Image Inpainting via Low-Rank Matrix Completion

Jared Spang and Marisa Lee

Abstract

We consider the problem of image reconstruction across a contaminated domain. This problem can be reformulated as the construction of a full data matrix given a partial matrix. Is it possible to reconstruct the missing elements of a matrix? Naturally, this is an under-constrained problem so we must make some sort of assumption about the result in order to apply any sort of technique for solving it. In this paper we make the assumption that the matrix we are attempting to complete is of low rank such that minimizing the rank of our matrix will yield a proper reconstruction of the missing data. Unfortunately, images are not naturally sparse so we sparsify the image through an application of the wavelet transform to the original image. After this we are left with the convex optimization problem of low rank matrix completion which we solve using a semi-supervised learning approach described in section 2.

1 Problems

The problem discussed in this paper is that of image inpainting; reconstruction of a image given a contaminated domain Ω^c in the initial image. We formulate the problem as follows [3]

$$\min_u R(u) \text{ , s.t. } u(x) = f(x), \forall x \in \Omega$$

Where $f : \mathbb{R}^2 \mapsto \mathbb{R}$ is the contaminated image. We define $R(u)$ as the wavelet transform applied to the original image x (see section 2.1). We also define Ω as the "clean" image domain. The method described in this paper to solve the problem of image reconstruction utilizes a technique known as low-rank matrix completion which we define below [1]:

$$\min_X \ rank(X) \text{ , s.t. } X_{ij} = M_{ij}, \forall (i,j) \in \Omega$$

Where M is the target fully completed matrix. Unfortunately, this problem is notoriously difficult to solve computationally (NP-Hard); however, it has been found that a close approximation to rank minimization is to solve the convex problem of minimizing the l_1 norm (in the case of a vector) or the l_* nuclear norm (in the case of a matrix). The intuition behind this stems from the fact that a matrix of rank n should have exactly n non-zero singular values. Thus, minimizing the nuclear norm (the sum of a matrix's singular values) results in close approximation to rank minimization [1]. We formally define the new problem below:

$$\min_X \ ||X||_* \text{ , s.t. } X_{ij} = M_{ij}, \forall (i,j) \in \Omega$$

Fortunately, there is a computationally efficient way to solve this new problem as we will discuss in the methodology. The only remaining problem we face is that this method is only a proper data reconstruction technique for sparse matrices but a pixel image array does not typically form a sparse array. The solution we chose to solve this problem was to sparsify the given pixel image array under some sort of transform for which, in this paper, we chose to use the wavelet transform. Utilizing the wavelet transform we can decompose a given image (x) into various sparse frequency sub-bands. From this we get our final minimization problem which is tackled by this paper:

$$\min_{x} ||Wx||_{*}, \text{ s.t. } Ax = b$$

Where Wx is the wavelet transform on the image x and A is a linear operator with the property $A : x \mapsto x_{\Omega}$. Here, b is simply the uncontaminated image domain. The methodology used to find the solution to this problem is discussed in section 2. The solution to this problem is not trivial because although the problem may be convex it is not smooth so the typical gradient descent approach would require some modifications to be used.

2 Methodology

2.1 Model

By nature, image inpainting is an ill posed problem because of the goal of reconstructing lost data. An image is modeled by a vector of numbers, b, representing the color at each point. This vector of numbers can be viewed as the result of a function, A. With this model, an uncontaminated image would have one solution to $Ax = b$, with $x = b$ representing the true image. A contaminated image would have many solutions to $Ax = b$ since only parts of the true image is known. Thus image inpainting is an under-determined problem with more variables than equations. To solve this ill posed problem, more constraints are added because without prior information there are infinite solutions. There are two common ways to add constraints. One may assume the function is piecewise continuous, resulting in a method using variation of PDEs [2], and one may assume that the image can be sparsified. It has been found that under a wavelet transformation the image becomes sparse.
Assuming sparsity under the wavelet transformation, image inpainting becomes a minimization problem of x under the wavelet transformation as discussed in the problem statement, namely

$$\min_{x} ||Wx||_{*}, \text{ s.t. } Ax = b \tag{1}$$

The method we implemented to solve this problem utilizes variable splitting and the Alternative Direction Minimization Method (ADMM) in which we begin by simply rewriting (1) as the following

$$\min_{Q,x} ||Q||_{*}, \text{ s.t. } Ax = b, \ Wx = Q \tag{2}$$

Using this modification we were able to minimize our target objective function (1) through the use of the augmented Lagrangian. Oftentimes constrained optimization problems are solved through the use of the method of Lagrangian multipliers where we define the Lagrangian for this problem as

$$\mathcal{L}(\boldsymbol{x}, Q, \boldsymbol{\lambda}) = ||Q||_{*} + < A\boldsymbol{x} - \boldsymbol{b}, \lambda_1 > + < Wx - Q, \lambda_2 > \tag{3}$$

However, note that this function is not differentiable and does not seem to yield any particularly "good" solutions to our problem. Thus, we use the method of the Augmented Lagrangian which is the same as (3) except we add penalty terms to simplify future calculations. We define the Augmented Lagrangian as so:

$$\mathcal{L}_A(\boldsymbol{x}, Q, \lambda_1, \lambda_2) = ||Q||_{*} + < A\boldsymbol{x} - \boldsymbol{b}, \lambda_1 > + \frac{r_1}{2}||A\boldsymbol{x} - \boldsymbol{b}||_2^2 + < Wx - Q, \lambda_2 > + \frac{r_2}{2}||Wx - Q||_2^2 \tag{4}$$

A good question to ask is whether or not this new formula will still minimize our objective function (1). It has been theoretically shown that $\mathcal{L}_A(\boldsymbol{x^*}, Q^*, \lambda_1, \lambda_2) < \mathcal{L}_A(\boldsymbol{x^*}, Q^*, \lambda_1^*, \lambda_2^*) < \mathcal{L}_A(\boldsymbol{x}, Q, \lambda_1^*, \lambda_2^*)\ \forall \boldsymbol{x}, Q, \lambda_1, \lambda_2$ [3]. Which shows that the saddle point of the Augmented Lagrangian is actually the minimizer of our target objective function (1). Thus, we can simply perform a combination of ascent and descent on the Augmented Lagrangian to get the solution.

2.2 Algorithm

The Alternative Direction Minimization Method was used to find the solution to (1) through the model described in 2.1 for which we will show the calculations resulting in our implemented algorithm.

Each variable is iteratively updated. First Q is updated by finding

$$\underset{Q}{\text{argmin}}\, \mathcal{L}(\boldsymbol{x^{k-1}}, Q^k, \lambda_1^{k-1}, \lambda_2^{k-1})$$

$$\underset{Q}{\text{argmin}}\, ||Q||_* + < A\boldsymbol{x} - \boldsymbol{b}, \lambda_1 > + \frac{r_1}{2}||A\boldsymbol{x} - \boldsymbol{b}||_2^2 + < Wx - Q, \lambda_2 > + \frac{r_2}{2}||Wx - Q||_2^2$$

$$\underset{Q}{\text{argmin}}\, ||Q||_* - < Q, \lambda_2 > + \frac{r_2}{2}||Wx - Q||_2^2$$

$$\underset{Q}{\text{argmin}}\, ||Q||_* + \frac{r_2}{2}||Wx - (Q - \frac{\lambda_2}{r_2}||_2^2$$

$$Q^k = shrink(Q - \frac{\lambda_2}{r_2}, 1)$$

The solution to this part of the minimization problem ends up simply being soft-thresholding. We can perform a similar method to solve for the minimizer with respect to x

$$\underset{\boldsymbol{x}}{\text{argmin}}\, \mathcal{L}(\boldsymbol{x^k}, Q^k, \lambda_1^{k-1}, \lambda_2^{k-1})$$

$$\underset{\boldsymbol{x}}{\text{argmin}}\, ||Q||_* + < A\boldsymbol{x} - \boldsymbol{b}, \lambda_1 > + \frac{r_1}{2}||A\boldsymbol{x} - \boldsymbol{b}||_2^2 + < Wx - Q, \lambda_2 > +$$

$$\underset{\boldsymbol{x}}{\text{argmin}} < A\boldsymbol{x}, \lambda_1 > + \frac{r_1}{2}||A\boldsymbol{x} - \boldsymbol{b}||_2^2 + < Wx, \lambda_2 > + \frac{r_2}{2}||Wx - Q||_2^2$$

$$0 = A^T\lambda_1 + r_1 A^T(Ax - b) + W^T\lambda_2 + r_2 W^T(Wx - Q)$$

$$(r_1 A^T A + r_2 I)x = A^T(-\lambda_1^{k-1} + r_1 b) + W^T(-\lambda_2^{k-1} + r_2 Q^k)$$

Solving this equation for x will then give us our next x value (Note: W^T is the inverse wavelet transform and $W^T W = I$). This can be efficiently solved using a method such as conjugate gradient. The values for λ_1 and λ_2 are then simply updated using gradient descent.

Algorithm 1 Alternative Direction Minimzation Method (ADMM)

Step 0. $Q^0 = \boldsymbol{0}, \boldsymbol{x}^0 = \boldsymbol{0}, \lambda_1 = \lambda_2 = 0$
while not converge **do**
  Step 1. $(r_1 A^T A + r_2 I)\boldsymbol{x}^k = A^T(r_1\boldsymbol{b} - \lambda_1^{k-1}) + W^T(r_2 Q^k - \lambda_2^{k-1})$
  Step 2. $Q^k = shrink(Wx^k - \frac{\lambda_2^{k-1}}{r_2}, 1)$
  Step 3. $\lambda_1^k = \lambda_1^{k-1} + A\boldsymbol{x}^k - \boldsymbol{b}$
  Step 4. $\lambda_2^k = \lambda_2^{k-1} + W\boldsymbol{x}^k - Q^k$
end while

---

3 Results

Below we consider the reconstruction of the classic Barbara image. We were given the image on the left and the contaminated domain (of text in this case) that we have to reconstruct. In order to solve this problem we applied Algorithm 1 from section 2 to minimize the rank of the sparse image matrix. The program (see Appendix) was able to reconstruct the image of Barbara on the right after 123 iterations with r_1 and r_2 values of both 150. The exit criteria chosen for this image reconstruction was for the program to finish after the $||Q^k - Q^{k-1}||_2$ fell below some threshold ε which was chosen to be 17 in our case but will vary depending on implementation. The results are shown below in Figure 1.

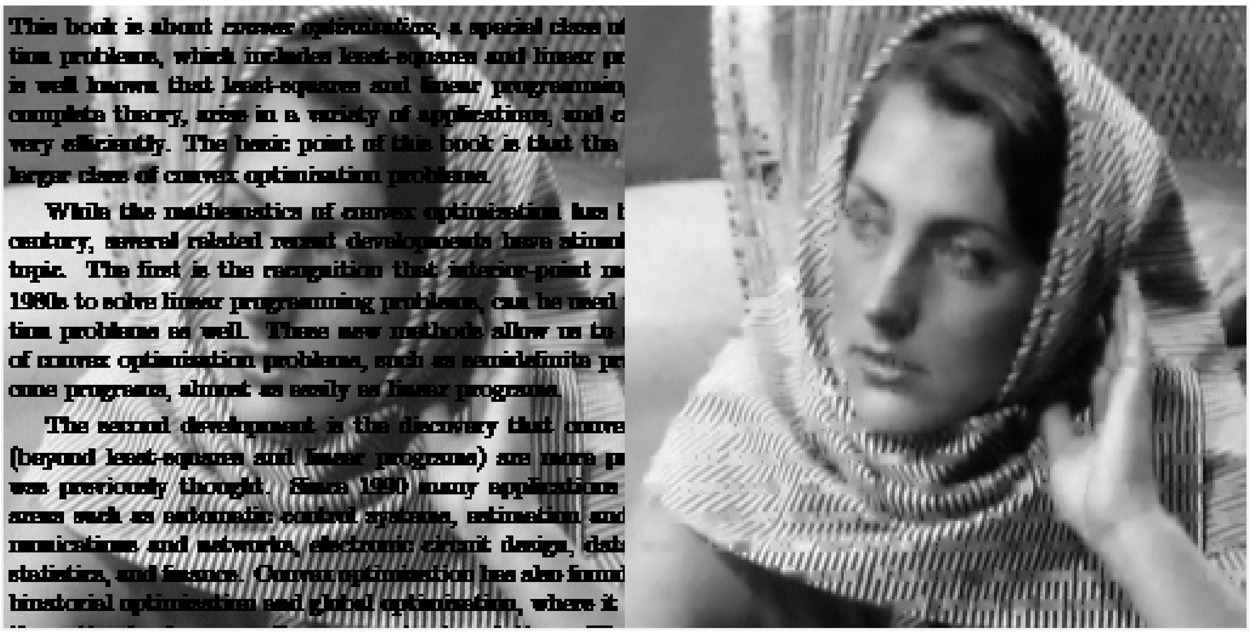

Figure 1: Barbara Contaminated (Left), Barbara Reconstructed (Right).

Using the same stopping condition and values for r_1 and r_2 (both 150), the described algorithm was applied to the contaminated Cameraman image which was reconstructed to within the exit criteria after 155 iterations producing the results shown in the right portion of Figure 2. After reconstructing the Cameraman we found it constructive to observe how quickly the root mean square error improves with each iteration as shown in Figure 3.

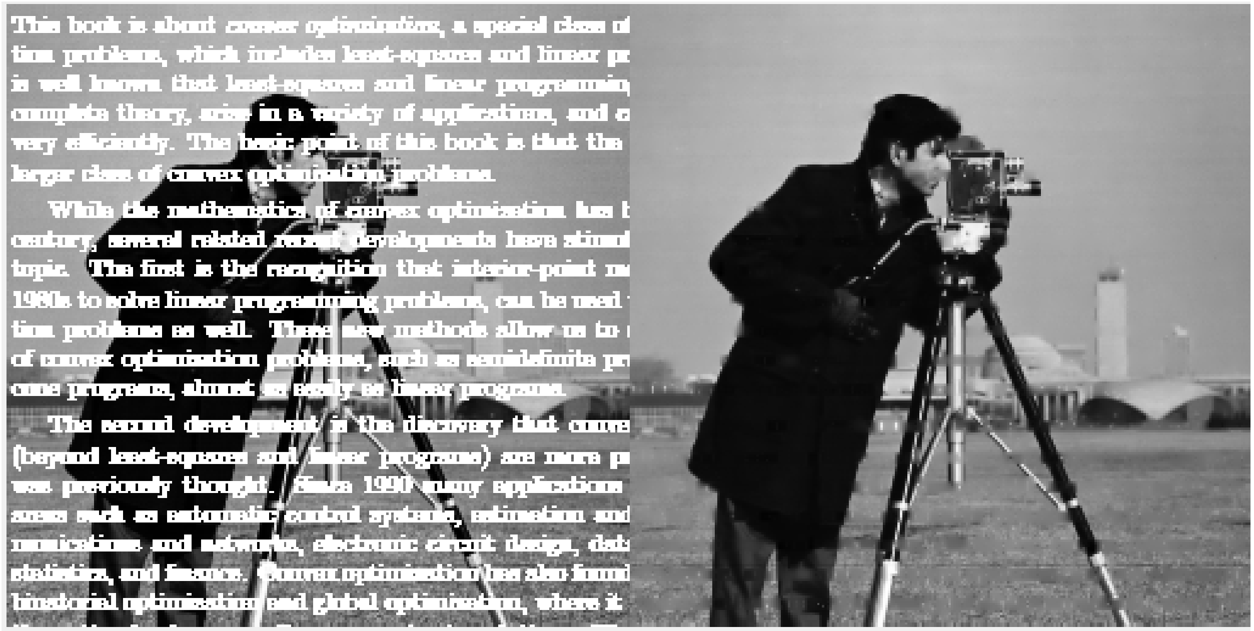

Figure 2: Cameraman Contaminated (Left), Cameraman Reconstructed (Right).

4 Observation and Conclusions

The image reconstructions for both the Cameraman and Barbara produced by this method intuitively appear to be good approximations of the original images (we show the exact error for Cameraman in 3). However, one may notice that using this method can result in minor artifacts resulting in somewhat blurry portions of the image that should actually be sleek straight lines. For instance, consider the lower right portion of the resultant Cameraman's jacket and notice slight blurry artifacts produced through this method. In order to produce an image reconstruction that avoids this issue a new method may need to be implemented such as the total variation approach. However, overall this method does a good job at reconstructing images with missing data.

After implementing our iterative approach to matrix completion we computed the Root Mean Square Error between our method's reconstruction of the Cameraman image and the ground truth Cameraman image. This error was calculated every iteration and plotted against the iteration number as shown in Figure 3. From this result we can see that the method we proposed does a good job at repairing images as the error is consistently reduced until we asymptotically approach a final error value which is simply the best approximation the method will generate. Future work in this area could include studying other methods of image inpainting such as minimizing total variation.

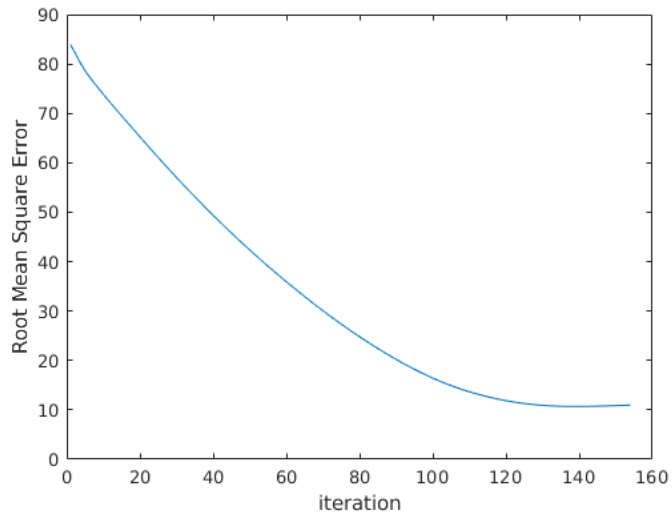

Figure 3: Iterative Root Mean Square Error for Cameraman

In this paper we explored the under constrained problem of image inpainting. From infinitely many solutions a method is devised to pick a solution that is a good approximation of the ground truth. Given an image contaminated in a domain, the image is reconstructed by utilizing the given wavelet tight frame transformation. The wavelet transform sparsifies the image in order to minimize the l_1 norm. This constrained convex optimization problem was solved using variable splitting and the Alternative Direction Minimization Method. Each variable was iteratively updated until the l_2 norm of the successive point difference of Q went below a determined threshold. Using this method, contaminated images of Barbara and the Cameraman were successfully restored given a contaminated image domain. Possible future extensions of this method include image super-resolution through the use of a Gaussian or Bernoulli random matrices instead of a known image domain and to perform low-rank matrix completion based on the chosen random matrix [4].

References

[1] Candes, Emmanuel J, and Benjamin Recht. ”Exact Matrix Completion via Convex Optimization.” pp. 15. ArXiv, 29 May 2008

[2] Zhang, F., Chen, Y., Xiao, Z., Geng, L., Wu, J., Feng, T., . . . Wang, J. (2015). Partial Differential Equation Inpainting Method Based on Image Characteristics. Lecture Notes in Computer Science Image and Graphics, 11-19.

[3] Lai, Rongjie. ”Convex Optimization and Low Rank Matrix Completion.” Rensselaer Polytechnic Institute. Computational Optimization, 2017-2018, Troy, New York.

[4] Divekar, Atul, and Deanna Needell. ”Using Correlated Subset Structure for Compressive Sensing Recovery.” 10 June 2013 1-4