

Recklessly Approximate Sparse Coding

by

Misha Denil

BSCS, Dalhousie University, 2010

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE

in

The Faculty of Graduate Studies

(Computer Science)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

December 2012

© Misha Denil 2012

Abstract

Introduction of the so called “K-means” or “triangle” features in Coates, Lee and Ng, 2011 [10] caused significant discussion in the deep learning community. These simple features are able to achieve state of the art performance on standard image classification benchmarks, outperforming much more sophisticated methods including deep belief networks, convolutional nets, factored RBMs, mcRBMs, convolutional RBMs, sparse autoencoders and several others. Moreover, these features are extremely simple and easy to compute.

Several intuitive arguments have been put forward to describe this remarkable performance, yet no mathematical justification has been offered. In Coates and Ng, 2011 [11], the authors improve on the triangle features with “soft threshold” features, adding a hyperparameter to tune performance, and compare these features to sparse coding. Both soft thresholding and sparse coding are found to often yield similar classification results, though soft threshold features are much faster to compute.

The main result of this thesis is to show that the soft threshold features are realized as a single step of proximal gradient descent on a non-negative sparse coding objective. This result is important because it provides an explanation for the success of the soft threshold features and shows that even very approximate solutions to the sparse coding problem are sufficient to build effective classifiers.

Table of Contents

Abstract	ii
Table of Contents	iii
List of Tables	v
List of Figures	vi
1 Introduction	1
1.1 Setting	1
1.2 Background	2
1.3 Related work	3
1.4 Structure of this thesis	5
2 Optimization	6
2.1 Setting	6
2.2 Proximal gradient	7
2.3 Dual ascent	8
2.4 Method of multipliers	9
2.5 Alternating direction method of multipliers	10
3 Sparse coding	11
3.1 Setting	11
3.2 Fast iterative soft thresholding	13
3.3 Sparse reconstruction by separable approximation	13
3.4 Alternating direction method of multipliers	14
3.5 Boosted lasso	14
4 A reckless approximation	16
4.1 Main result	16
4.2 Approximate FISTA	18
4.3 Approximate SpaRSA	18

4.4	Approximate ADMM	19
4.5	Approximate BLasso	20
5	Experiments	21
5.1	Experiment 1	21
5.1.1	Procedure	21
5.1.2	Results	23
5.2	Experiment 2	24
5.2.1	Procedure	25
5.2.2	Results	25
6	Conclusion	29
	Bibliography	31

List of Tables

5.1	Parameter values for each algorithm found by optimizing for one step classification accuracy (10 steps for BLasso). Dashes indicate that the parameter is not relevant to the corresponding algorithm.	23
5.2	Test set accuracy for the of the best set of parameters found in Experiment 1 on CIFAR-10 and STL-10 using features obtained by each different algorithm.	24

List of Figures

5.1	Classification accuracy versus computation time on CIFAR-10 using different sparse encoding algorithms. For FISTA, ADMM and SpaRSA the markers show performance measured with a budget of 1, 5, 10, 50, 100 iterations. The left-most marker on each of these lines shows performance using an implementation optimized to perform exactly one step of optimization. The line for BLasso shows performance measured with a budget of 10, 50, 200, 500 iterations. In all cases early stopping is allowed if a termination criterion has been met (which causes some markers to overlap).	26
5.2	Classification accuracy versus computation time on STL-10 using different sparse encoding algorithms. For FISTA, ADMM and SpaRSA the markers show performance measured with a budget of 1, 5, 10, 50, 100 iterations. The left-most marker on each of these lines shows performance using an implementation optimized to perform exactly one step of optimization. The line for BLasso shows performance measured with a budget of 10, 50, 200, 500 iterations. In all cases early stopping is allowed if a termination criterion has been met (which causes some markers to overlap).	27
5.3	Mean reconstruction error versus computation time on a small sample of patches from CIFAR-10. FISTA, ADMM and SpaRSA were run for 100 iterations each, while BLasso was run for 500 iterations. ADMM30 corresponds to ADMM run with parameters which gave the best one-step classification performance. ADMM1 was run with a different ρ parameter which leads to better reconstruction but worse classification.	28

Chapter 1

Introduction

1.1 Setting

Image classification is one of several central problems in computer vision. This problem is concerned with sorting images into categories based on the objects or types of objects that appear in them. An important assumption that we make in this setting is that the object of interest appears prominently in each image we consider, possibly in the presence of some background “clutter” which should be ignored. The related problem of object localization, where we predict the location and extent of an object of interest in a larger image, is not considered in this thesis.

Neural networks are a common tool for this problem and have been applied in this area since at least the late 80’s [26]. More recently the introduction of contrastive divergence [18] has led to an explosion of work on neural networks to this task. Neural network models serve two purposes in this setting:

1. They provide a method to design a dictionary of primitives to use for representing images. With neural networks the dictionary can be designed through learning, and thus tailored to a specific data set.
2. They provide a method to encode images using this dictionary to obtain a feature based representation of the image.

Representations constructed in this way can be classified using a standard classifier such as a support vector machine. Properly designed feature based representations can be classified much more accurately than using the raw pixels directly. Neural networks have proved to be effective tools for constructing these representations [17, 27, 34].

A major barrier to applying these models to large images is the number of parameters required. Designing features for an $n \times n$ image using these techniques requires learning $O(n^4)$ parameters which rapidly becomes intractable, even for small images. A common solution to this problem is

to construct features for representing image patches rather than full images, and then construct feature representations of full images by combining representations of their patches.

While much effort has been devoted to designing elaborate feature learning methods [13, 14, 25, 32, 33, 35, 36], it has been shown recently that provided the dictionary is reasonable then the encoding method has a far greater effect on classification performance than the specific choice of dictionary [11].

In particular, [10] and [11] demonstrate two very simple feature encoding methods that outperform a variety of much more sophisticated techniques. In the aforementioned works, these feature encoding methods are motivated based on their computational simplicity and effectiveness. The main contribution of this thesis is to provide a connection between these features and sparse coding; in doing so we situate the work of [10] and [11] in a broader theoretical framework and offer some explanation for the success of their techniques.

1.2 Background

In [10] and [11], Coates and Ng found that two very simple feature encodings were able to achieve state of the art results on several image classification tasks. In [10] they consider encoding image patches, represented as a vector in $x \in \mathbb{R}^N$ using the so called “K-means” or “triangle” features to obtain a K -dimensional feature encoding $z^{\text{tri}}(x) \in \mathbb{R}^K$, which they define elementwise using the formula

$$z_k^{\text{tri}}(x) = \max\{0, \mu(x) - \|x - w_k\|_2\} , \quad (1.1)$$

where $\{w_k\}_{k=1}^K$ is a dictionary of elements obtained by clustering data samples with K-means, and $\mu(x)$ is the average of $\|x - w_k\|_2$ over k . In [11] they consider the closely related “soft threshold” features, given by

$$z_k^{\text{st}}(x) = \max\{0, w_k^T x - \lambda\} , \quad (1.2)$$

with $\lambda \geq 0$ as a parameter to be set by cross validation. These feature encodings have proved to be surprisingly effective, achieving state of the art results on popular image classification benchmarks. However, what makes these features especially appealing is their simplicity. Given a dictionary, producing an encoding requires only a single matrix multiply and threshold operation.

We note that the triangle and soft threshold features are merely slight variations on the same idea. If we modify the triangle features to be defined in terms of squared distances, that is we consider

$$z_k^{\text{tri}}(x) = \max\{0, \mu_2(x) - \|x - w_k\|_2^2\} ,$$

in place of Equation 1.1, with $\mu_2(x)$ taking the average value of $\|x - w_k\|_2^2$ over k , we can then write

$$\mu_2(x) - \|x - w_k\|_2^2 = 2w_k^{\text{T}}x - \frac{2}{n} \sum_{i=1}^n w_i^{\text{T}}x - w_k^{\text{T}}w_k + \frac{1}{n} \sum_{i=1}^n w_i^{\text{T}}w_i .$$

If we constrain the dictionary elements w_i to have unit norm as in [11] then the final two terms cancel and the triangle features can be rewritten as

$$z_k^{\text{tri}}(x) = 2 \max\{0, w_k^{\text{T}}x - \lambda(x)\} ,$$

which we can see is just a scaled version of soft threshold features, where the threshold,

$$\lambda(x) = \frac{1}{n} \sum_{i=1}^n w_i^{\text{T}}x ,$$

is chosen as a function of x rather than by cross validation. In this thesis we consider only the soft threshold features, but their similarity with the triangle features is suggestive.

1.3 Related work

Similar approaches to feature encoding are well known in the computer vision literature under the name of vector quantization. In this approach, data are encoded by hard assignment to the nearest dictionary element. Van Gemert et al. [38] consider a softer version of this idea and find that using a kernel function for quantization rather than hard assignment leads to better performance. Bourdeau et al. [7] consider a similar soft quantization scheme but find that sparse coding performs better still.

Following the success of [10], the triangle and soft threshold features (or slight variants thereof) have been applied in several settings. Blum et al. [6] use triangle features for encoding in their work on applying unsupervised feature learning to RGB-D data using a dictionary designed by their own convolutional K-means approach. Knoll et al. [23] apply triangle features to

image compression using PAQ. An unthresholded version of triangle features was used in [37] as the low-level image features for a system which extracts and redesigns chart images in documents.

In [29] and [9] soft threshold features are used for the detection and recognition of digits and text (respectively) in natural images. The same features have also been employed in [12] as part of the base learning model in a system for selecting the receptive fields for higher layers in a deep network.

Jia et al. [22] consider both triangle and soft threshold features for encoding image patches and investigate the effects of optimizing the spatial pooling process in order to achieve better classification accuracy.

The work most similar to that found in this thesis is the work of Gregor and LeCun on approximations of sparse coding [15]. Like us, they are interested in designing feature encoders using approximations of sparse coding; however, their technique is very different than the one we consider here.

The chief difficulty with sparse coding is that the encoding step requires solving an ℓ_1 regularized problem, which does not have a closed form solution. For example, if we want to extract features from each frame of a video in real time then sparse coding is prohibitively slow. In [15] the authors design trainable fixed cost encoders to predict the sparse coding features. The predictor is designed by taking an iterative method for solving the sparse coding problem and truncating it after a specified number of steps to give a fixed complexity feed forward predictor. The parameters of this predictor are then optimized to predict the true sparse codes over a training set.

Both our work and that of [15] is based on the idea of approximating sparse coding with a fixed dictionary by truncating an optimization before convergence, but we can identify some key differences:

- The method of [15] is trained to predict sparse codes on a particular data set with a particular dictionary. Our method requires no training and is agnostic to the specific dictionary that is used.
- We focus on the problem of creating features which lead to good classification performance directly, whereas the focus of [15] is on predicting optimal codes. Our experiments show that, at least for our approach, these two quantities are surprisingly uncorrelated.
- Although there is some evaluation of classification performance of truncated iterative solutions without learning in [15], this is only done with a coordinate descent based algorithm. Our experiments suggest that

these methods are vastly outperformed by truncated proximal methods in this setting.

Based on the above points, the work in this thesis can be seen as complementary to that of [15].

1.4 Structure of this thesis

The remainder of this thesis is structured as follows:

- In Chapter 2 we give some general background on the proximal gradient method in optimization, and introduce some extensions of this method that operate in the dual space.
- In Chapter 3 we introduce sparse coding and outline four specific algorithms for solving the encoding problem.
- Chapter 4 contains the main result of this thesis, which is a connection between the soft threshold features and sparse coding through the proximal gradient algorithm. Using this connection we outline four possible variants of the soft threshold features, based on the sparse encoding algorithms from Chapter 3.
- In Chapter 5 we report on two experiments designed to assess the usefulness of these feature variants.
- In Chapter 6 we conclude and offer suggestions for future work.

Chapter 2

Optimization

This chapter provides the necessary optimization background to support the tools used in this thesis. Focusing on objective functions with a specific separable structure, we discuss the proximal gradient algorithm with a fixed step size and an extension of this method that uses a Barzilai-Borwein step size scheme. We also consider a variant of the proximal gradient algorithm which operates in the dual space, and an alteration of that method to make it tractable for the sparse coding problem.

Throughout this chapter we eschew generality in favour of developing tools which are directly relevant to the sparse coding problem. All of the methods we discuss have more general variants, which are applicable to a broader class of problems than we are concerned with. The citations in this chapter can be used to find expositions of these methods in more general settings.

The presentation in this chapter assumes familiarity with some common optimization tools, specifically the reader should be familiar with Lagrangian methods and dualization, which are used here without justification. Extensive discussions of the supporting theory of the Lagrangian dual can be found in any standard text on convex optimization such as [4] or [5].

2.1 Setting

In what follows, we concern ourselves with the function

$$f(x) = g(x) + h(x) \quad , \quad (2.1)$$

where $g : \mathbb{R}^n \rightarrow \mathbb{R}$ is differentiable with Lipschitz derivatives,

$$\|\nabla g(x) - \nabla g(y)\|_2 \leq L\|x - y\|_2 \quad ,$$

and $h : \mathbb{R}^n \rightarrow \mathbb{R}$ is convex. In particular we are interested in cases where h is not differentiable. The primary instance of that will concern us in this thesis is

$$f(x) = \frac{1}{2}\|Wz - x\|_2^2 + \lambda\|z\|_1 \quad ,$$

where g is given by the quadratic term and h handles the (non-differentiable) ℓ_1 norm.

We search for solutions to

$$\min_x f(x) \tag{2.2}$$

and we refer to Equation 2.2 as the *unconstrained problem*. It will also be useful to us to consider an alternative formulation,

$$\min_{x=z} g(x) + h(z) , \tag{2.3}$$

which has the same solutions as Equation 2.2. Introducing the dummy variable z gives us access to the dual space which will be useful later on. We refer to this variant as the *constrained problem*.

An object of central utility in this chapter is the $\text{prox}_{h,\rho}$ operator, which is defined, for a convex function h and a scalar $\rho > 0$ as

$$\text{prox}_{h,\rho}(x) = \arg \min_u \left\{ h(u) + \frac{\rho}{2} \|u - x\|_2^2 \right\} .$$

We are often interested in the case where $\rho = 1$, and write prox_h in these cases to ease the notation.

2.2 Proximal gradient

Proximal gradient is an algorithm for solving problems with the form of Equation 2.2 by iterating

$$\begin{aligned} x^{t+1} &= \text{prox}_{\alpha h}(x^t - \alpha \nabla g(x^t)) \\ &= \arg \min_u \left\{ \alpha h(u) + \frac{1}{2} \|u - (x^t - \alpha \nabla g(x^t))\|_2^2 \right\} \end{aligned} \tag{2.4}$$

for an appropriately chosen step size α . It can be shown that if $\alpha < 1/L$ then this iteration converges to an optimal point of f [39].

Alternative step size selection methods such as line search and iterate averaging [2] are also possible. One such scheme is the Barzilai-Borwein scheme [1] which picks the step size α^t so that $\alpha^t I$ approximates the inverse Hessian of g . This choice of step size can be motivated by the form of the proximal gradient updates. We consider approximating the solution to Equation 2.2 using a quadratic model of g with diagonal covariance $\alpha^{-1}I$.

Approximating g in this way, with a Taylor expansion about an arbitrary point x^0 , leads to the problem

$$\arg \min_x \{h(x) + g(x^0) + \nabla g(x^0)^\top (x - x^0) + \frac{1}{2\alpha} \|x - x^0\|_2^2\} .$$

Some algebra reveals that we can rewrite this problem as follows:

$$\begin{aligned} & \arg \min_x \{h(x) + \nabla g(x^0)^\top (x - x^0) + \frac{1}{2\alpha} \|x - x^0\|_2^2\} \\ &= \arg \min_x \{\alpha h(x) + \alpha \nabla g(x^0)^\top (x - x^0) + \frac{1}{2} \|x - x^0\|_2^2\} \\ &= \arg \min_x \{\alpha h(x) + \frac{1}{2} x^\top x - x^\top (x^0 - \alpha \nabla g(x^0)) + \frac{1}{2} \|x^0 - \alpha \nabla g(x^0)\|_2^2\} \\ &= \arg \min_x \{\alpha h(x) + \frac{1}{2} \|x - (x^0 - \alpha \nabla g(x^0))\|_2^2\} \\ &= \text{prox}_{\alpha h}(x^0 - \alpha \nabla g(x^0)) , \end{aligned}$$

showing that the proximal gradient update in Equation 2.4 amounts to minimizing h plus a quadratic approximation of g at each step. In the above calculation we have made repeated use of the fact that adding or multiplying by a constant does not affect the location of the argmax. The Barzilai-Borwein scheme adjusts the step size α^t to ensure that the model of g we minimize is as accurate as possible. The general step size selection rule can be found in [1, 40]. We will consider a specific instance of this scheme, specialized to the sparse coding problem, in Section 3.3.

2.3 Dual ascent

We now consider the constrained problem. The constraints in Equation 2.3 allow us to form the Lagrangian,

$$L(x, z, y) = g(x) + h(z) + y^\top (x - z) ,$$

which gives us access to the dual problem,

$$\max_y q(y) = \max_y \min_{x, z} L(x, z, y) .$$

It can be shown that if y^* is a solution to the dual problem then $(x^*, z^*) = \arg \min_{x, z} L(x, z, y^*)$ is a solution to the primal problem [5]. Assuming that $q(y)$ is differentiable, this connection suggests we compute a solution to the primal problem by forming the sequence

$$y^{t+1} = y^t + \alpha \nabla q(y^t)$$

and estimate the values of the primal variables as

$$(x^{t+1}, z^{t+1}) = \arg \min_{x,z} L(x, z, y^t) . \quad (2.5)$$

Forming this estimate at each step requires no extra computation since computing the update requires, $\nabla q(y) = x^{t+1} - z^{t+1}$. In our case the minimization in Equation 2.5 splits into separate minimizations in x and z ,

$$x^{t+1} = \arg \min_x \{g(x) + (y^t)^T x\} , \quad (2.6)$$

$$z^{t+1} = \arg \min_z \{h(z) - (y^t)^T z\} . \quad (2.7)$$

This method is known as dual ascent [8], and can be shown to converge under certain conditions. Unfortunately in the sparse coding problem these conditions are not satisfied. As we see in Chapter 3, we are often interested in problems where g is minimized on an entire subspace of \mathbb{R}^n . This is problematic because if the projection of y into this subspace is non-zero then the minimization in Equation 2.6 is unbounded and $\nabla q(y)$ is not well defined.

2.4 Method of multipliers

A tool to help us work around the shortcomings of dual ascent is the Augmented Lagrangian, which is a family of functions parametrized by $\rho \geq 0$,

$$L_\rho(x, z, y) = g(x) + h(z) + y^T(x - z) + \frac{\rho}{2} \|x - z\|_2^2 .$$

The function L_ρ is the Lagrangian of the problem

$$\min_{x=z} g(x) + h(z) + \frac{\rho}{2} \|x - z\|_2^2 , \quad (2.8)$$

which we see has the same solutions Equation 2.3. The quadratic term in the Augmented Lagrangian gives the dual problem nice behaviour. The augmented dual is given by

$$\max_y q_\rho(y) = \max_y \min_{x,z} L_\rho(x, z, y) .$$

We again consider gradient ascent of the objective q_ρ by forming the sequence

$$y^{t+1} = y^t + \rho \nabla q_\rho(y^t) , \quad (2.9)$$

where $\nabla q_\rho(y^t) = x^{t+1} - z^{t+1}$ with

$$(x^{t+1}, z^{t+1}) = \arg \min_{x,z} L_\rho(x, z, y^t) . \quad (2.10)$$

This algorithm is known as the method of multipliers. The quadratic term in L_ρ ensures that $\nabla q_\rho(y)$ always exists and the algorithm is well defined. The use of ρ as the step size is motivated by the fact that it guarantees that the iterates will be dual-feasible at each step [8].

It can be shown that Equation 2.9 can be written [4]

$$\begin{aligned} y^{t+1} &= \arg \min_{\lambda} \left\{ -q(\lambda) + \frac{1}{2\rho} \|\lambda - y^t\|_2^2 \right\} \\ &= \text{prox}_{-q, 1/\rho}(y^t) , \end{aligned}$$

where q is the Lagrangian of Equation 2.3. Comparing this to Equation 2.9 shows that in the dual space, proximal ascent and gradient ascent are equivalent. The actual derivation is somewhat lengthy and is not reproduced here (but see [5] pages 244–245). To make the connection with proximal gradient explicit we can write an iterative formula for every other element of this sequence

$$\begin{aligned} y^{t+2} &= \arg \min_{\lambda} \left\{ -q(\lambda) + \frac{1}{2\rho} \|\lambda - (y^t + \rho \nabla q(y^t))\|_2^2 \right\} \\ &= \text{prox}_{-q, 1/\rho}(y^t + \rho \nabla q(y^t)) . \end{aligned}$$

2.5 Alternating direction method of multipliers

The main difficulty we encounter with the method of multipliers is that for the sparse coding problem, the joint minimization in Equation 2.10 is essentially as hard as the original problem. We can separate 2.10 into separate minimizations over x and z by doing a Gauss-Seidel pass over the two blocks instead of carrying out the minimization directly. This modification leads to the following iteration:

$$x^{t+1} = \arg \min_x L_\rho(x, z^t, y^t) , \quad (2.11)$$

$$z^{t+1} = \arg \min_z L_\rho(x^{t+1}, z, y^t) , \quad (2.12)$$

$$y^{t+1} = y^t + \rho(x^{t+1} - z^{t+1}) ,$$

where we lose the interpretation as proximal gradient on the dual function since it is no longer true that $\nabla q(y^t) \neq x^{t+1} - z^{t+1}$. However, since this method is tractable for our problems we focus on it over the method of multipliers in the following chapters.

Chapter 3

Sparse coding

3.1 Setting

Sparse coding [31] is a feature learning and encoding process, similar in many ways to the neural network based methods discussed in the Introduction. In sparse coding we construct a dictionary $\{w_k\}_{k=1}^K$ which allows us to create accurate reconstructions of input vectors x from some data set. It can be very useful to consider dictionaries that are *overcomplete*, where there are more dictionary elements than dimensions; however, in this case minimizing reconstruction error alone does not provide a unique encoding. In sparse coding uniqueness is recovered by asking the feature representation for each input to be as sparse as possible.

As with the neural network methods from the Introduction, there are two phases to sparse coding:

1. A *learning* phase, where the dictionary $\{w_k\}_{k=1}^K$ is constructed, and
2. an *encoding* phase, where we seek a representation of a new vector x in terms of elements of the dictionary.

In this thesis we focus on the encoding phase, and assume that the dictionary $\{w_k\}_{k=1}^K$ is provided to us from some external source. This focus of attention is reasonable, since it was shown experimentally in [11] that as long as the dictionary is constructed in a reasonable way¹, then it is the encoding process that has the most effect on classification performance. When we want to make it explicit that we are considering only the encoding phase we refer to the sparse encoding problem.

Formally, the sparse encoding problem can be written as an optimization. If we collect the dictionary elements into a matrix $W = [w_1 | \dots | w_K]$ and

¹ What exactly “reasonable” means in this context is an interesting question, but is beyond the scope of this thesis. The results of [11] demonstrate that a wide variety of dictionary construction methods lead to similar classification performance, but do not offer conditions on the dictionary which guarantee good performance.

denote the encoded vector by \hat{z} we can write the encoding problem as

$$\hat{z} = \arg \min_z \frac{1}{2} \|Wz - x\|_2^2 + \lambda \|z\|_1 , \quad (3.1)$$

where $\lambda \geq 0$ is a regularization parameter that represents our willingness to trade reconstruction error for sparsity. This problem fits in the framework of Chapter 2 with $g(z) = 1/2 \|Wz - x\|_2^2$ and $h(z) = \lambda \|z\|_1$.

Often it is useful to consider a non-negative version of sparse coding which leads to the same optimization as Equation 3.1 with the additional constraint that the elements of z must be non-negative. There are a few ways we can formulate this constraint but the one that will be most useful to us in the following chapters is to add an indicator function on the positive orthant to the objective in Equation 3.1,

$$\hat{z} = \arg \min_z \frac{1}{2} \|Wz - x\|_2^2 + \lambda \|z\|_1 + \Pi(z) , \quad (3.2)$$

where

$$\Pi(z) = \begin{cases} 0 & \text{if } z_i \geq 0 \quad \forall i \\ \infty & \text{otherwise .} \end{cases}$$

In most studies of sparse coding both the learning and encoding phases of the problem are considered together. In these cases, one proceeds by alternately optimizing over z and W until convergence. For a fixed z , the optimization over W in Equation 3.1 is quadratic and easily solved. The optimization over z is the same as we have presented here, but the matrix W changes in each successive optimization. Since in our setting the dictionary is fixed, we need only consider the optimization over z .

This difference in focus leads to a terminological conflict with the literature. Since sparse coding often refers to both the learning and the encoding problem together, the term “non-negative sparse coding” typically refers to a slightly different problem than Equation 3.2. In Equation 3.2 we have constrained only z to be non-negative, whereas in the literature non-negative sparse coding typically implies that both z and W are constrained to be non-negative, as is done in [19]. We cannot introduce such a constraint here, since we treat W as a fixed parameter generated by an external process.

In the remainder of this chapter we introduce four algorithms for solving the sparse encoding problem. The first three are instances of the proximal gradient framework presented in Chapter 2. The fourth algorithm is based on a very different approach to solving the sparse encoding problem that works by tracking solutions as the regularization parameter varies.

3.2 Fast iterative soft thresholding

Iterative soft thresholding (ISTA) [2] is the name given to the proximal gradient algorithm with a fixed step size when applied to problems of the form of Equation 2.2, when the non-smooth part h is proportional to $\|x\|_1$. The name iterative soft thresholding arises because the proximal operator of the ℓ_1 norm is given by the soft threshold function (see [28] § 13.4.3.1):

$$\text{prox}_{\lambda\|\cdot\|_1}(x) = \text{soft}_\lambda(x) = \text{sign}(x) \max\{0, |x| - \lambda\} .$$

In the case of sparse encoding this leads to iterations of the form

$$z^{t+1} = \text{soft}_{\lambda/L}(z^t - \frac{1}{L}W^T(Wz^t - x)) .$$

The constant L here is the Lipschitz constant referred to in the statement of Equation 2.1 which, in the case of sparse coding, is the largest eigenvalue of $W^T W$. The “Fast” variant of iterative soft thresholding (FISTA) modifies the above iteration to include a specially chosen momentum term, leading to the following iteration, starting with $y^t = z^0$ and $k^1 = 1$:

$$z^t = \text{soft}_{\lambda/L}(y^t) , \tag{3.3}$$

$$k^{t+1} = \frac{1 + \sqrt{1 + 4(k^t)^2}}{2} , \tag{3.4}$$

$$y^{t+1} = z^t + \left(\frac{k^t - 1}{k^{t+1}}\right)(z^t - z^{t-1}) . \tag{3.5}$$

The form of the updates in FISTA is not intuitive, but can be shown to lead to a faster convergence rate than regular ISTA [2].

3.3 Sparse reconstruction by separable approximation

Sparse reconstruction by separable approximation (SpaRSA) [40] is an optimization framework designed for handling problems of the form considered in Chapter 2. This framework actually subsumes the ISTA and FISTA style algorithms discussed above, but we consider a specific instantiation of this framework which sets the step size using a Barzilai-Borwein [1] scheme, making it different from the methods described above. The development in [40]

discusses SpaRSA in its full generality, but specializing it to the sparse coding problem we get the following iteration:

$$\begin{aligned} z^{t+1} &= \text{soft}_{\lambda/\alpha^t}(x^t - \frac{1}{\alpha^t}W^T(Wz^t - x)) \ , \\ s^{t+1} &= z^{t+1} - z^t \ , \\ \alpha^{t+1} &= \frac{\|Ws^{t+1}\|_2^2}{\|s^{t+1}\|_2^2} \ . \end{aligned}$$

The SpaRSA family of algorithms shares an important feature with other Barzilai-Borwein methods, namely that it does not guarantee a reduction in the objective value at each step. In fact, it has been observed that forcing these methods to descend at each step (for example, by using a backtracking line search) can significantly degrade performance in practice [40]. In order to guarantee convergence of this type of scheme a common approach is to force the iterates to be no larger than the largest objective value in some fixed time window. This approach allows the objective value to occasionally increase, while still ensuring that the iterates converge in the limit.

3.4 Alternating direction method of multipliers

The alternating direction method of multipliers (ADMM) [8] was presented in Chapter 2. Its instantiation for the sparse encoding problem does not have its own name in the literature, but it can be applied nonetheless.

For the sparse encoding problem the minimizations in Equations 2.11 and 2.12 can be carried out in closed form. This leads to the updates:

$$\begin{aligned} x^{k+1} &= (W^TW + \rho I)^{-1}(W^Tx - \rho(z^k - \frac{1}{\rho}y^k)) \ , \\ z^{k+1} &= \text{soft}_{\lambda/\rho}(x^{k+1} + \frac{1}{\rho}y^k) \ , \\ y^{k+1} &= y^k + \rho(x^{k+1} - z^{k+1}) \ . \end{aligned}$$

3.5 Boosted lasso

Boosted Lasso (BLasso) [41] is a very different approach to solving the sparse encoding problem than those considered above. Rather than solving Equation 3.1 directly, BLasso works with an alternative formulation of the sparse

encoding problem,

$$\begin{aligned} \hat{z} &= \arg \min_z \frac{1}{2} \|Wz - x\|_2^2 \\ \text{st } & \|z\|_1 \leq \beta . \end{aligned} \tag{3.6}$$

For each value of λ in Equation 3.1 there is a corresponding value of β which causes Equation 3.6 to have the same solution, although the mapping between values of λ and β is problem dependent. BLasso works by varying the value of β and maintaining a corresponding solution to Equation 3.6 at each step.

As the name suggests BLasso draws on the theory of Boosting, which can be cast as a problem of functional gradient descent on the mixture parameters of an additive model composed of weak learners. In this setting the weak learners are elements of the dictionary and their mixing parameters are found in z . Similarly to the algorithms considered above, BLasso starts from the fully sparse solution but instead of applying proximal iterations, it proceeds by taking two types of steps: *forward* steps, which decrease the quadratic term in Equation 3.6 and *backward* steps which decrease the regularizer. In truth, BLasso also only gives exact solutions to Equation 3.6 in the limit as the step size $\epsilon \rightarrow 0$; however, setting ϵ small enough can force the BLasso solutions to be arbitrarily close to exact solutions to Equation 3.6.

BLasso can be used to optimize an arbitrary convex loss function with a convex regularizer; however, in the case of sparse coding the forward and backward steps are especially simple. This simplicity means that each iteration of BLasso is much cheaper than a single iteration of the other methods we consider, although this advantage is reduced by the fact that several iterations of BLasso are required to produce reasonable encodings. Another disadvantage of BLasso is that it cannot be easily cast in a way that allows multiple encodings to be computed simultaneously.

Chapter 4

A reckless approximation

4.1 Main result

In this chapter we present the main result of this thesis, which is a connection between the soft threshold features discussed in the Introduction, and the sparse encoding problem. Our key insight is to show how the soft threshold features, (as defined in Equation 1.2) can be viewed as an approximate solution to the non-negative sparse encoding problem (Equation 3.2).

We illustrate this connection through the framework of proximal gradient minimization. Using the tools presented in Chapter 2 we can demonstrate this connection by writing down a proximal gradient iteration for the sparse encoding problem and computing the value of the first iterate, starting from an appropriately chosen initial point. We summarize this result in a Proposition.

Proposition 1. *The soft threshold features*

$$z_k(x) = \max\{0, w_k^T x - \lambda\}$$

are given by a single step (of size 1) of proximal gradient descent on the non-negative sparse coding objective with regularization parameter λ and known dictionary W , starting from the fully sparse solution.

Proof. Casting the non-negative sparse coding problem in the framework of Chapter 2 we have

$$\min_z f(z) + g(z)$$

with

$$g(z) = \frac{1}{2} \|Wz - x\|_2^2, \\ h(z) = \lambda \|z\|_1 + \Pi(z).$$

The proximal gradient iteration for this problem (with $\alpha = 1$) is

$$z^{t+1} = \text{prox}_h(z^t - W^T(Wz^t - x)).$$

We now compute $\text{prox}_h(x)$

$$u^* = \text{prox}_h(x) = \arg \min_u \Pi(u) + \lambda \|u\|_1 + \frac{1}{2} \|u - x\|_2^2 .$$

This minimization is separable, and we can write down the solution for each element of the result independently:

$$u_k^* = \arg \min_{u_k \geq 0} \lambda u_k + \frac{1}{2} (u_k - x_k)^2 . \quad (4.1)$$

Each minimization is quadratic in u_k , and therefore the optimum of Equation 4.1 is given by $u_k^* = \max\{0, u_k^*\}$ where $u_k^* = x_k - \lambda$ is the unconstrained optimum. We set $z^0 = 0$ and compute

$$\begin{aligned} z_k^1 &= \text{prox}_h(z_k^0 - w_k^T(Wz^0 - x)) \\ &= \max\{0, w_k^T x - \lambda\} , \end{aligned}$$

which is the desired result. \square

Variants of the soft threshold features that appear in the literature can be obtained by slight modifications of this argument. For example, the split encoding used in [11] can be obtained by setting $W = \{w_k, -w_k\}$.

Once stated the proof of Proposition 1 is nearly immediate; however, this immediacy only appears in hindsight. In [11], soft threshold features and sparse coding features are treated as two separate and competing entities (see, for example, Figure 1 in [11]). In [30], triangle features and sparse coding are treated as two separate sparsity inducing objects.

From this Proposition we can draw two important insights:

1. Proposition 1 provides a nice explanation for the success of soft threshold features for classification. Sparse coding is a well studied problem and it is widely known that the features from sparse coding models are effective for classification tasks.
2. On the other hand, Proposition 1 tells us that even very approximate solutions to the sparse coding problem are sufficient to build effective classifiers.

Even optimizers specially designed for the sparse coding problem typically take many iterations to converge to a solution with low reconstruction error, yet here we see that a single iteration of proximal gradient descent is sufficient to give features which have been shown to be highly discriminative.

These insights open up three questions:

1. Is it possible to decrease classification error by doing a few more iterations of proximal descent?
2. Do different optimization methods for sparse encoding lead to different trade-offs between classification accuracy and computation time?
3. To what extent is high reconstruction accuracy a prerequisite to high classification accuracy using features obtained in this way?

We investigate the answers to these questions experimentally in Chapter 5, by examining how variants of the soft threshold features perform. We develop these variants by truncating other proximal descent based optimization algorithms for the sparse coding problem. The remainder of this chapter presents “one-step” features from each of the algorithms presented in Chapter 3.

4.2 Approximate FISTA

Fast iterative soft thresholding was described in Section 3.2. The first iteration of FISTA is a step of ordinary proximal gradient descent since the FISTA iteration (Equations 3.3, 3.4 and 3.5) requires two iterates to adjust the step size.

Starting from $z^0 = 0$ the one step FISTA features are given by

$$z^1 = \text{soft}_{\lambda/L} \left(\frac{1}{L} W^T x \right) = \frac{1}{L} \text{soft}_{\lambda} (W^T x) ,$$

which we see is equivalent to the soft threshold features, scaled by a factor of $1/L$.

4.3 Approximate SpaRSA

Sparse reconstruction by separable approximation was described in Section 3.3. SpaRSA, like FISTA, is an adaptive step size selection scheme for proximal gradient descent which chooses its step size based on the previous two iterates. For the first iteration this information is not available, and the authors of [40] suggest a step size of 1 in this case. This choice gives the following formula for one step SpaRSA features:

$$z^1 = \text{soft}_{\lambda} (W^T x) ,$$

which is exactly equivalent to the soft threshold features.

4.4 Approximate ADMM

The alternating direction method of multipliers was described in Section 3.4. Since ADMM operates in the dual space, computing its iterations requires choosing a starting value for the dual variable y^0 . Since this choice is essentially arbitrary (the value of y^0 does not effect the convergence of ADMM) we choose $y^0 = 0$ for simplicity. This leads to one step ADMM features of the following form:

$$z^1 = \text{soft}_{\lambda/\rho}((W^T W + \rho I)^{-1} W^T x) . \quad (4.2)$$

The parameter ρ in the above expression is the penalty parameter from ADMM. As long as we are only interested in taking a single step of this optimization, the matrix $(W^T W + \rho I)^{-1} W^T$ can be precomputed and the encoding cost for ADMM is the same as for soft threshold features. Unfortunately, if we want to perform more iterations of ADMM then we are forced to solve a new linear system in $W^T W + \rho I$ at each iteration, although we can cache an appropriate factorization in order to avoid the full inversion at each step.

The choice of $y^0 = 0$ allows us to make some interesting connections between the one step ADMM features and some other optimization problems. For instance, if we consider a second order variant of proximal gradient (by adding a Newton term to Equation 2.4) we get the following one step features for the sparse encoding problem:

$$z^1 = \text{soft}_{\lambda}((W^T W)^{-1} W^T x) .$$

We can thus interpret the one step ADMM features as a smoothed step of a proximal version of Newton's method. The smoothing in the ADMM features is important because in typical sparse coding problems the dictionary W is overcomplete and thus $W^T W$ is rank deficient. Although it is possible to replace the inverse of $W^T W$ in the above expression with its pseudoinverse we found this to be numerically unstable. The ADMM iteration smooths this inverse with a ridge term and recovers stability.

Taking a slightly different view of Equation 4.2, we can interpret it as a single proximal Newton step on the Elastic Net objective [42],

$$\arg \min_z \frac{1}{2} \|Wz - x\|_2^2 + \rho \|z\|_2^2 + \frac{\lambda}{\rho} \|z\|_1 .$$

Here the ADMM parameter ρ trades off the magnitude of the ℓ_2 term, which smooths the inverse, and the ℓ_1 term, which encourages sparsity.

4.5 Approximate BLasso

BLasso was described briefly in Section 3.5. Unlike the other algorithms we consider, each iteration of BLasso updates exactly one element of the feature vector z . This means that taking one step of BLasso leads to a feature vector with exactly one non-zero element (of magnitude ϵ). In contrast, the proximal methods described above update all of the elements of z at each iteration, meaning that the one step features can be arbitrarily dense.

This difference suggests that comparing one step features from the other algorithms to one step features from BLasso may not be a fair comparison. To accommodate this, in the sequel we use many more steps of BLasso than the other algorithms in our comparisons.

Writing out the one (or more) step features for BLasso is somewhat notationally cumbersome so we omit it here, but the form of the iterations can be found in [41].

Chapter 5

Experiments

5.1 Experiment 1

We evaluate classification performance using features obtained by approximately solving the sparse coding problem using the different optimization methods discussed in Chapter 3. We report the results of classifying the CIFAR-10² and STL-10³ data sets using an experimental framework similar to [10]. The images in STL-10 are 96×96 pixels, but we scale them to 32×32 to match the size of CIFAR-10 in all of our experiments. For each different optimization algorithm we produce features by running different numbers of iterations and examine the effect on classification accuracy.

5.1.1 Procedure

Training

During the training phase we produce a candidate dictionary W for use in the sparse encoding problem. We found the following procedure to give the best performance:

1. Extract a large library of 6×6 patches from the training set. Normalize each patch by subtracting the mean and dividing by the standard deviation, and whiten the entire library using ZCA [3].
2. Run K-means with 1600 centroids on this library to produce a dictionary for sparse coding.

We experimented with other methods for constructing dictionaries as well, including using a dictionary built by omitting the K-means step above and using whitened patches directly. We also considered a dictionary of normalized random noise, as well as a smoothed version of random noise obtained by convolving noise features with a Gaussian filter. However, we found that

²<http://www.cs.toronto.edu/~kriz/cifar.html>

³<http://www.stanford.edu/~acoates/stl10/>

the dictionary created using whitened patches and K-means together gave uniformly better performance, so we report results only for this choice of dictionary. An extensive comparison of different dictionary choices appears in [11].

Testing

In the testing phase we build a representation for each image in the CIFAR-10 and STL-10 data sets using the dictionary obtained during training. We build representations for each image using patches as follows:

1. Extract 6×6 patches densely from each image and whiten using the ZCA parameters found during training.
2. Encode each whitened patch using the dictionary found during training by running one or more iterations of each of the algorithms from Chapter 3.
3. For each image, pool the encoded patches in a 2×2 grid, giving a representation with 6400 features for each image.
4. Train a linear classifier to predict the class label from these representations.

This procedure involves approximately solving Equation 3.1 for each patch of each image of each data set, requiring the solution to just over 4×10^7 separate sparse encoding problems to encode CIFAR-10 alone, and is repeated for each number of iterations for each algorithm we consider. Since iterations of the different algorithms have different computational complexity, we compare classification accuracy against the time required to produce the encoded features rather than against number of iterations.

We performed the above procedure using features obtained with non-negative sparse coding as well as with regular sparse coding, but found that projecting the features into the positive orthant always gives better performance, so all of our reported results use features obtained in this way.

Parameter selection is performed separately for each algorithm and data set. For each algorithm we select both the algorithm specific parameters, ρ for ADMM and ϵ for BLasso, as well as λ in Equation 3.1, in order to maximize classification accuracy using features obtained from a single step of optimization.⁴ The parameter values we used in this experiment are shown in Table 5.1.

⁴For BLasso we optimized classification after 10 steps instead of 1 step, since 1 step of BLasso produced features with extremely poor performance.

Algorithm	λ	ρ	ϵ
BLasso	–	–	0.25
FISTA	0.1	–	–
ADMM	0.02	30	–
SPARSA	0.1	–	–

Table 5.1: Parameter values for each algorithm found by optimizing for one step classification accuracy (10 steps for BLasso). Dashes indicate that the parameter is not relevant to the corresponding algorithm.

5.1.2 Results

The results of this experiment on CIFAR-10 are summarized in Figure 5.1 and the corresponding results on STL-10 are shown in Figure 5.2. The results are similar for both data sets; the discussion below applies to both CIFAR-10 and STL-10.

The first notable feature of these results is that BLasso leads to features which give relatively poor performance. Although approximate BLasso is able to find exact solutions to Equation 3.1, running this algorithm for a limited number of iterations means that the regularization is very strong. We also see that for small numbers of iterations the performance of features obtained with FISTA, ADMM and SpaRSA are nearly identical. Table 5.2 shows the highest accuracy obtained with each algorithm on each data set over all runs.

Another interesting feature of BLasso performance is that, when optimized to produce one-step features, the other algorithms are significantly faster than 10 iterations of BLasso. This is unexpected, because the iterations of BLasso have very low complexity (much lower than the matrix-vector multiply required by the other algorithms).

The reason the BLasso features take longer to compute comes from the fact that it is not possible to vectorize BLasso iterations across different problems. To solve many sparse encoding problems simultaneously one can replace the vectors z and x in Equation 3.1 with matrices Z and X containing the corresponding vectors for several problems aggregated into columns. We can see from the form of the updates described in Chapter 3, that we can replace z and x with Z and X without affecting the solution for each problem. This allows us to take advantage of optimized BLAS libraries to compute the matrix-matrix multiplications required to solve these problems in batch. It is not possible to take advantage of this optimization with BLasso.

The most notable feature of Figures 5.1 and 5.2 is that for large numbers of iterations the performance of FISTA and SpaRSA actually drops below what we see with a single iteration, which at first blush seems obviously wrong. It is important to interpret the implications of these plots carefully. In these figures all parameters were chosen to optimize classification performance for one-step features. There is no particular reason one should expect the same parameters to also lead to optimal performance after many iterations. We have found that the parameters one obtains when optimizing for one-step classification performance are generally not the same as the parameters one gets by optimizing for performance after the optimization has converged. It should also be noted that the parameter constellations we found in Table 5.1 universally have a very small λ value. This contributes to the drop in accuracy we see with FISTA especially, since this algorithm becomes unstable after many iterations with a small λ .

This observation is consistent with the experiments in [11] which found different optimal values for the sparse coding regularizer and the threshold parameter in the soft threshold features in their comparisons. It should also be stated that, as reported in [11], the performance of soft threshold features and sparse coding is often not significantly different. We refer the reader to the above cited work for a discussion of the factors governing these differences.

CIFAR10	Accuracy	STL10	Accuracy
BLasso	66.7	BLasso	47.0
FISTA	77.5	FISTA	62.8
ADMM	77.3	ADMM	63.2
SpaRSA	77.5	SpaRSA	62.5

Table 5.2: Test set accuracy for the of the best set of parameters found in Experiment 1 on CIFAR-10 and STL-10 using features obtained by each different algorithm.

5.2 Experiment 2

This experiment is designed to answer our third question from Chapter 4, relating to the relationship between reconstruction and classification accuracy. In this experiment we measure the reconstruction accuracy of encodings obtained in the previous experiment.

5.2.1 Procedure

Training

The encoding dictionary used for this experiment was constructed using the method described in Section 5.1.1. In order to ensure that our results here are comparable to the previous experiment we actually use the same dictionary in both.

Testing

In order to measure reconstruction accuracy we use the following procedure:

1. Extract a small library of 6×6 patches from randomly chosen images in the CIFAR-10 training set and whiten using the ZCA parameters found during training.
2. Encode each whitened patch using the dictionary found during training by running one or more iterations of each of the algorithms from Chapter 3.
3. For each of the encoded patches, we measure the reconstruction error $\|Wz - x\|_2$ and report the mean.

5.2.2 Results

The results of this experiment are shown in Figure 5.3. Comparing these results to the previous experiment, we see that there is surprisingly little correlation between the reconstruction error and classification performance. In Figure 5.1 we saw one-step features give the best classification performance of all methods considered, here we see that these features also lead to the worst reconstruction.

Another interesting feature of this experiment is that the parameters we found to give the best features for ADMM actually lead to an optimizer which makes no progress in reconstruction beyond the first iteration. To confirm that this is not merely an artifact of our implementation we have also included the reconstruction error from ADMM run with an alternative setting of $\rho = 1$ which gives the lowest reconstruction error of any of our tested methods, while producing inferior performance in classification.

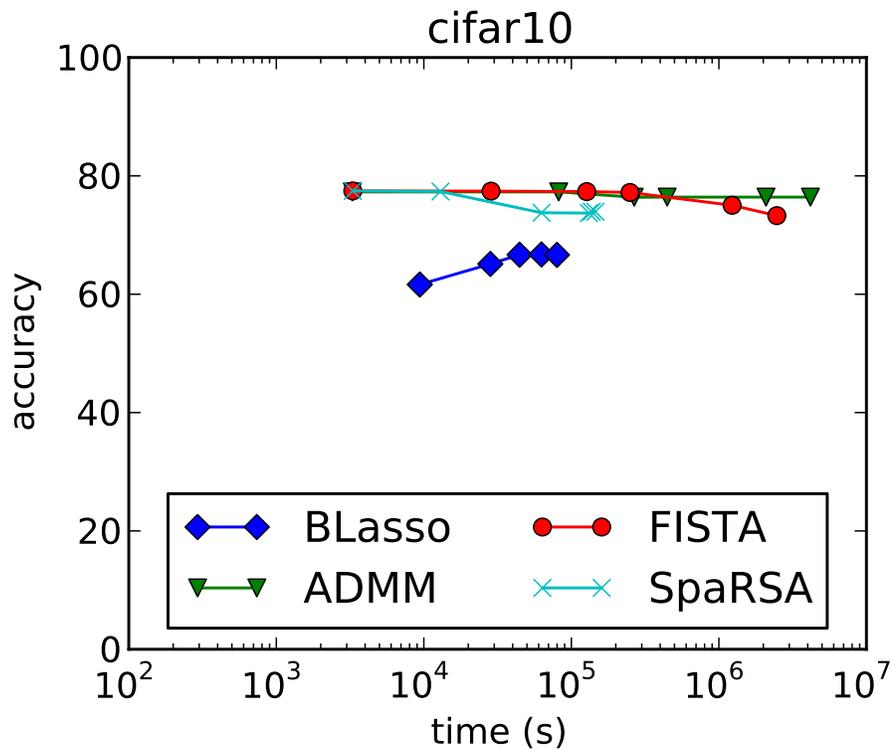


Figure 5.1: Classification accuracy versus computation time on CIFAR-10 using different sparse encoding algorithms. For FISTA, ADMM and SpaRSA the markers show performance measured with a budget of 1, 5, 10, 50, 100 iterations. The left-most marker on each of these lines shows performance using an implementation optimized to perform exactly one step of optimization. The line for BLasso shows performance measured with a budget of 10, 50, 200, 500 iterations. In all cases early stopping is allowed if a termination criterion has been met (which causes some markers to overlap).

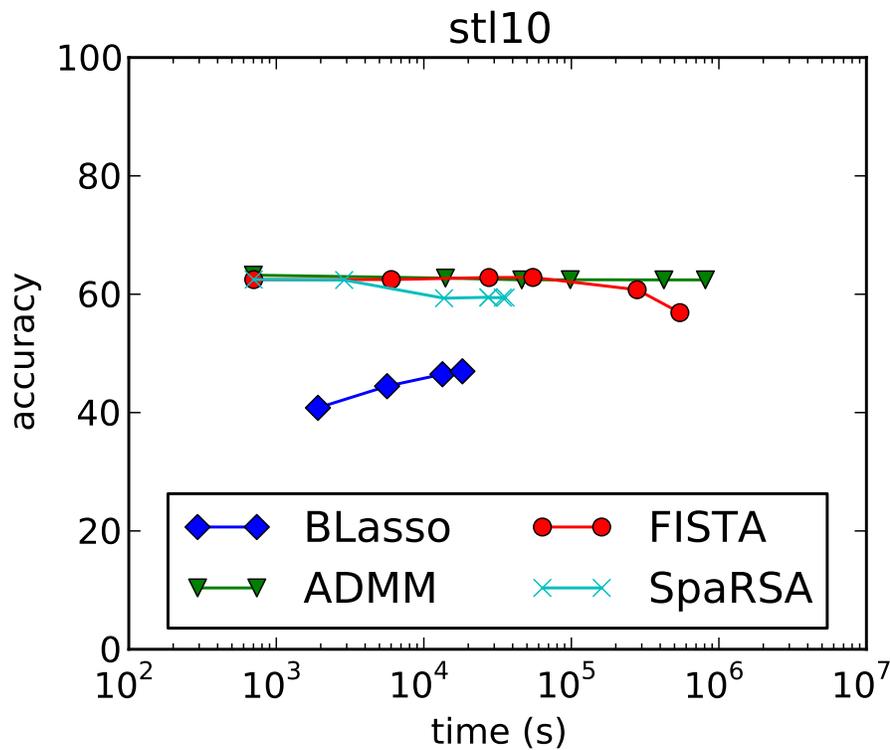


Figure 5.2: Classification accuracy versus computation time on STL-10 using different sparse encoding algorithms. For FISTA, ADMM and SpaRSA the markers show performance measured with a budget of 1, 5, 10, 50, 100 iterations. The left-most marker on each of these lines shows performance using an implementation optimized to perform exactly one step of optimization. The line for BLasso shows performance measured with a budget of 10, 50, 200, 500 iterations. In all cases early stopping is allowed if a termination criterion has been met (which causes some markers to overlap).

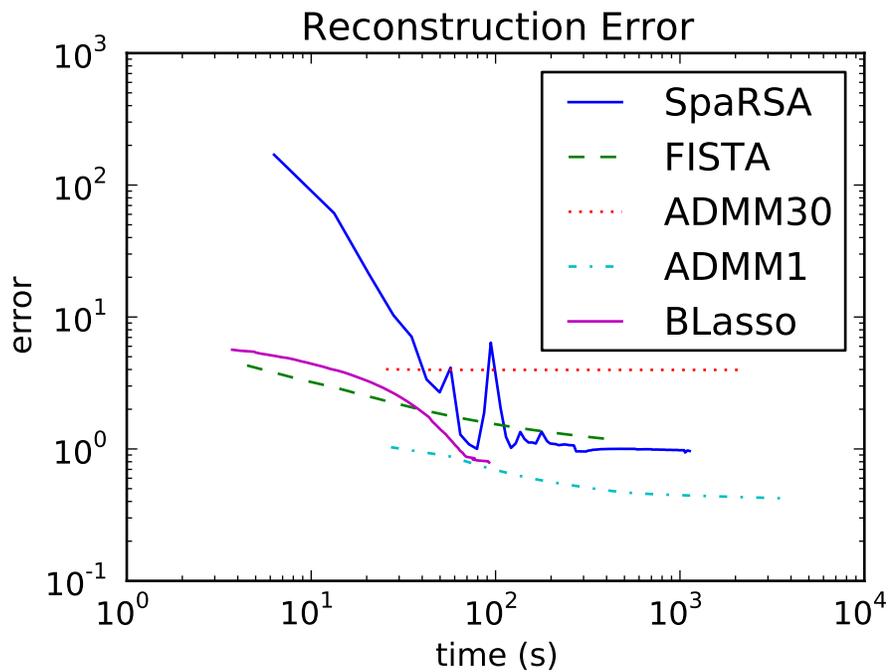


Figure 5.3: Mean reconstruction error versus computation time on a small sample of patches from CIFAR-10. FISTA, ADMM and SpaRSA were run for 100 iterations each, while BLasso was run for 500 iterations. ADMM30 corresponds to ADMM run with parameters which gave the best one-step classification performance. ADMM1 was run with a different ρ parameter which leads to better reconstruction but worse classification.

Chapter 6

Conclusion

In this thesis we have shown that the soft threshold features, which have enjoyed much success recently, arise as a single step of proximal gradient descent on a non-negative sparse encoding objective. This result serves to situate this surprisingly successful feature encoding method in a broader theoretical framework.

Using this connection we proposed four alternative feature encoding methods based on approximate solutions to the sparse encoding problem. Our experiments demonstrate that the approximate proximal-based encoding methods all lead to feature representations with very similar performance on two image classification benchmarks.

The sparse encoding objective is based around minimizing the error in reconstructing an image patch using a linear combination of dictionary elements. Given the degree of approximation in our techniques, one would not expect the features we find to lead to accurate reconstructions. Our second experiment demonstrates that this intuition is correct.

An obvious extension of this work would be to perform a more thorough empirical exploration of the interaction between the value of the regularization parameter λ and the degree of approximation. From our experimentation we can see only that such an interaction exists, but not glean insight into its structure. Some concrete suggestions along this line are:

1. Perform a full parameter search for multi-step feature encodings.
2. Evaluate the variation of performance across different dictionaries.

A full parameter search would make it possible to properly assess the usefulness of performing more than one iteration of optimization when constructing a feature encoding. In this thesis we evaluate the performance of different feature encoding methods using a single dictionary; examining the variability in performance across multiple dictionaries would lend more credibility to the results we reported in Chapter 5.

Another interesting direction for future work on this problem is an investigation of the effects of different regularizers on approximate solutions

to the sparse encoding problem. The addition of an indicator function to the regularizer in Equation 3.2 appears essential to good performance and proximal methods, which we found to be very effective in this setting, work by adding a quadratic smoothing term to the objective function at each step. The connection between one-step ADMM and the Elastic Net is also notable in this regard. Understanding the effects of different regularizers empirically, or better yet having a theoretical framework for reasoning about the effects of different regularizers in this setting, would be quite valuable.

In this work we have looked only at unstructured variants of sparse coding. It may be possible to extend the ideas presented here to the structured case, where the regularizer includes structure inducing terms [16]. Some potential launching points for this are [20] and [21], where the authors investigate proximal optimization methods for structured variants of the sparse coding problem.

Bibliography

- [1] J. Barzilai and J. Borwein. Two-point step size gradient methods. *IMA Journal of Numerical Analysis*, 8, 1988.
- [2] A. Beck and M. Teboulle. A fast iterative shrinkage-threshold algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, pages 183–202, 2009.
- [3] A. Bell and T. Sejnowski. The “independent components” of natural scenes are edge filters. *Vision Research*, 37(23):3327–3338, 1997.
- [4] D. Bertsekas. *Nonlinear Programming: Second Edition*. Athena Scientific, 1995.
- [5] D. Bertsekas and J. Tsitsiklis. *Parallel and distributed computation: numerical methods*. Prentice-Hall, Inc., 1989.
- [6] M. Blum, J. Springenberg, Wülfing J., and M. Riedmiller. On the applicability of unsupervised feature learning for object recognition in RGB-D data. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [7] Y-L. Boureau, F. Bach, LeCun Y., and J. Ponce. Learning mid-level features for recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2010.
- [8] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- [9] A. Coates, B. Carpenter, C. Case, S. Satheesh, B. Suresh, T. Wang, and A. Ng. Text detection and character recognition in scene images with unsupervised feature learning. In *Proceedings of the 11th International Conference on Document Analysis and Recognition*, pages 440–445. IEEE, 2011.

- [10] A. Coates, H. Lee, and A. Ng. An analysis of single-layer networks in unsupervised feature learning. In *International Conference on Artificial Intelligence and Statistics*, 2011.
- [11] A. Coates and A. Ng. The importance of encoding versus training with sparse coding and vector quantization. In *International Conference on Machine Learning*, 2011.
- [12] A. Coates and A. Ng. Selecting receptive fields in deep networks. In J. Shawe-Taylor, R.S. Zemel, P. Bartlett, F.C.N. Pereira, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 2528–2536. 2011.
- [13] A. Courville, J. Bergstra, and Y. Bengio. A spike and slab restricted Boltzmann machine. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15, 2011.
- [14] I. Goodfellow, A. Courville, and Y. Bengio. Large-scale feature learning with spike-and-slab sparse coding. In *International Conference on Machine Learning*, 2012.
- [15] K. Gregor and Y. LeCun. Learning fast approximations of sparse coding. In *Proceedings of the International Conference on Machine Learning*, 2010.
- [16] K. Gregor, A. Szlam, and Y. LeCun. Structured sparse coding via lateral inhibition. In *Advances in Neural Information Processing Systems*, volume 24, 2011.
- [17] G. Hinton. Learning to represent visual input. *Philosophical Transactions of the Royal Society, B*, 365:177–184, 2010.
- [18] G. Hinton, S. Osindero, and Y. Teh. A fast learning algorithm for deep belief nets. *Neural Computation*, 18:1527–1554, 2006.
- [19] P. Hoyer. Non-negative sparse coding. In *IEEE Workshop on Neural Networks for Signal Processing*, pages 557–565, 2002.
- [20] R. Jenatton, J. Mairal, G. Obozinski, and Bach F. Proximal methods for sparse hierarchical dictionary learning. In *Proceedings of the International Conference on Machine Learning*, 2010.
- [21] R. Jenatton, J. Mairal, G. Obozinski, and Bach F. Proximal methods for hierarchical sparse coding. *Journal of Machine Learning Research*, 12:2297–2334, 2011.

- [22] T. Jia, C. Huang, and T. Darrell. Beyond spatial pyramids: Receptive field learning for pooled image features. In *Computer Vision and Pattern Recognition*, 2012.
- [23] B. Knoll and N. de Freitas. A machine learning perspective on predictive coding with PAQ8. In *Data Compression Conference*, pages 377–386, 2012.
- [24] A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [25] Q. Le, M. Ranzato, R. Monga, M. Devin, G. Corrado, K. Chen, Dean J., and A. Ng. Building high-level features using large scale unsupervised learning. In *International Conference of Machine Learning*, 2012.
- [26] Y. LeCun, L. Jackel, B. Boser, J. Denker, H. Graf, I. Guyon, D. Henderson, R. Howard, and W. Hubbard. Handwritten digit recognition: Applications of neural net chips and automatic learning. In F. Fogelman, J. Herault, and Y. Burnod, editors, *Neurocomputing, Algorithms, Architectures and Applications*, Les Arcs, France, 1989. Springer.
- [27] V. Mnih and G. Hinton. Learning to detect roads in high-resolution aerial images. In *European Conference on Computer Vision*, 2010.
- [28] K. Murphy. *Machine Learning a Probabilistic Perspective*. MIT Press, 2012.
- [29] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [30] J. Ngiam, P. Koh, Z. Chen, S. Bhaskar, and A. Ng. Sparse filtering. In J. Shawe-Taylor, R.S. Zemel, P. Bartlett, F.C.N. Pereira, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 1125–1133. 2011.
- [31] B. Olshausen and D. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, 1996.
- [32] M. Ranzato and G. Hinton. Modelling pixel means and covariances using factored third-order boltzmann machines. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2010.

- [33] M. Ranzato, A. Krizhevsky, and Hinton G. Factored 3-way restricted boltzmann machines for modelling natural images. In *International Conference on Artificial Intelligence and Statistics*, 2010.
- [34] M. Ranzato, J. Susskind, V. Mnih, and G. Hinton. On deep generative models with applications to recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2011.
- [35] S. Rifai, Y. Dauphin, P. Vincent, Y. Bengio, and X. Muller. The manifold tangent classifier. In *NIPS'2011*, 2011. Student paper award.
- [36] S. Rifai, P. Vincent, X. Muller, X. Glorot, and Y. Bengio. Contracting auto-encoders: Explicit invariance during feature extraction. In *Proceedings of the Twenty-eight International Conference on Machine Learning (ICML'11)*, June 2011.
- [37] M. Savva, N. Kong, A. Chhajta, L. Fei-Fei, M. Agrawala, and J. Heer. ReVision: Automated classification, analysis and redesign of chart images. In *Proceedings of the 24th annual Symposium on User Interface Software and Technology*, 2011.
- [38] J. van Gemert, J. Geusebroek, C. Veenman, and A. Smeulders. Kernel codebooks for scene categorization. *Computer Vision–ECCV 2008*, pages 696–709, 2008.
- [39] L. Vandenberghe. EE236c – optimization methods for large-scale systems, 2011.
- [40] S. Wright, R. Nowak, and M. Figueiredo. Sparse reconstruction by separable approximation. *IEEE Transactions on Signal Processing*, 57(7):2479–2493, 2009.
- [41] P. Zhao and B. Yu. Boosted lasso. *Journal of Machine Learning Research*, 8(Dec):2701–2726, 2007.
- [42] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B*, 67(2):301–320, 2005.