\mathbf{A}^T$ or $\mathbf{A}^T\mathbf{A}$. So,

$$\|\mathbf{A}\|_{l_2}^2 = \sum_{i=1}^{m}\sum_{j=1}^{n}(a_{ij})^2 = \operatorname{trace}(\mathbf{A}\mathbf{A}^T) = \sum_{i=1}^{\min\{m,n\}} (\sigma_i(\mathbf{A}))^2.$$

Here we have used trace$(\cdot)$ to denote the trace of a matrix (which is equal to the sum of all diagonal entries of the matrix).

We will need the following simple and well-known fact about the nuclear norms of a matrix and that of its diagonal [15]: Let $D(\mathbf{A})$ denote the diagonal matrix using the diagonal of $\mathbf{A}$ then:

$$\|D(\mathbf{A})\|_* \le \|\mathbf{A}\|_*. \tag{4.5}$$

This inequality can be verified by using a SVD of $\mathbf{A} = \mathbf{U}\tilde{\mathbf{A}}\mathbf{V}^T$ as follows. For completeness we will provide a simple proof of (4.5). Write $\mathbf{U} = (u_{ij})$, $\mathbf{V} = (v_{ij})$, and $t = \min\{m, n\}$. Then

$$\|D(\mathbf{A})\|_* = \|D(\mathbf{U}\tilde{\mathbf{A}}\mathbf{V}^T)\|_* = \sum_{i=1}^{t}\left|\sum_{j=1}^{t}\sigma_j(\mathbf{A})u_{ij}v_{ij}\right| \le \sum_{j=1}^{t}\sigma_j(\mathbf{A})\sum_{i=1}^{t}|u_{ij}v_{ij}|$$

$$\le \sum_{j=1}^{t}\sigma_j(\mathbf{A})\cdot\left(\sum_{i=1}^{t}|u_{ij}|^2\right)^{1/2}\left(\sum_{i=1}^{t}|v_{ij}|^2\right)^{1/2} \le \sum_{j=1}^{t}\sigma_j(\mathbf{A}) = \|\mathbf{A}\|_*,$$

where we have used the Cauchy-Schwarz inequality in obtaining the second inequality, and the orthogonality of $\mathbf{U}$ and $\mathbf{V}$ (so that $\sum_{i=1}^{t}|u_{ij}|^2 \le 1$ and $\sum_{i=1}^{t}|v_{ij}|^2 \le 1$) in the last inequality.

We will also use the fact that for any orthogonal matrices $\mathbf{L}$ and $\mathbf{R}$, $\mathbf{LAR}$ and $\mathbf{A}$ have the same singular values, and therefore their Frobenius norms and nuclear norms are same:

$$\|\mathbf{LAR}\|_F = \|\mathbf{A}\|_F \text{ and } \|\mathbf{LAR}\|_* = \|\mathbf{A}\|_*.$$

This is known as the unitary invariance of the Frobenius norm and nuclear norm.

4.2. Solution to (4.4) via problem (2.1). We are ready to show how problem (4.4) is problem (2.1) in disguise. Let $\mathbf{U}\tilde{\mathbf{A}}\mathbf{V}^T$ be a SVD of $\mathbf{A}$. Given $\beta > 0$, using the unitary invariance of the Frobenius norm and the nuclear norm, we have

$$\min_{\mathbf{X}}\left[\|\mathbf{X}\|_* + \frac{\beta}{2}\|\mathbf{X}-\mathbf{A}\|_F^2\right] = \min_{\mathbf{X}}\left[\|\mathbf{X}\|_* + \frac{\beta}{2}\|\mathbf{X}-\mathbf{U}\tilde{\mathbf{A}}\mathbf{V}^T\|_F^2\right]$$
$$= \min_{\mathbf{X}}\left[\|\mathbf{U}^T\mathbf{X}\mathbf{V}\|_* + \frac{\beta}{2}\|\mathbf{U}^T\mathbf{X}\mathbf{V}-\tilde{\mathbf{A}}\|_F^2\right].$$

It can be seen from the last expression that the minimum occurs when $\tilde{\mathbf{X}} := \mathbf{U}^T\mathbf{X}\mathbf{V}$ is diagonal since both terms in that expression get no larger when $\tilde{\mathbf{X}}$ is replaced by its diagonal matrix (with the help of (4.5)). So, the matrix $\mathbf{E} = (e_{ij}) := \tilde{\mathbf{X}} - \tilde{\mathbf{A}}$ is a diagonal matrix. Thus,

$$\mathbf{X} = \mathbf{U}\tilde{\mathbf{X}}\mathbf{V}^T, \text{ with } \tilde{\mathbf{X}} = \tilde{\mathbf{A}} + \mathbf{E},$$

Electronic Journal of Linear Algebra, ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 32, pp. 163-171, May 2017.

which yields a SVD of $\mathbf{X}$ (using the same matrices $\mathbf{U}$ and $\mathbf{V}$ as in a SVD of $\mathbf{A}$!). Then,

$$\begin{aligned}\min_{\mathbf{X}}\left[\|\mathbf{X}\|_*+\frac{\beta}{2}\|\mathbf{X}-\mathbf{A}\|_F^2\right] &= \min_{\tilde{\mathbf{X}}\in\text{diag}}\left[\|\tilde{\mathbf{X}}\|_*+\frac{\beta}{2}\|\tilde{\mathbf{X}}-\tilde{\mathbf{A}}\|_F^2\right]\\ &= \min_{\tilde{\mathbf{X}}\in\text{diag}}\left[\sum_i|\tilde{x}_{ii}|+\frac{\beta}{2}\sum_i(\tilde{x}_{ii}-\sigma_i(\mathbf{A}))^2\right],\end{aligned} \tag{4.6}$$

where "diag" is the set of diagonal matrices in $\mathbb{R}^{m\times n}$. Note that (4.6) is an optimization problem similar to (2.1) (for vectors $(\tilde{x}_{11},\tilde{x}_{22},\cdots)^T$ as $\tilde{\mathbf{X}}$ varies) whose solution is given by

$$\tilde{x}_{ii}=S_{1/\beta}(\sigma_i(\mathbf{A})),\ i=1,2,\cdots.$$

To summarize, we have proven the following.

THEOREM 4.1. [3] *Suppose that $\mathbf{A}\in\mathbb{R}^{m\times n}$ and $\beta>0$ are given. Then the solution to the minimization problem (4.4) is given by $\hat{\mathbf{X}}=\mathbf{U}\tilde{\mathbf{X}}\mathbf{V}^T$ where the diagonal matrix $\tilde{\mathbf{X}}$ has diagonal entries*

$$\tilde{x}_{ii}=S_{1/\beta}(\sigma_i(\mathbf{A})),\ i=1,2,\ldots,\min\{m,n\},$$

where $\mathbf{U}\tilde{\mathbf{A}}\mathbf{V}^T$ is a SVD of $\mathbf{A}$.

REMARK 4.2. **1.** A recent proof of this theorem is given by Cai, Candes, and Shen in [3] where they give an advanced verification of the result. Our proof given above has the advantage that it is elementary and allows the reader to "discover" the result.
2. There are many earlier discoveries of related results ([14, 16]) where $\text{rank}(\mathbf{X})$ is used instead of the nuclear norm $\|\mathbf{X}\|_*$. We will examine one such variant in the next section.
3. A closely related (but harder) problem is *compressive sensing* ([4, 5]). In a more general set up where only a subset of the entries of the data matrix is observable, for example, matrix completion problem under low-rank penalties [3, 4, 12]:

$$\min_{\mathbf{X}}\text{rank}(\mathbf{X})\quad\text{subject to}\quad A_{ij}=X_{ij},\ (i,j)\in\Omega,$$

where $\Omega\subseteq\{(i,j):1\le i\le m,1\le j\le n\}$, is indeed NP-hard [6]. Note that, the matrix completion problem is a special case of the affinely constrained matrix rank minimization problem [12]:

$$\min_{\mathbf{X}}\text{rank}(\mathbf{X})\quad\text{subject to}\quad A(\mathbf{X})=\mathbf{b},$$

where $\mathbf{X}\in\mathbb{R}^{m\times n}$ is the decision variable and $\mathbf{A}:\mathbb{R}^{m\times n}\to\mathbb{R}^p$ is a linear map. One common idea used in such a situation is to consider a convex relaxation of the above problem by replacing the rank by its convex surrogate, the nuclear norm. Readers are strongly recommended to the recent survey by Bryan and Leise ([2]).

5. A variation. More problems can be solved by applying similar ideas. For example, let us consider a variant of a well-known result of Schmidt (see, e.g., [14, Section 5]), replacing the rank by the nuclear norm: For a fixed positive number τ, consider

$$\min_{\mathbf{X}\in\mathbb{R}^{m\times n}}\|\mathbf{X}-\mathbf{A}\|_F\quad\text{subject to}\quad\|\mathbf{X}\|_*\le\tau. \tag{5.7}$$

Electronic Journal of Linear Algebra, ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 32, pp. 163-171, May 2017.

Using similar methods as in Section 3, this problem can be transformed (see the derivation below) into the following:

$$\min_{\mathbf{u}\in\mathbb{R}^{\min\{m,n\}}} \|\mathbf{u}-\mathbf{v}\|_{\ell_2} \quad \text{subject to} \quad \|\mathbf{u}\|_{\ell_1} \le \tau. \tag{5.8}$$

The LASSO (Least absolute shrinkage and selection operator) is a vastly used regression technique in data mining and statistics [13, 18, 19]. It follows a simple model for variable selection and regularization. Let the predictor variables, $\mathbf{X} \in \mathbb{R}^{N\times p}$ and the responses, $y_i, i = 1, 2, \ldots, N$ are given. Assuming that $\sum_i x_{ij}/N = 0$ and $\sum_i x_{ij}^2/N = 1$, the LASSO estimate is given by:

$$\hat{\beta} = \arg\min_{\beta} \|\mathbf{y} - \beta_{\mathbf{0}} - \mathbf{X}\beta\|_{\ell_2}^2 \quad \text{subject to} \quad \|\beta\|_{\ell_1} \le \tau, \tag{5.9}$$

where $\beta_{\mathbf{0}}$ is the mean of the response vector $\mathbf{y}$. Note that, (5.8) shares some similarity with (5.9) [13, 18, 19]. As in [18], one can form a *Lagrangian* of (5.8) and solve:

$$\mathbf{u}^* = \arg\min_{\mathbf{u}\in\mathbb{R}^{\min\{m,n\}}} \left\{ \frac{1}{2}\|\mathbf{u}-\mathbf{v}\|_{\ell_2}^2 + \lambda\|\mathbf{u}\|_{\ell_1} \right\}, \text{ with } \|S_\lambda(\mathbf{v})\|_{\ell_1} = \tau, \tag{5.10}$$

which has a solution $\mathbf{u}^* = S_\lambda(\mathbf{v})$ according to Theorem 3.1. (The reason for us to use λ instead of β in (5.8) is nonessential. It is only for indicating the similarity with the LASSO formulation.) We now sketch the derivation of converting (5.7) to (5.8): As before, let $\mathbf{A} = \mathbf{U}\tilde{\mathbf{A}}\mathbf{V}^T$ be a SVD of $\mathbf{A}$. Then,

$$\min_{\mathbf{X}\in\mathbb{R}^{m\times n}} \|\mathbf{X}-\mathbf{A}\|_F = \min_{\mathbf{X}\in\mathbb{R}^{m\times n}} \|\mathbf{U^T X V} - \tilde{\mathbf{A}}\|_F.$$

Note that $\|\mathbf{X}\|_* = \|\mathbf{U}^T\mathbf{X}\mathbf{V}\|_*$, so (5.7) can be written as

$$\min_{\mathbf{X}\in\mathbb{R}^{m\times n}} \|\mathbf{X}-\tilde{\mathbf{A}}\|_F \quad \text{subject to} \quad \|\mathbf{X}\|_* \le \tau,$$

which, by using (4.5), can be further transformed to

$$\min_{\mathbf{X}\in\mathbb{R}^{m\times n}} \|\mathbf{X}-\tilde{\mathbf{A}}\|_F \quad \text{subject to} \quad \mathbf{X} \text{ being diagonal and } \|\mathbf{X}\|_* \le \tau. \tag{5.11}$$

Next, if we let $\mathbf{u}$ and $\mathbf{v}$ be two vectors in $\mathbb{R}^{\min\{m,n\}}$ consisting of the diagonal elements of $\mathbf{X}$ and $\tilde{\mathbf{A}}$, respectively, then (5.11) is equivalent to (5.8). Since $S_{\lambda_1}(\mathbf{v})$ solves (5.10) we have,

$$\frac{1}{2}\|S_{\lambda_1}(\mathbf{v})-\mathbf{v}\|_2^2 + \lambda_1\tau \le \frac{1}{2}\|\mathbf{u}-\mathbf{v}\|_{\ell_2}^2 + \lambda_1\|\mathbf{u}\|_{\ell_1},$$

for all $\mathbf{u} \in \mathbb{R}^{\min\{m,n\}}$. Which implies, for all $\mathbf{u} \in \mathbb{R}^{\min\{m,n\}}$,

$$\frac{1}{2}\|S_{\lambda_1}(\mathbf{v})-\mathbf{v}\|_2^2 \le \frac{1}{2}\|\mathbf{u}-\mathbf{v}\|_{\ell_2}^2 + \lambda_1(\|\mathbf{u}\|_{\ell_1} - \tau).$$

Therefore,

$$\frac{1}{2}\|S_{\lambda_1}(\mathbf{v})-\mathbf{v}\|_2^2 \le \left\{ \frac{1}{2}\|\mathbf{u}-\mathbf{v}\|_{\ell_2}^2 \right\},$$

for all $\mathbf{u} \in \mathbb{R}^{\min\{m,n\}}$, such that $\|\mathbf{u}\|_{\ell_1} \le \tau$. Hence $\mathbf{u}^* = S_\lambda(\mathbf{v})$ solves (5.8). Thus we have established the following result.

Electronic Journal of Linear Algebra, ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 32, pp. 163-171, May 2017.

THEOREM 5.1. [8] *With the notations above, the solution to problem* (5.7) *is given by*

$$\hat{\mathbf{X}} = \mathbf{U}S_\lambda(\tilde{\mathbf{A}})\mathbf{V}^T,$$

for some λ *such that* $\|S_\lambda(\tilde{\mathbf{A}})\|_{\ell_1} = \tau$.

Acknowledgment. Toby Boas, Kathryn P. Mercier, and Eric Niederman contributed to this work as undergraduate students.

REFERENCES

[1] R. Bracewell. The Impulse Symbol (Section 5). In: *The Fourier Transform and Its Applications*, third edition. McGraw-Hill Book Co., New York, 74–104, 1986.
[2] K. Bryan and T. Leise. Making do with less: an introduction to compressed sensing. *SIAM Review*, 55:547–566, 2013.
[3] J. Cai, E.J. Candés, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20:1956–1982, 2010.
[4] E.J. Candés and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56:2053–2080, 2010.
[5] E.J. Candés, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52:489–509, 2006.
[6] A.L. Chistov and D.Yu. Grigor'ev. Complexity of quantifier elimination in the theory of algebraically closed fields. In: M.P. Chytil and V. Koubek (editors), *Mathematical Foundations of Computer Science, Lecture Notes in Computer Science*, 11th Symposium Proceedings, Springer-Verlag, Berline, 176:17–31, 1984
[7] D.L. Donoho and I.M. Johnstone. Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81:425–455, 1994.
[8] A. Dutta. *Weighted Low-Rank Approximation of Matrices: Some Analytical and Numerical Aspects*. Ph.D. Dissertation, Department of Mathematics, University of Central Florida, 2016.
[9] M. Fazel. *Matrix Rank Minimization with Applications*. Ph.D. Dissertation, Department of Electrical Engineering, Stanford University, 2002.
[10] E.T. Hale, W. Yin, and Y. Zhang. Fixed-point continuation for l_1-minimization: methodology and convergence *SIAM Journal on Optimization*, 19:1107–1130, 2008.
[11] Z. Lin, M. Chen, and Y. Ma. The augmented Lagrange multiplier method for exact recovery of corrupted low-rank matrices. Preprint, arXiv:1009.5055v2, 2011.
[12] S. Ma, D. Goldfarb, and L. Chen. Fixed point and Bergman iterative methods for matrix rank minimization. *Mathematical Programming, Series A*, 128:321–353 2011.
[13] M.R. Osborne, B. Presnell, and B.A. Turlach. On the LASSO and its dual. *Journal of Computational and Graphical Statistics*, 9:319–337, 2000.
[14] G.W. Stewart. On the early history of the singular value decomposition. *SIAM Review*, 35:551–566, 1993.
[15] G.W. Stewart and J. Sun. *Matrix Perturbation Theory*. Computer Science and Scientific Computing, Academic Press, Inc., Boston, 1990.
[16] G.W. Strang. *Introduction to Linear Algebra*, third edition. Wellesley-Cambridge Press, Wellesley, 2003.
[17] M. Tao and X. Yuan. Recovering low-rank and sparse components of matrices from incomplete and noisy observations. *SIAM Journal on Optimization*, 21:57–81, 2011.
[18] R. Tibshirani. Regression shrinkage and selection via the LASSO. *Journal of the Royal Statistical Society. Series B. Methodological*, 58:267–288, 1996.
[19] R.J. Tibshirani and J. Taylor. The solution path of the generalized LASSO. *The Annals of Statistics*, 39:1335–1371, 2011.
[20] G.A. Watson. Characterization of the subdifferential of some matrix norms. *Linear Algebra and its Applications*, 170:33–45, 1992.
[21] X. Yuan and J. Yang. Sparse and low-rank matrix decomposition via alternating direction method. Preprint, available at http://www.optimization-online.org/DB_HTML/2009/11/2447.html, 2009.