# ORT Braude College 

Department of Mathematics<br>B.SC. FINAL YEAR PROJECT

# First order gradient methods 

Author:
Shira Bar-Dov

Supervisor:
Assoc. Prof. Aviv Gibali

July 19, 2019

## Contents

1 Abstract ..... 2
2 Introduction ..... 3
3 Preliminaries ..... 6
4 Data Fitting ..... 8
4.1 ISTA ..... 9
4.2 FISTA ..... 12
4.3 Numerical experiments ..... 13
5 Classification ..... 19
5.1 Pegasos ..... 21
5.2 Numerical experiments ..... 22
6 Conclusion ..... 25
7 Acknowledgements ..... 25
8 Matlab Codes ..... 26
8.1 Data Fitting ..... 26
8.2 Classification ..... 29

## 1 Abstract

In this project we describe and analyze different kinds of iterative first order gradient algorithms for solving optimization problems arising in image processing and support vector machines.
This class of methods, which can be viewed as an extension of the classical gradient algorithms, are attractive due to their simplicity and thus they adequate for solving large-scale problems even with dense matrix data. However, such methods are also known to converge quite slowly.
In this project we present several powerful iterative descent method for finding a local minimum of a multivariable function, that have been proven to converge faster than classical gradient algorithms.
Numerical experiments for image deblurring and classification problems are given.

## 2 Introduction

In this project we are concerned with optimization models for data fitting and data classification. Specifically, we focus on first order gradient descent methods such as ISTA- Iterative Shrinkage Thresholding Algorithm and FISTA- Fast Iterative Shrinkage Thresholding Algorithm for data fitting and a stochastic method called Pegasos for data classification.
Gradient descent was invented by the french mathematician Louis Augustin Cauchy in 1847. Since then, gradient descent methods kept developing, subgradients methods and stochastic subgradient methods from convex optimization were discovered during 1960-1970.
Gradient descent is a first-order iterative optimization algorithm for finding the minimum of a function. It is a common optimization method in machine learning algorithms, which are based on a convex function. Through an iterative process and the use of partial differential equations, gradient descent constructs a set of parameters, in order to minimize a given cost function to its local minimum.
Gradient descent is based on the observation that if the multivariable function $F(x)$ is defined and differentiable in a neighborhood of a point $x$, then $F(x)$ decreases fastest if one goes from $x$ in the direction of the negative gradient of $F$ at $x,-\nabla F(x)$. It follows that, if

$$
x_{n+1}=x_{n}-\gamma \nabla F\left(x_{n}\right)
$$



Figure 1: The iterates of the gradient method along with the contour lines of the objective function
for $\gamma \in \mathbb{R}_{+}$small enough, the term $\gamma \nabla F(x)$ is subtracted from $x$ because we want to move against the gradient, toward the minimum.

The stepsize or the learning rate $\gamma$ is a positive number that determines the size of each step in the process. If it is too small, the process could be slow. On the other hand, if it is too large the process can skip the minimum and may not converge.
With this in mind, consider the unconstrained nonsmooth convex optimization problem

$$
\begin{equation*}
\min _{x}\left\{F(x) \equiv f(x)+g(x): x \in \mathbb{R}^{n}\right\} \tag{1}
\end{equation*}
$$

where $f, g$ are convex function, with $g$ possibly nonsmooth.
The gradient algorithm is one of the simplest methods for solving (1). It generates a sequence $\left\{x_{k}\right\}$ via

$$
x_{0} \in \mathbb{R}^{n}, \quad x_{k}=x_{k-1}-t_{k}\left(\nabla f\left(x_{k-1}\right)+\nabla g\left(x_{k-1}\right)\right),
$$

where $t_{k}>0$ is a suitable stepsize.

This problem arises at many different applications among them signal and image processing, optics, speech tagging and music identification.
When solving large-scale problems, first-order methods are often the only practical option, but the sequence $x_{k}$ converges quite slowly to a solution.

In this project, we focus on two specific models, one for data fitting with and without regulators, and one for classification.
Data fitting and classifying data are common tasks in machine learning. If a parametrized model function meant to explain some phenomena is given, the goal is to adjust the numerical values for the model so that it most closely matches some data.

In data fitting problems, a classical approach to the familiar basic linear problem (1) where

$$
f(x)=A x-b, g(x)=0
$$

where $A \in \mathbb{R}^{m \times n}, x \in \mathbb{R}^{n}, b \in \mathbb{R}^{m}, R(A)$ the column space of $A$.
is the least squares, which estimate the solution by minimize the data error. In other words, if we define the error $r=A x-b$ then we find $x=x_{l s}$ that minimize $\|r\| . x_{l s}$ is called least-squares approximate solution of $A x=b$. $A x_{l s}$ is point in $R(A)$ closest to $b$, and the projection of $b$ onto $R(A)$.


Note that if $b \in R(A)$ then $x_{l s}$ solves $A x_{l s}=b$, or if $A$ is square it is invertible (nonsingular) so $x_{l s}=A^{-1} b$.

In classification problems, data points are given, and each belong to one of two or more classes. When a new point is given, the goal is to determine in which class it will be. In support vector machines, each data point is a $p$ dimensional vector, and the goal is to determine whether it could be separated with a $p-1$ dimensional hyperplane. Note that there are many hyperplanes that might classify the data, but the best hyperplane is the one that represents the largest separation (margin) between the two classes. Hence, the optimal hyperplane is the one that maximized the distance from it to the nearest data point on each side.

## 3 Preliminaries

Convex function: A function $f$ is convex if for every $u, v$ in the domain, and for every $\lambda \in[0,1]$ we have

$$
f(\lambda u+(1-\lambda) v) \leq \lambda f(u)+(1-\lambda) f(v)
$$

In general, the necessary condition for $x$ to be a minimum for the function $f$ is $\nabla f(x)=0$. For convex functions, this is both necessary and sufficient.


Geometrically, the inequality in the definition means the graph of $f$ between $u$ and $v$ is below the segment which joins the points $(u, f(u))$ and $(v, f(v))$.
$l_{1}$ Norm: $l_{1}$ Norm is the sum of the magnitudes of the vectors in a space. A vector norm defined for a vector $\vec{x}=\left(x_{1}, \ldots, x_{n}\right)$ as $\|\vec{x}\|_{1}=\sum_{i=1}^{n}\left|x_{i}\right|$.

Dense matrix: A matrix in which most of the elements are nonzero.
Lipschitz continuous gradient: A function $f$ that uphold

$$
\|\nabla f(x)-\nabla f(y)\| \leq L(f)\|x-y\| \quad \text { for every } x, y \in \mathbb{R}^{n}
$$

where $\|\cdot\|$ denotes the standard Euclidean norm and $L(f)>0$ is the Lipschitz constant of $\nabla f$.

Hyperplane: A hyperplane in an $n$ dimensional Euclidean space, is a flat $n-1$ dimensional subset of that space that divides the space into two disconnected parts.

Inner product: The inner product of two vectors $\vec{a}=\left(a_{1}, a_{2}, \ldots, a_{n}\right)$ and $\vec{b}=\left(b_{1}, b_{2}, \ldots, b_{n}\right)$ is defined as:

$$
a \cdot b=a^{T} b=\sum_{i=1}^{n} a_{i} b_{i}=a_{1} b_{1}+a_{2} b_{2}+\ldots+a_{n} b_{n} .
$$

Gradient: The gradient of a scalar multivariable function $f\left(x_{1}, x_{2}, \ldots, x_{n}\right)$ is denoted $\nabla f$, packages all its partial derivative information into a vector: $\nabla f(\vec{x})=\left(\frac{\partial f}{\partial x_{1}}, \frac{\partial f}{\partial x_{2}}, \ldots, \frac{\partial f}{\partial x_{n}}\right)$.

Smooth function: A smooth function is a function that has derivatives of all orders everywhere in its domain.

## 4 Data Fitting

In this section we discuss first order methods of gradient descent iterative algorithms, for solving linear inverse problems arising in image processing. A basic linear inverse "image deblurring" problem $A x=b$

where $A \in \mathbb{R}^{m \times n}$ represents a two-dimensional convolution blur operator, and $b \in \mathbb{R}^{m}$ represents the blurred image, are known. $x \in \mathbb{R}^{n}$ is the "true" and unknown image to be estimated, its size is assumed to be the same as that of $b(n=m)$. Both $x$ and $b$ are formed by stacking the columns of their corresponding two-dimensional images.
The blur operator or the Gaussian operator is a two-dimensional convolution operator that is used to blur images and remove detail and noise. The idea of Gaussian smoothing is to use this two-dimensional distribution as a 'pointspread' function, and this is achieved by convolution of a Gaussian mask h and the true image $x$.


### 4.1 ISTA

The class of Iterative Shrinkage Thresholding Algorithms (ISTA) is an extension of the classical gradient algorithm. It is attractive due to its simplicity and thus is adequate for solving large-scale problems even with dense matrix data.

Adopting this same gradient idea to the problem formulation (1) where the following assumptions are made:
$g: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is a continuous convex function which is possibly nonsmooth.
$f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is a continuously differentiable with Lipschitz continuous gradient $L(f)$.

Note that when $g(x) \equiv 0,(1)$ is the general unconstrained smooth convex minimization problem.
When $f(x)=\|A x-b\|^{2}, g(x) \equiv\|x\|_{1},(1)$ is the $l 1$ regularization problem

$$
\min \left\{f(x)+\lambda\|x\|_{1}: x \in \mathbb{R}^{n}\right\}
$$

To calculate the Lipschitz constant of the gradient $\nabla f$ we first write $f$ as follows:

$$
\begin{aligned}
f(x) & =\|A x-b\|^{2}=(A x-b)^{T}(A x-b)=x^{T} A^{T} A x-x^{T} A^{T} b-b^{T} A x+b^{T} b= \\
& =x^{T} A^{T} A x-\left(b^{T} A x\right)^{T}-b^{T} A x+b^{T} b=x^{T} A^{T} A x-2 b^{T} A x+b^{T} b .
\end{aligned}
$$

The last transaction is because the product $b^{T} A x$ is a scalar product of vector $b$ and $A x$, and therefore $\left(b^{T} A x\right)^{T}=b^{T} A x$.
The gradient of $\nabla f(x)$ :
$\nabla f(x)=\nabla\left(x^{T} A^{T} A x-2 b^{T} A x+b^{T} b\right)=\left(A^{T} A+\left(A^{T} A\right)^{T}\right) x-2\left(b^{T} A\right)^{T}=2 A^{T} A x-2 A^{T} b$.
We define the Lipschitz continuous gradient at the preliminaries section,
$\|\nabla f(x)-\nabla f(y)\|=\left\|2 A^{T} A x-2 A^{T} b-\left(2 A^{T} A y-2 A^{T} b\right)\right\|=\left\|2 A^{T} A(x-y)\right\| \leq 2\left\|A^{T} A\right\|\|x-y\|$.
The inequality is a result of the Cauchy Bunyakovsky Schwarz inequality, therefore the Lipschitz constant of the gradient is $2\left\|A^{T} A\right\|$.

Theorem 1 Let $A$ be a symmetric $n \times n$ matrix, $\|A\|=\max _{j}\left|\lambda_{j}\right|$, where $\lambda_{j}$ are the eigenvalues of $A$.

Proof. The norm of a matrix is defined as

$$
\|A\|=\max _{\|u\|=1}\|A u\|
$$

Taking the singular value decomposition of the matrix $A$, we have

$$
A=V D W^{T}
$$

where $V$ and $W$ are re orthonormal and $D$ is a diagonal matrix. Since $V$ and $W$ are orthonormal, we have $\|V\|=1$ and $\|W\|=1$. Then $\|A v\|=\|D v\|$ for any vector $v$. Then we can maximize the norm of $A v$ by maximizing the norm of $D v$. By the definition of singular value decomposition, $D$ will have the singular values of $A$ on its main diagonal and will have zeros everywhere else. Let $\lambda_{1}, \ldots, \lambda_{n}$ denote these diagonal entries so that

$$
D=\left(\begin{array}{cccc}
\lambda_{1} & 0 & \cdots & 0 \\
0 & \lambda_{2} & \cdots & 0 \\
\vdots & & \ddots & \vdots \\
0 & 0 & \cdots & \lambda_{n}
\end{array}\right)
$$

Taking some $v=\left(v_{1}, v_{2}, \cdots, v_{n}\right)^{T}$, the product $D v$ takes the form

$$
D v=\left(\begin{array}{c}
\lambda_{1} v_{1} \\
\vdots \\
\lambda_{n} v_{n}
\end{array}\right)
$$

Maximizing the norm of this is the same as maximizing the norm squared. Then we are trying to maximize the sum

$$
S=\sum_{i=1}^{n} \lambda_{i}^{2} v_{i}^{2}
$$

under the constraint that $v$ is a unit vector (i.e., $\sum_{i} v_{i}^{2}=1$ ). The maximum is attained by finding the largest $\lambda_{i}^{2}$ and setting its corresponding $v_{i}$ to 1 and then setting each other $v_{j}$ to 0 . Then the maximum of $S$ (which is the norm squared) is the square of the absolutely largest eigenvalue of $A$. Taking the square root, we get the absolutely largest eigenvalue of $A$.
$A^{T} A$ is a symmetric matrix, hence, the smallest Lipschitz constant of the gradient $\nabla f$ is $L(f)=2\left\|A^{T} A\right\|=2 \lambda_{\max }\left(A^{T} A\right)$, where $\lambda_{\max }$ is the maximal eigenvalue of $A^{T} A$.

One of the methods for solving (1) is in the class of ISTA. The ISTA method which generates a sequence $\left\{x_{k}\right\}$ in this case would be:

$$
\begin{equation*}
x_{0} \in \mathbb{R}^{n}, \quad x_{k}=x_{k-1}-\frac{1}{L}\left(\nabla f\left(x_{k-1}\right)+\nabla g\left(x_{k-1}\right)\right) . \tag{2}
\end{equation*}
$$

The ISTA iteration (2) can be viewed as a proximal regularization of the linearized function $f$ at $x_{k-1}$, and written equivalently as

$$
x_{k}=\underset{x}{\operatorname{argmin}}\left\{f\left(x_{k-1}\right)+\left\langle x-x_{k-1}, \nabla f\left(x_{k-1}\right)\right\rangle+\frac{L}{2}\left\|x-x_{k-1}\right\|^{2}+g(x)\right\} .
$$

For any $L>0$, consider the quadratic approximation of (1) at a given point $y$ :

$$
Q_{L}(x, y):=f(y)+\langle x-y, \nabla f(y)\rangle+\frac{L}{2}\|x-y\|^{2}+g(x),
$$

which admits a unique minimizer

$$
p_{L}(y):=\operatorname{argmin}\left\{Q_{L}(x, y): x \in \mathbb{R}^{n}\right\} .
$$

Simple Algebra shows that

$$
\begin{gathered}
Q_{L}(x, y):=f(y)+\langle x-y, \nabla f(y)\rangle+\frac{L}{2}\langle x-y, x-y\rangle+g(x)= \\
=g(x)+\frac{L}{2}\left\langle x-y, \frac{2}{L} \nabla f(y)\right\rangle+\frac{L}{2}\langle x-y, x-y\rangle+f(y)= \\
=g(x)+\frac{L}{2}\left(\langle x-y, x-y\rangle+2\left\langle x-y, \frac{1}{L} \nabla f(y)\right\rangle\right)+f(y)= \\
=g(x)+\frac{L}{2}\left(\left\langle x-y+\frac{1}{L} \nabla f(y), x-y+\frac{1}{L} \nabla f(y)\right\rangle-\left\langle\frac{1}{L} \nabla f(y), \frac{1}{L} \nabla f(y)\right\rangle\right)+f(y)= \\
=g(x)+\frac{L}{2}\left\|x-y+\frac{1}{L} \nabla f(y)\right\|^{2}-\frac{1}{2 L}\|\nabla f(y)\|+f(y) .
\end{gathered}
$$

ignoring constant terms in $y$

$$
p_{L}(y)=\underset{x}{\operatorname{argmin}}\left\{g(x)+\frac{L}{2}\left\|x-\left(y-\frac{1}{L} \nabla f(y)\right)\right\|^{2}\right\} .
$$

Then, the gradient algorithm ISTA generates a sequence $\left\{x_{k}\right\}$ via

$$
x_{0} \in \mathbb{R}^{n}, x_{k}=p_{L}\left(x_{k-1}\right),
$$

where $L>0$ plays the role of a stepsize.

To conclude, the basic iteration of ISTA for solving problem (1) with constant stepsize is:

```
Input: \(L:=L(f)-\mathrm{A}\) Lipschitz constant of \(\nabla f\).
Initialize: Take \(x_{0} \in \mathbb{R}^{n}\)
Step k: For each \(k \geq 1\) compute
```

$$
x_{k}=p_{L}\left(x_{k-1}\right) .
$$

### 4.2 FISTA

A Fast Iterative Shrinkage-Thresholding Algorithm (FISTA) is an improvement of the ISTA method for solving the general problem (1). While FISTA keeps the simplicity of ISTA, the sequence $x_{k}$ that FISTA generates converges more quickly to a solution than other gradient methods.
The main difference between the ISTA method and the FISTA method is that the iterative shrinkage operator $p_{L}()$ is not employed on the previous point $x_{k-1}$, but rather at the point $y_{k}$, which uses a very specific linear combination of the previous two points $x_{k 1}, x_{k 2}$. Obviously, the main computational effort in both ISTA and FISTA remains the same, because the requested additional computation for FISTA in the steps (4) and (5) is clearly marginal.

The basic iteration of FISTA for solving problem (1) with constant stepsize is:

Input: $L:=L(f)$ - A Lipschitz constant of $\nabla f$.
Initialize: Take $y_{1}=x_{0} \in \mathbb{R}^{n}, t_{1}=1$
Step k: For each $k \geq 1$ compute

$$
\begin{gather*}
x_{k}=p_{L}\left(x_{k-1}\right)  \tag{3}\\
t_{k+1}=\frac{1+\sqrt{1+4 t_{k}^{2}}}{2}  \tag{4}\\
y_{k+1}=x_{k}+\left(\frac{t_{k}-1}{t_{k+1}}\right) \tag{5}
\end{gather*}
$$

### 4.3 Numerical experiments

In this section we illustrate an image deblurring problem. We estimate an image $\boldsymbol{x}$ from blurred image $\boldsymbol{b}$ using the iterative gradient descent algorithms ISTA and FISTA. We compare ISTA to FISTA and show the difference between the performance of those methods.
Both methods were used with a constant stepsize rule and applied once when

$$
\begin{equation*}
f(x)=\|A x-b\|^{2}, \quad g(x)=0 \tag{6}
\end{equation*}
$$

and once with the $l_{1}$ regulator when

$$
\begin{equation*}
f(x)=\|A x-b\|^{2}, \quad g(x)=\lambda\|x\|_{1} . \tag{7}
\end{equation*}
$$



Figure 2: Deblurring of the cameraman

In this example we look at the $250 \times 250$ cameraman image. The image went through a Gaussian blur of size $9 \times 9$ and standard deviation 4 (applied by the MATLAB functions conv2 and fspecial). The original and blurred images are given in Figure (2).
We tested ISTA and FISTA for solving problems (6) and (7), where $\boldsymbol{b}$ represents the vectorized blurred image, $\boldsymbol{A}$ represents the blur operator, which is a two-dimensional convolution operator, and $\boldsymbol{x}$ represent the unknown true image. The regularization parameter has been chosen to be $\lambda=1$. The Lipschitz constant was computable in this example, since the maximum eigenvalue of $A^{T} A$ can be calculated.
Iterations 5, 50, 100 and 200 of ISTA algorithm are described in Figure (3). Iterations 5, 50, 100 and 200 of FISTA algorithm are described in Figure (4).

The peak signal to noise ratio (PSNR) for the reconstructed image is calculated, with the true image as the reference. This ratio denotes how close are we to the true image, where higher values would denote that the image is closer to the true image.
The function value that we want to minimize is given at those iterations, in order to track its value as it gets lower with the iterations.

## ISTA's iterations:




Figure 3: Iterations of ISTA for deblurring of the cameraman.

## FISTA's iterations:




Figure 4: Iterations of FISTA for deblurring of the cameraman.
The data is summarized in the following tables:

| Function value of $f(x)+g(x)=\\|A x-b\\|$ |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- |
|  | 5 iterations | 50 iterations | 100 iterations | 200 iterations |
| ISTA | $3.453 \mathrm{e}+05$ | $2.662 \mathrm{e}+04$ | $1.292 \mathrm{e}+04$ | $5.535 \mathrm{e}+03$ |
| FISTA | $2.033 \mathrm{e}+05$ | $2.282 \mathrm{e}+03$ | $4.887 \mathrm{e}+02$ | $9.660 \mathrm{e}+01$ |
| Function value of $f(x)+g(x)=\\|A x-b\\|+\\|x\\|_{1}$ |  |  |  |  |
|  | 5 iterations | 50 iterations | 100 iterations | 200 iterations |
| ISTA | $5.079 \mathrm{e}+06$ | $4.766 \mathrm{e}+06$ | $4.752 \mathrm{e}+06$ | $4.745 \mathrm{e}+06$ |
| FISTA | $4.940 \mathrm{e}+06$ | $4.742 \mathrm{e}+06$ | $4.740 \mathrm{e}+06$ | $4.739 \mathrm{e}+06$ |

It is clear that FISTA gives better results than ISTA, because the function values of FISTA are consistently lower than the function values of ISTA. Note that the function value of FISTA after only 50 iterations is better than the function value of ISTA after 200 iterations.

| PSNR: peak signal to noise ratio for $f(x)+g(x)=\\|A x-b\\|$ |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- |
|  | 5 iterations | 50 iterations | 100 iterations | 200 iterations |
| ISTA | $2.172+01$ | $2.367+01$ | $2.446+01$ | $2.533+01$ |
| FISTA | $2.201+01$ | $2.618+01$ | $2.786+01$ | $2.993+01$ |


| PSNR: peak signal to noise ratio for $f(x)+g(x)=\\|A x-b\\|+\\|x\\|_{1}$ |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- |
|  | 5 iterations | 50 iterations | 100 iterations | 200 iterations |
| ISTA | $2.171+01$ | $2.368+01$ | $2.448+01$ | $2.535+01$ |
| FISTA | $2.200+01$ | $2.628+01$ | $2.810+01$ | $3.025+01$ |

Here too, it is clear that FISTA gives better results than ISTA, because the PSNR of the FISTA method compared with the original image is consistently higher than the PSNR of the ISTA method compared with the original image.
Note that the PSNR of function $f(x)+g(x)=\|A x-b\|+\|x\|_{1}$ is consistently higher than the PSNR of function $f(x)+g(x)=\|A x-b\|$, so we can say that the $l_{1}$ regulator helps the iterations to converge faster.

## 5 Classification

In this section we discuss Stochastic Sub-Gradient Descent algorithms (SGD) for solving the optimization problem cast by Support Vector Machines.
SGD is a modification of the basic gradient descent algorithm, which allows us scaling these algorithms to much bigger training sets.
The problem with gradient descent is that if the number of the training examples is much bigger, computing the gradient could be very expensive or even impossible because each step requires storing all the data, and calculates the derivative. SGD doesn't need to look at all of the training set in every single iteration, but only at a single training example.
Support vector machine is a linear model for classification and regression problems that constructs a hyperplane that separates the data into classes. First, it randomly reorders the data examples, a pre-processing step which ensures that when scanning through the training set, the order of visiting the training examples would be in randomly sorted order, that speeds up the convergence.
Then it finds the points closest to the hyperplane from both the classes. These points are called support vectors. Then, the distance between the hyperplane and the support vectors is computed. This distance is called the margin. The goal is to maximize the margin, so the hyperplane for which the margin is maximized is the optimal hyperplane.
The main difference between gradient descent and SGC is that in SGC the algorithm modifies the parameters a little bit for every iteration on one training example to fit just the specific training example a little bit better.

Formally, given a training set $S=\left\{\left(\vec{x}_{i}, y_{i}\right)\right\}_{i=1}^{m}$ where $\vec{x}_{i} \in \mathbb{R}^{n}$ and $y_{i} \in$ $\{+1,1\}$, we would like to find the minimizer of the problem

$$
\begin{equation*}
\min _{\vec{w}} \frac{\lambda}{2}\|\vec{w}\|^{2}+\frac{1}{m} \sum_{(\vec{x}, y) \in S} l(\vec{w} ;(\vec{x}, y)), \tag{8}
\end{equation*}
$$

where

$$
\begin{equation*}
l(\vec{w} ;(\vec{x}, y))=\max \{0,1-y\langle\vec{w}, \vec{x}\rangle\}, \tag{9}
\end{equation*}
$$

$\langle u, v\rangle$ denotes the standard inner product between the vectors $u$ and $v, \lambda$ is the regularization parameter of SVM, and $l$ is called the Loss function.

We denote that

$$
f(w)=\frac{\lambda}{2}\|\vec{w}\|^{2}+\frac{1}{m} \sum_{(\vec{x}, y) \in S} l(\vec{w} ;(\vec{x}, y))
$$

hence, the gradient decent iterative algorithm in this case would be:

$$
w_{t+1}=w_{t}-\eta_{t} \nabla f\left(w_{t}\right)
$$

where $\eta_{t}>0$ is called the learning rate. The gradient $\nabla f$ would be:

$$
\nabla f(w)=\lambda\|w\|+\frac{\partial}{\partial w} \max \{0,1-y \vec{w} \cdot \vec{x}\}
$$

where the sub gradient is

$$
\frac{\partial}{\partial w}=\left\{\begin{align*}
-y_{i} x_{i}, & y_{i} \vec{w} \cdot \overrightarrow{x_{i}}<1  \tag{10}\\
0, & y_{i} \vec{w} \cdot \overrightarrow{x_{i}}=1 \\
0, & y_{i} \vec{w} \cdot \overrightarrow{x_{i}}>1
\end{align*}\right.
$$

if $y_{i} \vec{w} \cdot \overrightarrow{x_{i}}<1$ than

$$
\begin{equation*}
w_{t+1}=w_{t}-\eta_{t}\left(\lambda\left\|w_{t}\right\|-y_{i} x_{i}\right) \tag{11}
\end{equation*}
$$

else $\left(y_{i} \vec{w} \cdot \overrightarrow{x_{i}} \geq 1\right)$ than

$$
\begin{equation*}
w_{t+1}=w_{t}-\eta_{t} \lambda\left\|w_{t}\right\| \tag{12}
\end{equation*}
$$

Any hyperplane can be written as the set of points $\vec{x}$ satisfying $\vec{w} \cdot \vec{x}-b=0$ where $\vec{w}$ is the normal vector to the hyperplane. The following constraint were added to prevent data points from falling into the margin, for each $i$ either:

$$
\vec{w} \cdot \vec{x}_{i}-b \geq 1, \text { if } y_{i}=1
$$

or

$$
\vec{w} \cdot \vec{x}_{i}-b \leq 1, \text { if } y_{i}=-1 .
$$

These constraints state that each data point must lie on the correct side of the margin. It could also be written as:

$$
\begin{equation*}
y_{i}\left(\vec{w} \cdot \vec{x}_{i}-b\right) \geq 1, \quad \text { for all } 1 \leq i \leq n \tag{13}
\end{equation*}
$$



Figure 5: Maximum-margin hyperplane and margins for an SVM trained with samples from two classes

### 5.1 Pegasos

Pegasos performs stochastic gradient descent on the primal objective (8) with a carefully chosen stepsize. Pegasos gets as input $S$ set of data examples, $\lambda$ the regularization parameter of SVM and $T$ the number of iterations.

Input: $S, \lambda, T$
Initialize: Set $w_{1}=0$
For $t=1,2, \ldots, T$
Choose $i_{t} \in\{1, \ldots,|S|\}$ uniformly at random.
Set $\eta_{t}=\frac{1}{\lambda t}$
If $y_{i_{t}}\left\langle w_{t}, x_{i_{t}}\right\rangle<1$, then:
Set $w_{t+1} \leftarrow\left(1-\eta_{t} \lambda\right) w_{t}+\eta_{t} y_{i_{t}} x_{i_{t}}$
Else (if $y_{i_{t}}\left\langle w_{t}, x_{i_{t}}\right\rangle \geq 1$ ):
Set $w_{t+1} \leftarrow\left(1-\eta_{t} \lambda\right) w_{t}$
Output: $w_{T+1}$
On each iteration Pegasos operates in the following manner. Initially, it sets $w_{1}$ to the zero vector. On iteration $t$, it first chooses a random training example $\left(x_{i_{t}}, y_{i_{t}}\right)$ by picking an index $i_{t} \in 1, \ldots, m$ uniformly at random. Then, it replaces the objective in (8) with an approximation based on the
training example $\left(x_{i_{t}}, y_{i_{t}}\right)$, yielding:

$$
f\left(w ; i_{t}\right)=\frac{\lambda}{2}\|w\|^{2}+l\left(w ;\left(x_{i_{t}}, y_{i_{t}}\right)\right) .
$$

The sub gradient of the above approximate objective is given by:

$$
\nabla_{t}=\lambda w_{t}-\mathbb{1}\left[y_{i_{t}}\left\langle w_{t}, x_{i_{t}}\right\rangle<1\right] y_{i_{t}} x_{i_{t}},
$$

where $\mathbb{1}[y\langle w, x\rangle<1]$ is the indicator function which takes a value of one if its argument is true ( $w$ yields non-zero loss on the example $(x, y)$ ), and zero otherwise. Then it updates $w_{t+1} \leftarrow w_{t}-\eta_{t} \nabla t$ using a step size of $\eta_{t}=\frac{1}{\lambda t}$. Note that this update can be written as:

$$
w_{t+1} \leftarrow\left(1-\frac{1}{t}\right) w_{t}+\eta_{t} \mathbb{1}\left[y_{i_{t}}\left\langle w_{t}, x_{i_{t}}\right\rangle<1\right] y_{i_{t}} x_{i_{t}} .
$$

After $T$ iterations, the last iterate $w_{T+1}$ is returned.

### 5.2 Numerical experiments

In this section we illustrate a data classification problem. We estimate a hyperplane $w^{T} x-b$ from a set of data and its labels $\left\{\left(\vec{x}_{i}, y_{i}\right)\right\}_{i=1}^{m}$ using the a built in SVM classifier in MATLAB and a stochastic gradient descent algorithm Pegasos. We compare those two algorithms and show the hyperplane each function finds.
In this example we would like to classify recipes as cupcakes or muffins. When given a new recipe, our model could determine if it's a cupcake or a muffin. We took 82 different recipes of cupcakes and muffins and normalized the data to percentage of the whole batter. Then, we chose 2 ingredientsbutter and sugar that we've noticed that are different between those two types of recipes.


As you can see muffins have little sugar and little butter and cupcakes have a lot of sugar and a lot of butter, therefore it's a good example for SVM classify.
First, the data went through the SVM classifier (applied by the MATLAB function fitcsvm), the hyperplane and the support vectors solution for this problem are given in Figure (6).


Figure 6: SVM hyperplane and support vector
The hyperplane defined by SVM is $(0.93,1.47) x-10.71$.
Then, we have implemented the Pegasos algorithm on MATLAB, the hyperplane and its margins are given in Figure (7)


Figure 7: Pegasos hyperplane and its margins

The hyperplane defined by Pegasos is $(0.94,1.46) x-10.54$.
The results we got are approximately the same as the result of the built in SVM classifier function.
Another interesting information is the minimized function value, because we used a stochastic method we wont see a stable descent, but stochastic steps that their general target is the minimum.


## 6 Conclusion

In this project, we presented and described the gradient descent algorithm. We presented two extensions of the algorithm one at the field of data fitting and one on the field of classification.
We have shown that solving optimization problems is a common task in machine learning that arises at many different applications. For large-scale problems, those methods are often the only practical option, so extending those methods with better convergence rate could yield a tremendous influence at this field. We have also been able to analyze and implement the data fitting methods and classification method in an extremely simple and efficient way.

## 7 Acknowledgements

I would like to thank Assoc. Prof. Aviv Gibali for giving me the opportunity to work on this fascinating and instructive project. I would also like to thank him for the time and help he granted me to help me move forward with this project, while pushing me to be as independent as possible.

## 8 Matlab Codes

### 8.1 Data Fitting

## FISTA Algorithm:

clear; close all; clc;

```
[A,b,m,n, Gray_Image]=Conv_Blur_Img('images.jpg');
[AT,ATA, Lipschitz,S_A] = Elements_Calc(A);
lambda = 1;
%Initial step
X0 = zeros(S_A(2),1); X02 = X0;
y0 = zeros(S_A (2),1); y02 = y0;
t0 = 1;
```

NumOfIter $=200$;

```
for i = 1:NumOfIter
    %Iterative step
    Xk = GradStep(Lipschitz, y0,AT,ATA, b);
    Xk2 = GradX1Step(Lipschitz,y02,AT,ATA,b,lambda); %
        with norm1
    tk}=0.5*(1+\operatorname{sqrt}(1+4*\textrm{t}0^2))
    yk=Xk+((t0 - 1)/tk)*(Xk-X0);
    yk2=Xk2+((t0-1)/tk)*(Xk2-X02);
    y0=yk; y02=yk2; t0=tk; X0=Xk; X02=Xk2;
    if i==5
        Img_Per_Iter(Xk, Xk2,m,n, Gray_Image, i )
        Fi(A,Xk, Xk2,b,i)
        elseif i==50
            Img_Per_Iter(Xk, Xk2,m,n, Gray_Image, i )
            Fi(A,Xk,Xk2,b,i)
        elseif i==100
            Img_Per_Iter(Xk, Xk2,m,n, Gray_Image, i )
            Fi(A,Xk,Xk2,b,i)
        end
```

end

Img_Per_Iter (Xk, Xk2, m, n, Gray_Image, i )
Fi (A, Xk, Xk2, b, i)

## ISTA Algorithm:

clear; close all; clc;

```
[A, b ,m, n, Gray_Image]=Conv_Blur_Img('images.jpg');
[AT,ATA, Lipschitz,S_A] = Elements_Calc(A);
lambda = 1;
%Initial step
X0 = zeros(S_A(2),1); X02 = X0;
NumOfIter = 200;
```

```
for i = 1:NumOfIter
    %Iterative step
    Xk = GradStep(Lipschitz,X0,AT,ATA, b);
    Xk2 = GradX1Step(Lipschitz,X02,AT,ATA,b,lambda); %
        with norm1
    X0=Xk; }\textrm{X}02=\textrm{Xk}2
    if i==5
            Img_Per_Iter(Xk, Xk2,m,n,Gray_Image, i )
            Fi(A,Xk, Xk2,b,i)
        elseif i==50
            Img_Per_Iter(Xk, Xk2,m,n, Gray_Image, i )
            Fi(A, Xk, Xk2,b,i)
        elseif i==100
            Img_Per_Iter(Xk, Xk2,m,n, Gray_Image, i )
            Fi(A,Xk, Xk2,b,i)
    end
end
```

Img_Per_Iter (Xk, Xk2, m, n, Gray_Image, i )
Fi(A, Xk, Xk2, b, i)

Image Blur:
\%This function reads an image, blur it via convolution

```
%and return this action as a system of linear equations
function [A,b,m, n,Image]=Conv_Blur_Img(image_name)
%Convert to gray image
Image=double (rgb2gray (imread (image_name)) );
%Generate a Gaussian kernel of size 9*9 std=4, and
    convulate it with the image
h = fspecial('gaussian', 9,4);
Blur_Img=conv2(h,Image);
%reshape a matrix as a vector column
m=size(Image,1);
n=size(Image,2);
p=size(Blur_Img,1);
q=size(Blur_Img,2);
b=double(reshape(Blur_Img , p*q,1));
%Generate a matrix operator of Blurr
A=convmtx2(h,m,n);
subplot(1,2,1)
imshow(uint8(Image))
title('Original Image');
subplot(1,2,2)
imshow(uint8(Blur_Img))
title('Blurred Image')
```


## Elements Calculator:

function [AT,ATA, Lipschitz, S_A] = Elements_Calc (A)
$\mathrm{AT}=\mathrm{A} .{ }^{\prime}$;
$\mathrm{ATA}=\mathrm{AT} * \mathrm{~A}$;
S_A=size (A);
Lipschitz $=2 * \operatorname{igs}(A T A, 1) ; \quad \% \% L i p=2 *$ max_eighenvalue end

## Gradient step:

function Xk=GradStep (Lipschits, X0,AT,ATA, b)
$\mathrm{Xk}=\mathrm{X} 0-(1 /$ Lipschits $) *(2 * \mathrm{ATA} * \mathrm{X} 0-2 * \mathrm{AT} * \mathrm{~b}) ;$
end

## Gradient step with norm1:

```
function Xk=GradX1Step(Lipschits,X0,AT,ATA,b,lambda)
```



```
    sign(X0));
end
```


## Show image in current iteration:

```
function Img_Per_Iter(Xk,Xk2,m,n,Image,i )
    figure
    subplot(1, 2,1);
    imshow(uint8(reshape(Xk,m,n)));
    t = sprintf('PSNR after %d ietrations: ', i);
    str1 = sprintf('|Ax-b|= %d', psnr(uint8(reshape(Xk,m
        ,n)), uint8(Image)));
    title({t;str1})
    subplot(1,2,2);
    imshow(uint8(reshape(Xk2,m,n)));
    str2 = sprintf(' |Ax-b | + |x| _ 1= %d', psnr(uint8(
        reshape(Xk2,m,n)),uint8(Image)));
    title(str2)
```

end
Value of function in current iteration:

```
function Fi(A, Xk, Xk2,b,i)
fprintf('Iteration: %d, |A*X-b| = %d',i, norm(A*Xk-b) ^2);
fprintf(', |A*X-b | + x| | 1= %d\n', norm(A*Xk2-b)`2 + norm(
    Xk2,1));
end
```


### 8.2 Classification

## main function SVM VS Pegasos:

clear; close all; clc;
data $=$ xlsread ('data');
\% define variables

$$
\begin{aligned}
& \mathrm{X}=\operatorname{data}(:, 4: 5) ; \\
& \mathrm{y}=\operatorname{data}(:, 6) ; \\
& \mathrm{T}=10000 ; \\
& \text { lambda }=0.015 ;
\end{aligned}
$$

mdl $=$ SVM_solver (X, y);
$[\mathrm{wT}, \operatorname{loss}, \mathrm{b}]=\mathrm{Peg}$ asos_solver $(\mathrm{X}, \mathrm{y}$, lambda, T$)$;
SVM solver:
function $\quad[\mathrm{mdl}]=$ SVM_solver (X, y)
\%\%\%\%\%\%\% it SVM model $1 \% \% \% \% \% \% \% \% \% \% \% \% \% \% \% \% \% \% \% \% \% \% \% \% \% \% ~$
$\mathrm{mdl}=\operatorname{fitcsvm}(\mathrm{X}, \mathrm{y})$;
$\mathrm{sv}=\mathrm{mdl}$.SupportVectors;
$\% \% \% \% \% \% \%$ presenting results of svm model $1 \% \% \% \% \% \% \%$
figure
gscatter (X (: , 1) , X (: , 2) , y)
hold on
plot (sv(:, 1) , sv (: , 2) , 'ko ', 'MarkerSize', 10)
$\mathrm{d}=0.02$;
$[\mathrm{X} 1 \mathrm{Grid}, \mathrm{X} 2 \mathrm{Grid}]=\operatorname{meshgrid}(\min (\mathrm{X}(:, 1)): \mathrm{d}: \max (\mathrm{X}(:, 1))$, $\min (\mathrm{X}(:, 2)): \mathrm{d}: \max (\mathrm{X}(:, 2)))$;
XGrid $=[\operatorname{X1Grid}(:), \mathrm{X} 2 G r i d(:)] ;$
[~, scores] $=$ predict (mdl, XGrid) ;
contour (X1Grid, X2Grid, reshape ( $\operatorname{scores}(:, 2)$, size (X1Grid)) , $\left.\left[\begin{array}{ll}0 & 0\end{array}\right],{ }^{\prime} \mathrm{k}^{\prime}\right)$;
legend ('Cupcake', 'Muffin ', 'Support Vector', 'Optimal hyperplane')
fprintf('The hyperplane defined by SVM is [(\%.2f, \%.2f)X $\% .2 \mathrm{f} \backslash \mathrm{n}$ ', mdl. Beta, mdl. Bias ) ;
end

## Pegasos solver:

```
function [w, lossval, b]=Pegasos_solver(X, y, lambda,T)
[n,~}\mp@subsup{~}{}{~}=\operatorname{size}(X)
%%%initial values%%%
w = zeros (2,1);
lossval=zeros(T,1);
```

```
b=sum (y-X*w)/n;
for t= 1:T
    eta = 1/(lambda*t);
    permi = randperm(n);
    for i = 1:n
        j=permi(i);
        if y(j)*(dot(w,X(j,:)')+b)< < 
            w = (1-eta*lambda)*w + eta*y(j)*X(j,:)';
            b}=\textrm{b}+\operatorname{eta*y(j);
        else
            w}=(1-\textrm{eta}*lambda)*\textrm{w}
        end
    end
    l=max(0,1 - (y.*(X*w)));
    lossval(t) = 0.5*lambda*w'*w + (1/n)*sum(l);
end
```

```
figure
```

figure
step = 100;
step = 100;
semilogy( 1:step:T, lossval(1:step:end), 'r-.','
semilogy( 1:step:T, lossval(1:step:end), 'r-.','
LineWidth',1.5 );
LineWidth',1.5 );
xlabel('iteration');
xlabel('iteration');
ylabel('Loss function value');
ylabel('Loss function value');
axis([1 T 0 max(lossval)]);
axis([1 T 0 max(lossval)]);
title('Convergence of Pegasos algorithm');
title('Convergence of Pegasos algorithm');
% visualize
% visualize
figure
figure
xp = linspace(min(X(:,1)), max(X(:, 1)));
xp = linspace(min(X(:,1)), max(X(:, 1)));
yp =- (w(1)*xp +b )}/\textrm{w}(2)
yp =- (w(1)*xp +b )}/\textrm{w}(2)
yp1 = - (w(1)*xp +b - 1)/w(2); % margin boundary for
yp1 = - (w(1)*xp +b - 1)/w(2); % margin boundary for
support vectors for y=1 (muffin)
support vectors for y=1 (muffin)
yp0 = - (w(1)*xp +b + 1)/ w(2); % margin boundary for
yp0 = - (w(1)*xp +b + 1)/ w(2); % margin boundary for
support vectors for y=-1 (cupcake)
support vectors for y=-1 (cupcake)
i_cupcake = find(y==-1);% index of cupcake samples
i_cupcake = find(y==-1);% index of cupcake samples
i_muffin = find (y==1);% index of muffin samples

```
i_muffin = find (y==1);% index of muffin samples
```

plot (X(i_cupcake, 1), X(i_cupcake, 2), 'ro');
hold on
plot (X(i_muffin, 1), X(i_muffin, 2), 'bo');
plot (xp, yp, '-k', xp, yp1, '--g', xp, yp0, '——'); title (sprintf('Pegasos classifier with its boundary, lambda $=\%{ }^{\prime}{ }^{\prime}$, lambda) $) ;$
fprintf('The hyperplane defined by Pegasos is (\%.2f, \%.2 f) $\left.\mathrm{X} \% .2 \mathrm{f} \backslash \mathrm{n}^{\prime}, \mathrm{w}, \mathrm{b}\right)$;
end

## References

[1] Shalev-Shwartz, Shai, et al. "Pegasos: Primal estimated sub-gradient solver for svm." Mathematical programming 127.1 (2011): 3-30.
[2] Beck, Amir, and Marc Teboulle. "A fast iterative shrinkage-thresholding algorithm for linear inverse problems." SIAM journal on imaging sciences 2.1 (2009): 183-202.
[3] Beck, Amir. Introduction to nonlinear optimization: Theory, algorithms, and applications with MATLAB. Vol. 19. Siam, 2014.
[4] https://see.stanford.edu/materials/lsoeldsee263/05-ls.pdf
[5] https://homepages.inf.ed.ac.uk/rbf/HIPR2/gsmooth.htm

