

# Convex Coding

David M. Bradley, J. Andrew Bagnell

CMU-RI-TR-09-22

May, 2009

Robotics Institute  
Carnegie Mellon University  
Pittsburgh, Pennsylvania 15213

© Carnegie Mellon University

### Abstract

Inspired by recent work on convex formulations of clustering (Lashkari & Golland, 2008; Nowozin & Bakir, 2008) we investigate a new formulation of the *Sparse Coding Problem* (Olshausen & Field, 1997). In sparse coding we attempt to simultaneously represent a sequence of data-vectors sparsely (i.e. *sparse approximation* (Tropp et al., 2006)) in terms of a “code” defined by a set of basis elements, while also finding a code that enables such an approximation. As existing alternating optimization procedures for sparse coding are theoretically prone to severe local minima problems, we propose a convex relaxation of the sparse coding problem and derive a boosting-style algorithm, that (Nowozin & Bakir, 2008) serves as a convex “master problem” which calls a (potentially non-convex) sub-problem to identify the next code element to add. Finally, we demonstrate the properties of our boosted coding algorithm on an image denoising task.

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Sparse Coding</b>	<b>2</b>
2.1	Notation . . . . .	2
2.2	Generative Model . . . . .	2
2.3	Convex Relaxation . . . . .	3
2.3.1	Compositional Norms . . . . .	3
2.4	Connection to Multi-task Learning . . . . .	4
<b>3</b>	<b>A Boosting Approach to Coding</b>	<b>4</b>
3.1	Fenchel Conjugate . . . . .	5
3.2	Boosting With Fenchel Conjugates . . . . .	5
3.3	The Regularization Conjugate $\Phi^*(Z)$ . . . . .	7
3.3.1	Dual of a Compositional Norm . . . . .	8
3.4	The Subgradient $\partial_Z \Phi^*(Z)$ . . . . .	9
3.4.1	Solving the Infimal Convolution . . . . .	10
<b>4</b>	<b>Oracles for Infinite Bases</b>	<b>10</b>
4.1	$L_1$ Regularization . . . . .	10
4.2	$L_{2,1}$ Regularization . . . . .	11
4.3	$L_{2,1} + \gamma L_1$ Regularization . . . . .	11
<b>5</b>	<b>Results on Image Denoising</b>	<b>13</b>
5.1	Boosted Coding . . . . .	14
5.2	Alternating Optimization . . . . .	18

## 1 Introduction

A crucial part of many machine learning applications is representing the raw input in terms of a “code”, i.e. a set of features which captures the aspects of the input examples that are relevant to prediction. Unsupervised techniques such as clustering and sparse coding (Olshausen & Field, 1997) learn codes which capture the structure of unlabeled data, and have shown to be useful for a variety of machine learning problems (Raina et al., 2007; Bradley & Bagnell, 2009c; Mairal et al., 2009). In these techniques the input is represented as a combination of the features (also known as basis vectors or dictionary elements) that make up the code. The traditional approach to clustering and coding problems is to alternate between optimizing over the elements of the code and the combinations of elements used to represent the raw input. However, this “alternating optimization” approach is non-convex with many local minima, leading to recent work on alternative, convex versions of clustering and coding (Lashkari & Golland, 2008; Nowozin & Bakir, 2008; Bach et al., 2008).

This work adds several main contributions: we present a regularization function based on *compositional norms* that implements a convex version of sparse coding, we derive the *Fenchel conjugate* of these compositional norms, and we show how Fenchel conjugates can be used to construct an efficient boosting algorithms for convex (including non-differentiable) regularization functions.

Clustering and coding can be viewed as matrix factorization problems (Ding et al., 2005; Singh & Gordon, 2008), which seek to approximate a set of input signals  $X \approx f(BW)$  with the product of a dictionary matrix  $B$  a coefficient (or weight) matrix  $W$  and an elementwise transfer function  $f$ . When  $B$  is known and fixed and a regularization function is used to encourage  $W$  to be “sparse”, this is the sparse approximation technique developed in engineering and the sciences. Sparse approximation relies on an *optimization* algorithm to infer the Maximum A-Posteriori (MAP) weights  $\hat{W}$  that best reconstruct the signal. In this notation, each input signal forms a column of an input matrix  $X$ , and is generated by multiplying the dictionary (or basis matrix)  $B$  by a column from  $W$ , and (optionally) applying a transfer function  $f$ . This relationship is only approximate, as the input data is assumed to be corrupted by random noise. Priors which produce sparse solutions for  $W$ , especially  $L_1$  regularization, have gained attention because of their usefulness in ill-posed engineering problems (Tropp, 2006), elucidating neuro-biological phenomena, (Olshausen & Field, 1997; Karklin & Lewicki, 2005), face recognition (Wright et al., 2009), and semi-supervised and transfer learning (Raina et al., 2007; Bradley & Bagnell, 2009c; Mairal et al., 2009).

Sparse coding (Olshausen & Field, 1997) extends sparse approximation by also learning the basis matrix  $B$ . Clustering can also be viewed as a restricted form of the coding matrix factorization problem (Ding et al., 2005); a special case of coding where the basis vectors are the cluster centroids and the  $W$  matrix is the cluster membership of each example. Recently (Lashkari & Golland, 2008) showed that the clustering problem can be made convex by considering a fixed set of possible cluster centroids. Exemplars are a natural choice for these candidate cluster centroids, but (Nowozin & Bakir, 2008) show that for some problems this can be overly restrictive, and better results can be achieved by defining the problem in terms of a convex “master” problem, and a subproblem where new centroid candidates are generated.

We extend this “convex clustering” approach to the coding setting. Starting from a convex but intractable version of sparse coding in Section 2, we derive in Section 3 a boosting-style approach that is convex except for a subproblem. In Section 4 we give an efficient algorithm for solving the subproblem that will only improve on a fully convex exemplar-based approach, which is demonstrated on an image denoising task in Section 5. A similar convex formulation of sparse coding was independently developed by (Bach et al., 2008), who present an interesting optimization approach based on convex relaxations. Our work complements theirs by providing a novel optimization strategy applicable to all convex regularization functions.

## 2 Sparse Coding

### 2.1 Notation

Uppercase letters,  $X$ , denote matrices and lowercase letters,  $x$ , denote vectors. For matrices, superscripts and subscripts denote rows and columns respectively.  $X_j$  is the  $j$ th column of  $X$ ,  $X^i$  is the  $i$ th row of  $X$ , and  $X_j^i$  is the element in the  $i$ th row and  $j$ th column. Elements of vectors are indicated by subscripts,  $x_j$ .  $X^T$  is the transpose of matrix  $X$ . The Fenchel conjugate of a function  $f$  is denoted by  $f^*$ . The notation  $(x)_+$  means the larger of 0 or  $x$ .

### 2.2 Generative Model

From a probabilistic viewpoint, sparse coding fits a generative model (1) to unlabeled data, which factorizes a matrix of input examples,  $X \in \mathbb{R}^{m \times n}$ , in terms of latent variable matrices  $B \in \mathbb{R}^{m \times d}$  and  $W \in \mathbb{R}^{d \times n}$ . The matrix  $B$  is referred to as the basis, code, or dictionary, and is shared across all  $n$  examples (columns of  $X$ ). Given  $B$  the examples are assumed to be independent of each other. The matrix  $W$  is commonly referred to as the coefficients or the activations of the basis vectors.

$$P(X) = \int_B \int_W P(X|BW)P(W)P(B)dWdB \quad (1)$$

Applying the Maximum A Posteriori (MAP) approximation replaces the integration over  $W$  and  $B$  in (1) with its maximum value  $P(X|\hat{B}\hat{W})P(\hat{W})P(\hat{B})$ , where the values of the latent variables at the maximum,  $\hat{W}$  and  $\hat{B}$ , are referred to as the MAP estimates. Finding  $\hat{W}$  given  $B$  is an approximation problem; solving for  $\hat{W}$  and  $\hat{B}$  simultaneously over a set of examples is a coding problem.

We will focus on the case examined in (Lee et al., 2007) where the input  $X$  is assumed to be the matrix product  $BW$  corrupted by additive i.i.d Gaussian noise on each element, the prior  $P(W)$  is assumed to be Laplacian with mean zero, and the columns of the basis matrix are constrained to unit length. For numerical stability we minimize the negative log probability instead of maximizing (1):

$$\begin{aligned} \hat{W}, \hat{B} &= \arg \min_{W, B} \|BW - X\|_F^2 + \lambda \|W\|_1 \\ \text{s.t. } &\|B_i\|_2 = 1, \forall i. \end{aligned} \quad (2)$$

The  $L_1$ -norm<sup>1</sup>, is a common choice for the regularization function  $\Phi(W)$  because it tends to produce  $\hat{W}$  which are “sparse”—contain a small number of non-zero elements—even when the basis matrix has infinitely many columns (Bengio et al., 2006). This preference has been shown to be useful for applications such as prediction (Raina et al., 2007; Bradley & Bagnell, 2009c; Mairal et al., 2009) and denoising of images and video (Mairal et al., 2008).

## 2.3 Convex Relaxation

This formulation, and similar variants, is commonly solved by alternating between optimization over  $W$  and optimization over  $B$ , as both problems are convex when the other matrix is constant. However, a common objection to this approach is that the joint optimization problem is non-convex.

As noted by (Bengio et al., 2006) for the related problem of learning neural networks, the non-convexity can be removed if  $B$  is a fixed, infinite basis matrix containing all unit-length vectors as columns, and the optimization is only with respect to  $W$ . They go on to show that if  $L_1$  regularization is placed on  $W$ , it will have optimal solutions with only a finite number of non-zero weights, even if the number of basis vectors is infinite. Hence, the matrix  $B$  in (2) can be interpreted as the small set of basis vectors that have non-zero weight in  $W$ , and the fixed number of columns of  $B$  can be written as a *compositional norm* constraint on  $W$ .

### 2.3.1 Compositional Norms

A compositional norm is a norm composed of norms<sup>2</sup> over disjoint sets of variables (Bradley & Bagnell, 2009a). A useful and notationally convenient example of a compositional norm is a block norm<sup>3</sup>. Define a block norm of the matrix  $W$ ,  $\|W\|_{p,q}$  to be the  $L_q$  norm of the  $L_p$  norms of every row  $W^i$ :

$$L_{p,q}(W) = \|W\|_{p,q} = \left( \sum_i \left( \sum_j |W_j^i|^p \right)^{\frac{q}{p}} \right)^{\frac{1}{q}} \quad (3)$$

Since in our setting  $W$  is defined so that each row corresponds to a basis vector and each column corresponds to an example, a block norm can encourage all examples to use a subset of the basis vectors. For instance, the fixed size of  $B$  in (2) is equivalent to a hard constraint on  $W$  in terms of the non-convex  $L_{2,0}$  block semi-norm.  $L_{2,0}$  is the  $L_0$  semi-norm<sup>4</sup> of the  $L_2$  norm of each row of  $W$ , which counts how many basis vectors have non-zero entries in  $W$ .

Our convex coding formulation relaxes this non-convex  $L_{2,0}$  constraint by substituting regularization with the convex (but still sparse)  $L_{2,1}$  constraint:

<sup>1</sup>The  $L_p$  norm of a vector  $x$  is:  $\|x\|_p = (\sum_i |x_i|^p)^{1/p}$ , for  $p \geq 1$

<sup>2</sup>The component norms of a compositional norm can also be compositional norms, allowing hierarchical arrangements of three or more norms.

<sup>3</sup>It can be easily verified that (3) satisfies the definition of a norm.

<sup>4</sup>For  $0 \leq p < 1$ , the  $L_p$  norm of a vector  $x$  is redefined as:  $\|x\|_p = \sum_i |x_i|^p$

$$\arg \min_W \frac{1}{2} \|BW - X\|_F^2 + \lambda \left( \frac{1}{2} \|W\|_{2,1}^2 + \frac{\gamma}{2} \|W\|_1^2 \right) \quad (4)$$

The fact that the norms are squared in (4) will be mathematically convenient later, and is equivalent to scaling the regularization constant<sup>5</sup>.

The  $L_{2,1}$  block norm has been advocated recently as a regularization function for *multi-task learning* (Obozinski et al., 2006; Tropp et al., 2006), and the combination of the  $L_{2,1}$  block norm with the  $L_1$  norm was independently used for sparse coding by (Bach et al., 2008), although presented quite differently in terms of decomposition norms. Their work provides an interesting alternative framework and optimization strategy to the boosting approach presented here.

## 2.4 Connection to Multi-task Learning

The neural network learning setting considered in (Bengio et al., 2006) is equivalent to coding a single input vector in our setting.  $X$  corresponds to a vector  $y \in \mathbb{R}^m$ , containing the target outputs for  $m$  examples. Each column of  $B$  corresponds to the output of a non-linear classifier across all  $m$  examples, and  $W \in \mathbb{R}^d$  is a vector of weights which reconstruct the target outputs using the weak-learners in  $B$ .

The problem of coding a set of examples is equivalent to a multi-task learning problem, where each example (column of  $X$ ) is a task<sup>6</sup>. If the tasks are related, constraining them to use a smaller shared set of basis vectors should exploit the similarities between tasks and reduce the overall sample complexity of learning  $W$ .

For this purpose  $L_1$  regularization alone is unsatisfactory because it is indifferent between using separate sets of basis vectors for each task, or using a shared set of basis vectors for all tasks. The  $L_{2,1}$  block-norm regularization function used in (4) encourages information sharing across different tasks (examples).

## 3 A Boosting Approach to Coding

With a finite basis matrix  $B$ , (4) can be solved directly with various convex optimization techniques, including, e.g., subgradient descent. (Zinkevich, 2003). The regularization term effectively “selects” a small active set from the full basis (by encouraging multiple examples to share the same bases) to form the coded representation of the input. However, by exploiting properties of the  $L_{2,1}^2 + \gamma L_1^2$  regularization function it is also possible to handle an infinitely large  $B$ . We show how to solve (4) with an efficient boosting algorithm (Algorithm 1) that adds one new basis vector to the active basis matrix in each step.

This approach is motivated by the view of boosting as functional gradient descent in the space of weak learners (Mason et al., 2000; Friedman, 2001). In our case, each

<sup>5</sup>At the minimum  $\hat{W}$  of (4),  $\lambda \|\hat{W}\|_{2,1}^2 = \lambda' \|W\|_{2,1}$ , where  $\lambda' = \lambda \|\hat{W}\|_{2,1}$  is the original regularization constant scaled by the  $L_{2,1}$ -norm of  $\hat{W}$ .

<sup>6</sup>Test examples in any of the tasks are analogous to missing values in the examples, hence the input inpainting problem in coding is equivalent to prediction on test examples, and the input denoising problem corresponds to correcting label noise.

weak learner is a vector with unit length. Each boosting step, attempts to maximize the correlation between the negative loss gradient and a “small” change in  $W$ , as measured by a regularization function  $\Phi(W)$ . For infinite sets of potential basis vectors, the maximization at each step of this boosting approach is a non-convex sub-problem which must be solved (or approximated) by an oracle (Algorithm 2).

### 3.1 Fenchel Conjugate

A useful tool in our analysis will be the the Fenchel-Legendre conjugate (5), also known as the conjugate function (Boyd & Vandenberghe, 2004), which generalizes Legendre duality to include non-differentiable functions:

$$f^*(z) = \sup_x \{x^T z - f(x)\}. \quad (5)$$

$f^*(z)$  is the conjugate of the function  $f(x)$ , and the variable  $z$  is the dual variable of  $x$ . When the supremum in (5) is achieved, every maximal value  $\hat{x}$  of (5) is a subgradient<sup>7</sup> with respect to  $z$  of the conjugate function  $f^*(z)$ :

$$\frac{\partial f^*(z)}{\partial z} = \hat{x} = \arg \max_x \{x^T z - f(x)\}. \quad (6)$$

### 3.2 Boosting With Fenchel Conjugates

Each step of a gradient boosting style algorithm (Mason et al., 2000; Friedman, 2001) seeks a descent direction which provides the greatest reduction of loss for a small increase in the regularization function. If  $w$  is a vector of weights over all of the possible weak learners, we wish to find the step  $\Delta \hat{w}$  that is both maximally correlated with the negative loss gradient  $\nabla L(w)$ , and smaller than  $\epsilon$  as measured by the regularization function  $\Phi(w)$ :

$$\Delta \hat{w} = \arg \max_{\Phi(\Delta w) \leq \epsilon} -\nabla L(w)^T \Delta w. \quad (7)$$

This section will show that if the regularization function  $\Phi$  is convex, defined on  $\mathfrak{R}^d$ , and the constraint is strictly feasible<sup>8</sup>, every subgradient of its Fenchel conjugate evaluated on the loss gradient,  $\Delta \hat{w} \in \partial \Phi^*(-\frac{1}{\lambda} \nabla L(w))$ , is an optimal boosting step according to (7). This provides a useful method for constructing boosting algorithms for a wide class of regularization functions, and we will apply it to (4).

First we upper bound (7) by the minimum of the unconstrained Lagrange dual function, assuming the step size constraint is feasible (i.e. there exists a  $\Delta w$  such that  $\Phi(\Delta w) \leq \epsilon$ ):

$$\min_{\lambda \geq 0} \max_{\Delta w} -\nabla L(w)^T \Delta w - \lambda (\Phi(\Delta w) - \epsilon). \quad (8)$$

<sup>7</sup>A vector  $\phi \in R^n$  is a subgradient,  $\phi \in \partial_x f(x)$ , of a function  $f : R^2 \rightarrow (-\infty, \infty]$  at  $x \in R^n$  if  $\phi^T y \leq f(x+y) - f(x)$ ,  $\forall y \in R^n$ . Consider the function  $f = \max(x_1, x_2)$ . If  $x_1 > x_2$ , then there is a unique gradient  $\frac{\partial f}{\partial x} = [1 \ 0]$ . However, if  $x_1 = x_2$ , then  $[1 \ 0]$  and  $[0 \ 1]$  (and any convex combination of the two) are subgradients.  $\partial_x f(x)$  is the set of all subgradients.

<sup>8</sup>Slater’s Condition, which in this case means  $\Phi(\Delta w) < \epsilon$  for some vector  $\Delta w \in \mathfrak{R}^d$

If  $\Phi$  meets the conditions above, then the pair  $(\Delta\hat{w}, \hat{\lambda})$  that optimizes the upper bound (8), will also be optimal for the primal (7), iff the KKT conditions (generalized to sub-differential functions) are satisfied (Borwein & Lewis, 2005). In this case the conditions state that the negative loss gradient must be parallel to a subgradient of the regularization function (9), and all of the active constraints on  $\Delta\hat{w}$  must be tight (10):

$$-\nabla L(w) \in \hat{\lambda} \partial \Phi(\Delta\hat{w}) \quad (9)$$

$$\hat{\lambda} (\Phi(\Delta\hat{w}) - \epsilon) = 0. \quad (10)$$

Since  $\nabla L(w)^T \Delta w$  is a linear function of  $\Delta w$  and there is only one constraint, it must be active (i.e.  $\lambda > 0$ ) whenever  $\nabla L(w) \neq 0$ . Since boosting would stop if  $\nabla L(w) = 0$ , dividing (8) by  $\lambda$  and adding  $\epsilon$  does not change the optimal values of  $(\Delta\hat{w}, \hat{\lambda})$ , and produces the definition of the Fenchel conjugate of the regularization function, with the dual variable  $z = -\frac{1}{\lambda} \nabla L(w)$ :

$$\max_{\Delta w} \left\{ -\frac{1}{\lambda} \nabla L(w)^T \Delta w - \Phi(\Delta w) \right\} = \Phi^*(z). \quad (11)$$

Further, (6) means that all subgradients  $\Delta\hat{w} \in \partial_z \Phi^*(z)$  are boosting steps which will optimize (7). Hence for convex functions we find a boosting update rule:

$$\Delta\hat{w} = \partial \Phi^* \left( -\frac{1}{\lambda} \nabla L(w) \right) \quad (12)$$

that is a natural subgradient generalization of the mirror-descent rule for Legendre regularization functions (Cesa-Bianchi & Lugosi, 2006):

$$\Delta\hat{w} = \nabla \Phi^* \left( -\frac{1}{\lambda} \nabla L(w) \right). \quad (13)$$

By extending the mirror-descent rule to convex but non-differentiable regularization functions, we gain the ability to use mirror descent to optimize over infinite-dimensional spaces, while only ever storing a finite number of non-zero entries in  $w$ . The key is to employ regularization functions like  $L_1$  that can have extremely sparse conjugate subgradients, and to find a computational trick (oracle) that can compute a subgradient (12) without ever explicitly computing the full gradient of the loss (which will be infinitely large for an infinite set of possible basis vectors).

In the following we show how a practical boosting algorithm can be constructed for the  $L_{2,1}^2 + \gamma L_1^2$  regularization function used in (4), by deriving the conjugate of the regularization function,  $\Phi^*$ , computing a sparse subgradient over finite sets, and providing a tractable oracle heuristic for boosting from infinite sets of basis vectors (Algorithm 2). For this regularization function there are always subgradients consisting of a single new basis vector at each step, and we employ a step-wise fitting approach in Algorithm 1 to boost a basis matrix for sparse coding by adding one basis vector in each boosting step. Note that either  $\epsilon$  or  $\lambda$  is assumed to be a known hyper-parameter. Here we assume  $\lambda$  is a known constant as this leads to greater deflation between boosting steps, and a smaller, less coherent basis matrix.

**Algorithm 1** Boosted Coding

**Input:** Data matrix  $X \in \mathbb{R}^{m \times n}$ , scalars  $d, \lambda \in \mathbb{R}^+$ , convex functions  $L(BW, X)$ ,  $\Phi(W)$ , and a function  $b = \text{oracle}(-1/\lambda \nabla L(BW, X))$  which returns a new basis vector  $b$  corresponding to a non-zero row of  $\Delta w$ . (12).

**Output:** Active basis matrix  $B \in \mathbb{R}^{m \times d}$  and coefficients  $W \in \mathbb{R}^{d \times n}$ .

**Initialize:**  $W = 0^{d \times n}$ ,  $B = 0^{m \times d}$ .

**for**  $t = 1$  to  $d$  **do**

    Add a new basis vector with zero weight:

**1:**  $B_t = \text{oracle}(-1/\lambda \nabla L(BW, X))$ ,  $W^t = \vec{0}^T$

    Optimize  $W$ :

**2:**  $W = \arg \min_W L(BW, X) + \lambda \Phi(W)$

**if**  $\|W^t\|_2 = 0$  **then**

        return

**end if**

**end for**

**3.3 The Regularization Conjugate  $\Phi^*(Z)$** 

This section will prove the following lemma:

**Lemma 1.** *The Fenchel conjugate of:*

$$\Phi(W) = \frac{1}{2} \|W\|_{2,1}^2 + \frac{\gamma}{2} \|W\|_1^2$$

is the minimization over  $A$  of:

$$\Phi^*(Z) = \inf_A \frac{1}{2} \|Z - A\|_{2,\infty}^2 + \frac{1}{2\gamma} \|A\|_\infty^2 \quad (14)$$

The infimum over the additional variable  $A$  in (14) is known as the *infimal convolution* of the  $L_1^2$  and  $L_{2,1}^2$  terms. Since in the coding problem the dual variable  $Z$  is the negative loss gradient,  $Z = -\frac{1}{\lambda} \nabla L(W)$ , we will see in Section 4.3 that  $A$  has the effect of focusing the boosting step on the examples with the highest loss.

Lemma 1 follows from the the fact that the Fenchel conjugate of the sum of two functions is the *infimal convolution* of their conjugates (Rifkin & Lippert, 2007). The conjugate of the  $L_1^2$  squared norm term is well known: the dual of a squared norm  $f(x) = \frac{1}{2} \|x\|^2$  is simply the square of its *dual norm*. Let  $\|\cdot\|$  be a norm on  $\mathbb{R}^n$ . The dual norm,  $\|\cdot\|_*$  is defined as:

$$\|z\|_* = \sup_{\|x\| \leq 1} z^T x \quad (15)$$

For norms, the dual of the dual norm is the original norm  $\|x\|_{**} = \|x\|$ . The dual of the  $L_1$ -norm  $\|x\|_1$  is the  $L_\infty$ -norm  $\|z\|_\infty$ , and the dual of an  $L_p$ -norm  $\|x\|_p$  is the  $L_q$ -norm  $\|z\|_q$ , where  $1/p + 1/q = 1$  (Boyd & Vandenberghe, 2004). If  $f(x)$  is a squared norm multiplied by a scalar  $\gamma$ , then the conjugate will be multiplied by  $\gamma^{-1}$ ,  $(\gamma f(x))^* = \gamma^{-1} f^*(z)$  (Boyd & Vandenberghe, 2004)

Showing that the conjugate of the  $L_{2,1}^2$  term is  $L_{2,\infty}^2$ , requires a lemma about the dual norm of a block norm. We establish that the dual norm of the  $\|W\|_{2,1}$  block norm is  $\|Z\|_{2,\infty}$ , and finish the proof of Lemma 1 by proving the following lemma:

**Lemma 2.** *The dual of the  $L_{p,q}$  block norm is a  $L_{p^*,q^*}$  block norm where  $1/p + 1/p^* = 1/q + 1/q^* = 1$ ,  $p, p^*, q, q^* \in [1, \infty]$ .*

### 3.3.1 Dual of a Compositional Norm

Block norms are an example of a *compositional norm* which is a norm composed of norms over disjoint groups of variables. Given a column vector  $x \in \mathfrak{R}^n$ , a collection of disjoint sets of indices  $I = \{s_1, \dots, s_d\}$  such that each integer  $i \in \{1, \dots, n\}$  is in exactly one set  $s_j$ , a set of norms  $S = \{\|\cdot\|_{s_1}, \dots, \|\cdot\|_{s_d}\}$ , and a top-level norm  $\|\cdot\|$ , the compositional norm  $\|x\|_{C(I,S,\|\cdot\|)}$  is defined as:

$$\begin{aligned} \|x\|_{C(I,S,\|\cdot\|)} &= \|\alpha\| \\ \text{where } \alpha_j &= \|x_{i \in s_j}\|_{s_j}, \quad \forall j \in \{1, \dots, d\} \end{aligned} \quad (16)$$

A block norm  $\|X\|_{p,q} = (\sum_i \|Z^i\|_p^q)^{1/q}$ , can be written as a compositional norm where  $x$  is formed by concatenating the rows of  $X$  into a column vector. Each set  $s_i \in I$  is the indices of the variables in  $x$  corresponding to row  $i$  of  $X$ . Each norm  $\|\cdot\|_{s_i} \in S$  is the  $L_p$  norm  $\|\cdot\|_p$ , and the top level norm is  $\|\cdot\|_q$ . Lemma 2 is a special case of a more general lemma about the dual of a compositional norm:

**Lemma 3.** *The dual of a compositional norm,  $\|x\|_{C(I,S,\|\cdot\|)}^*$  is a compositional norm  $\|z\|_{C(I,S^*,\|\cdot\|)^*}$ , where the top-level norm and each norm in  $S$  has been replaced by its dual-norm:*

$$\begin{aligned} \|x\|_{C(I,S,\|\cdot\|)}^* &= \|z\|_{C(I,S^*,\|\cdot\|)^*} = \|\beta\|^* \\ \text{where } \beta_j &= \|z_{i \in s_j}\|_{s_j}^*, \quad \forall j \in \{1, \dots, d\} \end{aligned} \quad (17)$$

The proof of this lemma follows readily from the definition of the dual norm (15) and the fact that each norm in  $S$  operates on a disjoint set of variables. The dot product in (15) can be rewritten as the sum of the dot products of each set of variables (18). Next, by the definition of the compositional norm (16), the constraint  $\|x\|_{C(I,S,\|\cdot\|)} \leq 1$  can be replaced by a constraint on each set of variables,  $\|x_{i \in s_j}\|_{s_j} \leq \alpha_j$ , and a constraint on  $\alpha$ ,  $\|\alpha\| \leq 1$  (19). By the positive scalability property of norms, the constraint  $\|x_{i \in s_j}\|_{s_j} \leq \alpha_j$  can be rewritten as  $\|\alpha_j x_{i \in s_j}\|_{s_j} \leq 1$  (20). However, (20) is the definition of the dual norm of  $\alpha_j z_{i \in s_j}$ , and by using positive scalability again the  $\alpha_j$  factor can be moved outside the dual norm (21). Finally we substitute  $\beta_j = \|z_{i \in s_j}\|_{s_j}^*$ , recognize (22) as the dual of the top-level norm and recover (17).

$$\begin{aligned} \|z\|_{C(I,S,\|\cdot\|)}^* &= \sup_{\|x\|_{C(I,S,\|\cdot\|)} \leq 1} z^T x \\ &= \sup_{\|x\|_{C(I,S,\|\cdot\|)} \leq 1} \sum_j z_{i \in s_j}^T x_{i \in s_j} \end{aligned} \quad (18)$$

$$= \sup_{\|\alpha\| \leq 1} \sum_j \sup_{\|x_{i \in s_j}\|_{s_j} \leq \alpha_j} z_{i \in s_j}^T x_{i \in s_j} \quad (19)$$

$$= \sup_{\|\alpha\| \leq 1} \sum_j \sup_{\|x_{i \in s_j}\|_{s_j} \leq 1} \alpha_j z_{i \in s_j}^T x_{i \in s_j} \quad (20)$$

$$= \sup_{\|\alpha\| \leq 1} \sum_j \alpha_j \|z_{i \in s_j}\|_{s_j}^* \quad (21)$$

$$= \sup_{\|\alpha\| \leq 1} \sum_j \alpha_j \beta_j \quad (22)$$

$$= \|\beta\|^*, \quad \beta_j = \|z_{i \in s_j}\|_{s_j}^* \quad \forall j \in \{1, \dots, d\}$$

As an aside, this proof can be easily extended to compositional norms with more than two layers of norms, as long as each input variable has only one "path" to the top-level norm (i.e. the norms in each layer use disjoint subsets of the variables used by the norm above them). Since a compositional norm satisfies all of the properties of a standard norm, compositional norms can be used for some or all of the first-level norms in  $S$ . In this way, Lemma 3 can be used to find the dual of a compositional norm with an arbitrary number of layers.

### 3.4 The Subgradient $\partial_Z \Phi^*(Z)$

In order to apply our boosting approach (Algorithm 1) to the sparse coding problem (4), we must compute a subgradient  $\phi \in \partial_Z \Phi^*(Z)$  of (14), where the dual variable  $Z$  is the negative gradient of the loss:  $Z = -\frac{1}{\lambda} \nabla L(W)$ . Fortunately  $\partial_Z \Phi^*(Z)$  is generally very sparse, and we show that it equals:

$$\begin{aligned} \frac{\partial \Phi^*(Z)}{\partial Z_j^m} &= \begin{cases} 0, & \text{if } |Z_j^m| < \alpha \\ 0, & \text{if } \|(|Z_j^m| - \alpha)_+\|_2^2 < \kappa \\ Z_j^m - \text{sign}(Z_j^m)\alpha, & \text{otherwise} \end{cases} \\ \hat{\alpha} &= \|\hat{A}\|_\infty \\ \kappa &= \max_i \sum_k ((|Z_k^i| - \hat{\alpha})_+)^2 \end{aligned} \quad (23)$$

where  $\hat{\alpha}$  is the magnitude of the largest element of the matrix  $|\hat{A}|$  (derived below), which minimizes the infimal convolution in (14), and  $\kappa$  is equal to the maximal squared  $L_2$  norm of any row in the matrix  $Z - \hat{A}$ . Deriving this subgradient requires analyzing some details of the infimal convolution, but is necessary for understanding the subproblem involved in boosting from an infinite set of basis vectors.

### 3.4.1 Solving the Infimal Convolution

The solution  $\hat{A}$  to the infimal convolution in (14) is:

$$\hat{A}_j^i = \begin{cases} Z_j^i & |Z_j^i| \leq \alpha \\ \text{sign}(Z_j^i)\alpha & |Z_j^i| > \alpha \end{cases} \quad (24)$$

where  $\alpha = \|\hat{A}\|_\infty$ . The infimal convolution is effectively a one-dimensional search over  $\alpha \in [0, \|Z\|_\infty]$  which seeks to minimize the max of a set of piecewise quadratic functions ( $\|Z^i - A^i\|_2^2, \forall i$ ). Replacing the  $\|Z - A\|_{2,\infty}^2$  term in the infimal convolution with  $\max_i \sum_j \left( (|Z_j^i| - \alpha)_+ \right)^2$  produces:

$$\Phi^*(Z) = \min_\alpha \max_i \frac{1}{2\gamma} \alpha^2 + \frac{1}{2} \sum_j \left( (|Z_j^i| - \alpha)_+ \right)^2 \quad (25)$$

Note that this optimization problem is convex with respect to  $\alpha$ , since it is a max over convex functions. The minimizer,  $\hat{\alpha}$ , can be found to any desired accuracy by an interval bisection search<sup>9</sup> over  $\alpha \in [0, \|Z\|_\infty]$ . The subgradient (23) is the partial derivative<sup>10</sup> of (25).

## 4 Oracles for Infinite Bases

The convex exemplar-based approach to clustering formulated by (Lashkari & Goland, 2008) can be applied to sparse coding to create a simple oracle from a set of exemplars. In this case, the finite set of exemplars is the set of possible basis vectors, and the subgradient (23) derived above can be found efficiently by solving the infimal convolution. However, this solution can be improved by Algorithm 2, which optimizes over an infinite set of possible basis vectors. This section defines the optimization problem that must be solved to boost from an infinite set of possible basis vectors. We start by considering the two limiting cases of  $L_{2,1} + \gamma L_1$  regularization,  $L_1$  regularization and  $L_{2,1}$  regularization. We then present Algorithm 2 as a heuristic for finding a good solution in the general case.

### 4.1 $L_1$ Regularization

Assume the loss function is the squared reconstruction error,  $L(W, B) = \frac{1}{2} \|BW - X\|_F^2$ , and the regularization function is the  $L_1$ -norm  $\Phi(W) = \frac{1}{2} \|W\|_1^2$ , with conjugate  $\Phi^*(Z) = \frac{1}{2} \|Z\|_\infty^2$ . In this case, a subgradient of  $\Phi^*(Z)$  is a matrix with one non-zero element, corresponding to a maximal element of  $|Z|$ . Therefore, to find the best possible basis vector (with unit  $L_2$  norm) we must solve for the basis vector  $b_m$  that

<sup>9</sup>For a fixed level of accuracy the computational complexity of this search is at most  $O(Nd)$ , where  $N$  is the number of examples (columns of  $Z$ ) and  $d$  is the number of rows of  $Z$ .

<sup>10</sup>One might expect that since  $\alpha$  is a function of  $Z_j^m$ , we should be interested in the total derivative:  $\frac{d\Phi^*(Z)}{dZ_j^m} = \frac{\partial\Phi^*(Z)}{\partial Z_j^m} + \frac{\partial\Phi^*(Z)}{\partial\alpha} \frac{\partial\alpha}{\partial Z_j^m}$ . However, since  $\frac{\partial\Phi^*(Z)}{\partial\alpha} = 0$ , the total derivative is equal to the partial derivative.

produces the largest element of  $Z$ . Since  $Z = -\frac{1}{\lambda}\nabla L(W) = -\frac{1}{\lambda}B^T(BW - X)$ ,  $b_m$  is given by:

$$b_m = \frac{E_m}{\|E_m\|_2}, \text{ where:} \quad (26)$$

$$E = BW - X \quad (27)$$

$$m = \arg \max_j \|E_j\|_2^2.$$

$E$  is the reconstruction loss matrix, and  $m$  is the index of the example with the largest reconstruction loss. Hence  $L_1^2$  regularized boosting corresponds to adding the  $L_2$  projection of the loss gradient from the highest loss example to the active basis at each step.

## 4.2 $L_{2,1}$ Regularization

If instead we regularize with the  $L_{2,1}^2$  block norm  $\Phi(W) = \|W\|_{2,1}^2$ , with conjugate  $\Phi^*(Z) = \|Z\|_{2,\infty}^2$ , any row  $Z^m$  with maximal  $L_2$  norm is a subgradient of  $\Phi^*(Z)$ . Hence we can optimize over all possible basis vectors  $\vec{b}$ , by finding the basis vector  $\vec{b}_m$  best correlated with the loss gradients on all the examples (measured by  $L_2$ -norm):

$$\begin{aligned} b_m &= \arg \max_{\|b\|_2 \leq 1} \|b^T(BW - X)\|_2^2 \\ &= \arg \max_{\|b\|_2 \leq 1} b^T E E^T b. \end{aligned} \quad (28)$$

Again  $E$  is the reconstruction error (27). Although (28) is not convex, it is well known that  $b_m$  is the eigenvector associated with the maximum eigenvalue of the matrix  $E E^T$  (Boyd & Vandenberghe, 2004). Running Algorithm 1 with this choice of  $\Phi$  is very related to PCA, and can be interpreted as PCA with incomplete deflation.

## 4.3 $L_{2,1} + \gamma L_1$ Regularization

The  $L_{2,1} + \gamma L_1$  regularization used in (4) interpolates between the behavior of  $L_1$  and  $L_{2,1}$  based on the value of  $\gamma$ . The optimal new basis vector  $b_m$  will maximize the infimal convolution:

$$b_m = \arg \max_{\|b\|_2 \leq 1} \left\{ \min_{\alpha} \frac{\alpha^2}{\gamma} + \sum_j (|b^T E_j| - \alpha)_+^2 \right\} \quad (29)$$

The introduction of the minimization over  $\alpha$  makes (29) significantly more difficult to solve than the  $L_{2,1}$  case (28). However, we can solve the infimal convolution for a finite set of candidate basis vectors, and we have already seen the solution for the limiting cases of  $\gamma \rightarrow 0$  and  $\gamma \rightarrow \infty$ . To review: as  $\gamma \rightarrow 0$ ,  $\hat{\alpha}$  will converge to 0, and  $b_m$  will be equal to the solution of the  $L_{2,1}$  only case (28). When  $\gamma \rightarrow \infty$ ,  $\hat{\alpha}$  converges to  $\max_j |b_m^T E_j|$ , and  $L_{2,1} + \gamma L_1$  regularization reduces to  $L_1$  regularization. In this case<sup>11</sup>,  $b_m \propto E_j$ , where  $\|E_j\|_2 = \max_i \|E_i\|_2$ , i.e.  $b_m$  is guaranteed to be proportional to some column of the reconstruction error  $E$ .

<sup>11</sup>The  $L_2$  projection of any column of  $E$  with maximal  $L_2$  norm is a valid choice for  $b_m$

**Algorithm 2**  $L_{2,1} + \gamma L_1$  Heuristic

**Input:** Scalars  $\gamma$ ,  $\eta$  and  $N$ , reconstruction error matrix  $E \in \mathfrak{R}^{m \times n}$  where  $E = BW - X$ .

**Output:** New basis vector  $b$ .

**Initialization:** Compute a matrix  $\tilde{B} \in \mathfrak{R}^{m \times N+1}$  containing candidate basis vectors as columns:

**1:** Set  $\tilde{B}_1$  to the solution for  $L_{2,1}$  regularization (28).

**2:** Set  $\tilde{B}_2$  through  $\tilde{B}_N + 1$  equal to the  $L_2$  projections of the largest (in  $L_2$  norm)  $N$  columns of  $E$ :

$$\tilde{B}_i = E_i / \|E_i\|_2, \forall \|E_i\|_2 > C$$

**3:** Compute the dual variable:  $Z = -\frac{1}{\lambda} \tilde{B}^T E$

**4:** Find the index  $m$  of a maximal row by solving the infimal convolution:

$$m = \arg \max_i \min_{\alpha} \frac{1}{2\gamma} \alpha^2 + \frac{1}{2} \sum_j \left( (|Z_j^i| - \alpha)_+ \right)^2$$

**5:** Assign  $b$  to the best candidate:  $b = \tilde{B}_m$ .

**6:** Improve  $b$  by gradient ascent using (30):

**repeat**

**7:**  $z = -\frac{1}{\lambda} b^T E$

**8:**  $b = b + \eta \frac{\partial \Phi^*(z)}{\partial b}$

**until**  $\left\| \frac{\partial \Phi^*(z)}{\partial b} \right\|_2 < \epsilon$

**Return**  $b$

Algorithm 2 provides an empirically effective method for estimating  $\beta$  and  $b_m$ . It starts by combining the solutions for the  $L_{2,1}$  and  $L_1$  regularization cases discussed above into a matrix of candidate basis vectors  $\tilde{B}$ . Then a promising choice for  $b_m$  is selected by finding the basis vector in  $\tilde{B}$  which produces a maximal row of the infimal convolution over  $Z = -\frac{1}{\lambda} \tilde{B}^T E$ . Finally  $b_m$  is improved by gradient ascent to maximize the conjugate of the regularization,  $\Phi^*(Z^m)$ , where  $Z^m = -\frac{1}{\lambda} b_m^T E$  and the gradient is:

$$\begin{aligned} \frac{\partial \Phi^*(Z^m)}{\partial b_m} &= \frac{\partial \Phi^*(Z^m)}{\partial Z^m} \frac{\partial Z^m}{\partial b_m} \\ &= -\frac{1}{\lambda} \sum_j E_j \text{sign}(Z_j^m) (|Z_j^m| - \alpha)_+ \end{aligned} \quad (30)$$

where: 
$$\alpha = \min_{\alpha} \frac{\alpha^2}{\gamma} + \sum_j \left( (|b^T E_j| - \alpha)_+ \right)^2$$

## 5 Results on Image Denoising

We apply boosted coding (Algorithm 1) to the task of image denoising in order to evaluate its performance on a real-world task that is well-suited to sparse coding (Elad & Aharon, 2006), and lends itself particularly well to visualizing the behavior of the algorithm. The performance of alternating optimization<sup>12</sup> and boosted coding turn out to be quite similar on this task, with a slight advantage for the boosted approach. This result provides reassuring evidence that the non-convex but simple alternating optimization algorithm is not seriously impaired by inferior local minima on this task.

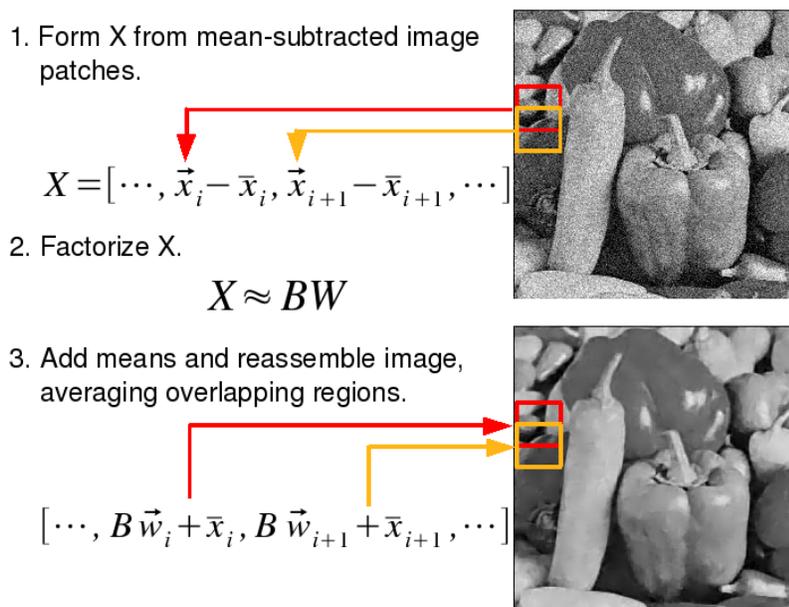


Figure 1: Image denoising proceeds by extracting overlapping patches from a noisy input image. Each patch is rearranged to form a column of the data matrix  $X$ . After  $X$  is approximated as  $B*W$ , the denoised image is reconstructed by averaging the overlapping patches.

Our approach is modeled on the K-SVD algorithm presented in (Elad & Aharon, 2006). As shown in Figure 1, overlapping patches are extracted from a noisy input image. Each patch is rearranged into a vector  $\vec{x}_i$ , the mean  $\bar{x}_i$  of the patch is subtracted, and the result becomes a column of the data matrix  $X$ .  $X$  is factorized into the product of  $B$  and  $W$  using sparse coding. Non-zero components of  $W$  are then refit without regularization. Finally, the denoised image is reconstructed by adding back the mean of each patch, and averaging areas of the image where multiple patches overlap. For

<sup>12</sup>Used in many past works such as (Raina et al., 2007).

these experiments we used 8x8 pixel patches with an overlap of four pixels between neighboring patches. The alternating optimization approach has two hyperparameters—the regularization constant  $\lambda$  and  $d$ , the number of columns of  $B$ —and boosted coding has two regularization constant hyperparameters  $\lambda$  and  $\gamma$ . The hyper-parameters of both algorithms were independently tuned for maximal performance, in order to isolate the effect of the different optimization strategies.

Independent and identically distributed Gaussian noise ( $\sigma = 0.1$ ) was added to five common benchmark images to match the assumed noise model of the  $L_2$  loss function (Figure 2, left-most column). Subtracting the mean of each patch removes low-frequency components of the image, leaving the coding problem to focus on identifying which high-frequency components of the image are contained in each patch. If  $W$  is simply set to zero, the result is to average each 8x8 patch of the image, which improves the Signal-to-Noise Ratio (SNR) of the low-frequency components of the image at the expense of the high-frequency details in the image (Figure 2, second column).

Coding  $X$  with alternating optimization or the convex up-to-an oracle approach presented in Algorithm 1 (Figure 2, right side) restores high-frequency detail by finding basis vectors that describe shared patterns across all patches. Figure 3 provides a zoomed-in view where the detail enhancement provided by sparse coding is more obvious. In our experiments both algorithms produce roughly equivalent results in terms of SNR, with a slight advantage for the convex approach. The parameters of both algorithms were tuned for maximal performance, in order to isolate the effect of the different optimization strategies.

## 5.1 Boosted Coding

The effect of relaxing the non-convex rank constraint on  $W$  by substituting  $L_{2,1}$  regularization can be clearly seen by comparing the basis vectors selected by the two algorithms. Boosted coding selects the most important basis vectors first, and those are used by many image patches due to the “group discount” provided by the  $L_2$  norm in  $L_{2,1}$  regularization (Figure 4). This causes the signal to noise ratio to rise quickly at the beginning of the process and then level off once most of the underlying signal can be represented by the basis vectors. As shown in Figure 5, images with just a few dominant high-frequency patterns (such as “house” and “boat”) converge quickly, while images with a variety of repetitive texture require more steps.

The first basis vectors chosen are also smoother in appearance than basis vectors chosen at later steps. This follows from the fact (shown in Section 4.3) that each step of boosting with  $L_{2,1} + \gamma L_1$  regularization will find a basis vector that is maximally correlated with the reconstruction error on a subset of the image patches (Figure 6). Therefore, in later rounds of boosting much of the structure of the image patches is already explained, and the reconstruction error on each patch will consist largely of noise. Additionally, the basis selected by boosted coding is less coherent (i.e. the basis vectors are less correlated with each other) than the basis selected by alternating optimization as shown in Table 1. Basis coherency is useful when denoising with only  $L_1$  regularization, as the image structure lies mainly in a low-dimensional subspace of the 64-dimensional space of 8x8 pixel patches, while the noise is spread uniformly throughout the space. However, because of the group discount on heavily used features



Figure 2: Results on five benchmark images (from top: Barbara, Boat, Lena, House, Peppers) show that both alternating optimization and boosted coding produce similar results for image denoising. The range reported for alternating optimization is the best and worst performance from 20 randomly initialized trials.

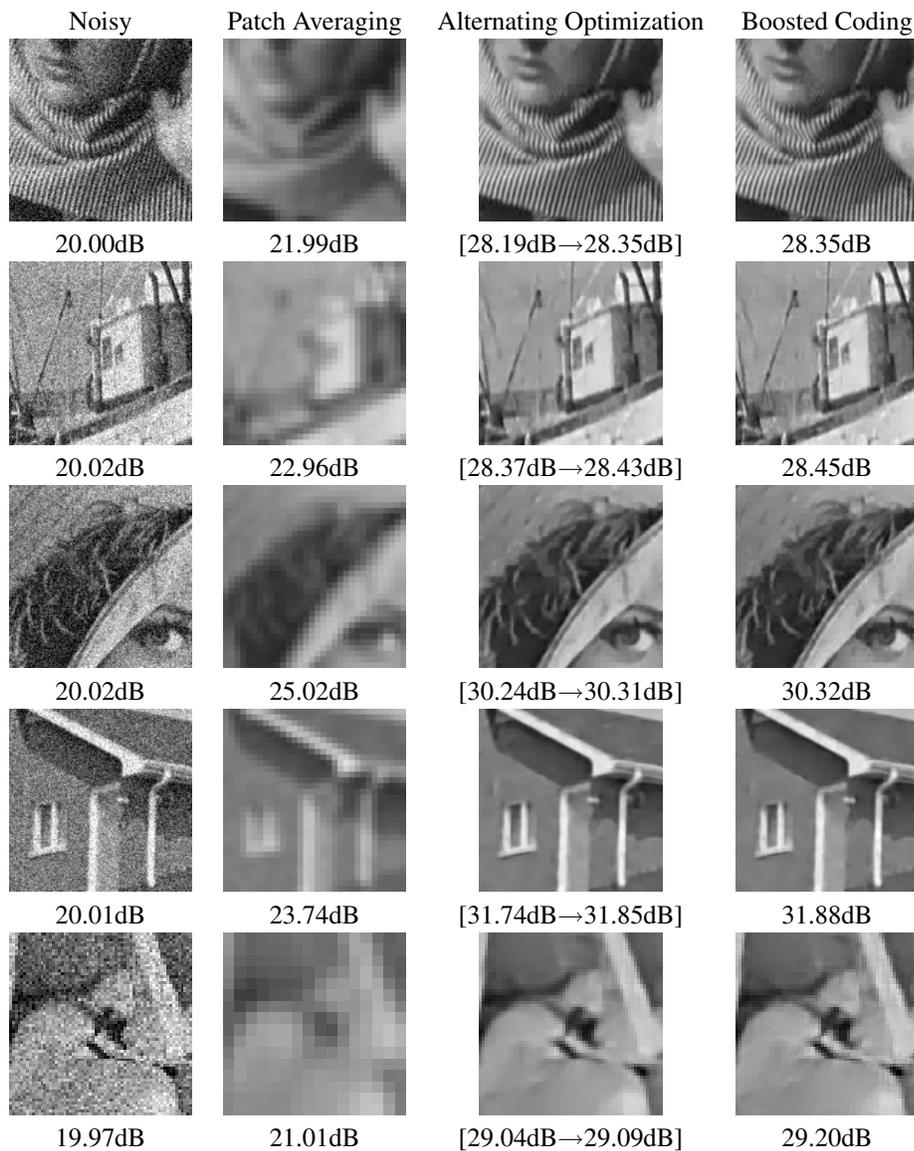


Figure 3: Close-ups of the denoising results for the five benchmark images, along with the Signal-to-Noise ratio after denoising as compared to the original image before Gaussian noise was added. The range reported for alternating optimization is the best and worst performance from 20 randomly initialized trials.

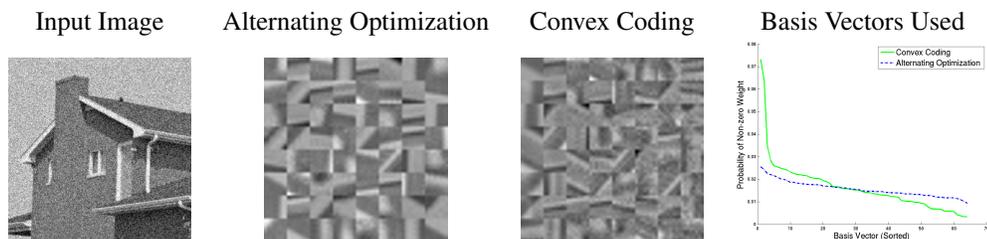


Figure 4: **Left:** Noisy input image. **Center:** 8x8 basis vectors learned by the alternating optimization and boosted coding algorithms. The boosted coding basis vectors are displayed in the order they were selected (top to bottom and left to right). In some cases boosted coding chose less than 64 basis vectors. **Right:** Fraction of image patches that each basis vector participated in. Boosted coding uses the first basis vectors selected much more often than later basis vectors.

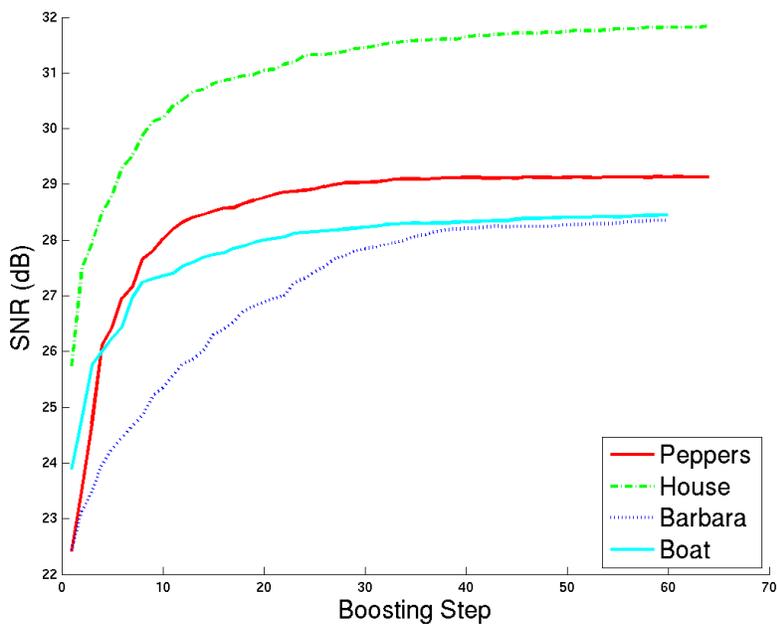


Figure 5: The basis vectors selected by the oracle quickly maximize denoising performance.

afforded by  $L_{2,1}$  regularization, Boosted Coding is able to perform well with a less coherent basis, as basis vectors that are used by fewer patches in the image are penalized more heavily.

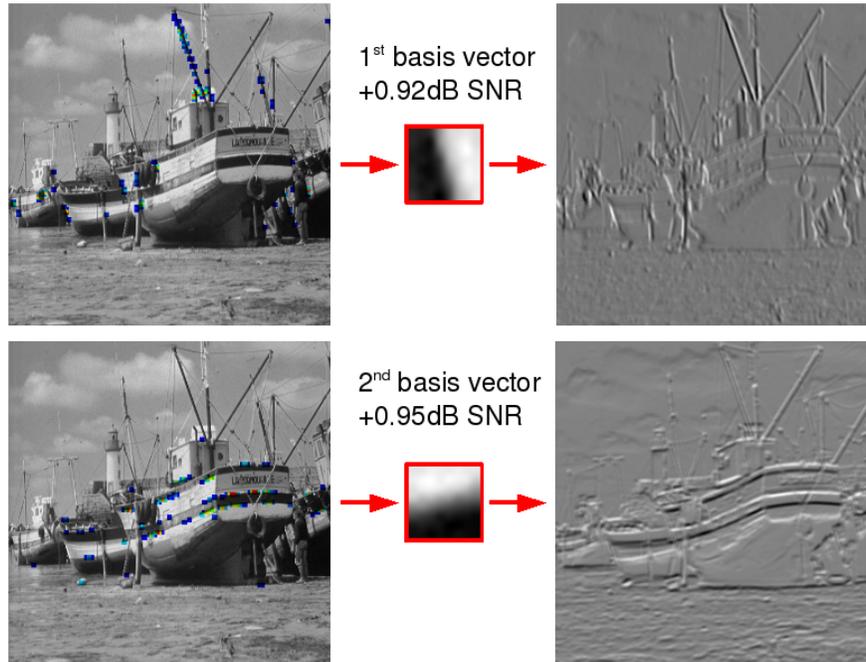


Figure 6: The basis vector selected by the oracle is chosen to maximize its correlation with a weighted set of high-error patches. **Left:** patches used by the oracle to compute the new basis vector (**center**) are colored according to weight, where blue is low and red is high. **Right:** filtering the image with the new basis vector indicates where it is most useful.

## 5.2 Alternating Optimization

A common objection to the traditional, alternating optimization approach to sparse coding (2) is that the non-convex rank constraint on  $B$  could result in the algorithm returning inferior local minima. Anecdotal evidence suggests this problem should be most acute for relatively small basis sizes, where there are only a small number of randomly-initialized basis vectors. In this case there are fewer degrees of freedom available to let alternating optimization escape a local minima. In our experience, inferior local minima, while they do occur, have a relatively small effect on the performance of the alternating optimization algorithm. Table 2 quantifies this assertion by showing the results of repeatedly running the alternating optimization image denoising

Image	Correlation between basis vectors	
	Boosted Coding	Alternating Optimization
Barbara	0.57 (0.12)	0.92 (0.16)
Boat	0.55 (0.21)	0.88 (0.30)
Lena	0.66 (0.24)	0.90 (0.35)
House	0.63 (0.23)	0.93 (0.33)
Peppers	0.60 (0.21)	0.91 (0.33)

Table 1: Alternating optimization tends to produce very coherent (correlated) basis vectors. Because of the additional  $L_{2,1}$  regularization, boosted coding performs just as well with a less coherent basis. Shown is the maximum absolute correlation between basis vectors for the best-performing basis on each image, with the mean correlation given in parentheses.

algorithm from different random initializations of the basis vectors. The optimization alternated between  $B$  and  $W$  20 times. This represents a stress-test for the alternating optimization algorithm as the basis contains only eight basis vectors. Even in this challenging case alternating optimization performs reasonably consistently, although Figure 7 provides a detailed look at a case where one local minima was superior to the others.

Image	SNR Improvement (dB)		
	Min	Max	Mean
Lena	4.60	4.78	$4.73 \pm 0.06\text{dB}$
Barbara	3.10	3.38	$3.20 \pm 0.11\text{dB}$
Boat	4.57	4.65	$4.62 \pm 0.03\text{dB}$
Peppers	6.55	6.84	$6.71 \pm 0.1\text{dB}$
House	6.20	6.25	$6.22 \pm 0.02\text{dB}$

Table 2: Variance observed when using alternating optimization  $L_1$ -regularized sparse coding to denoise grayscale images. Shown is the mean and standard deviation of the improvement (vs. patch averaging) in then Signal-to-Noise ratio of the image observed during 5 random initializations of  $B$ .

## Acknowledgements

We would like to thank the UAI reviewers of the conference version of this paper (Bradley & Bagnell, 2009b) for their helpful comments.

## References

Bach, F., Mairal, J., & Ponce, J. (2008). Convex sparse matrix factorizations. *CoRR*, *abs/0812.1869*.

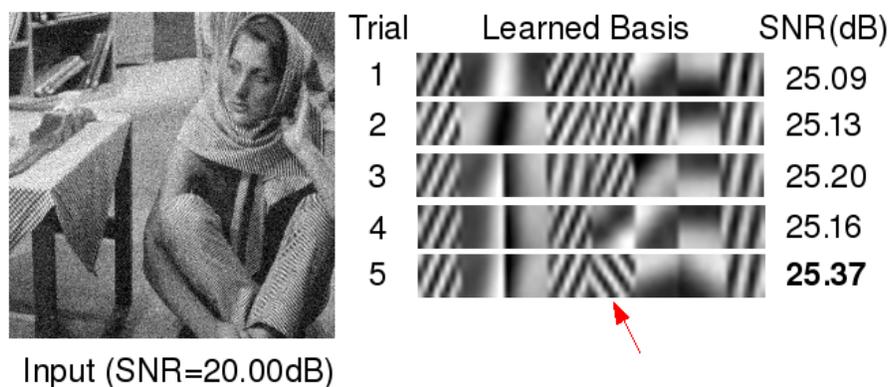


Figure 7: **Left:** input image. **Right:** five bases learned by alternating optimization starting from random initializations. The basis that performed best found an important pattern (indicated) that was not representable in the other bases, which were stuck in worse local minima.

Bengio, Y., Le Roux, N., Vincent, P., Delalleau, O., & Marcotte, P. (2006). Convex neural networks. *Advances in Neural Information Processing Systems*, 18, 123.

Borwein, J., & Lewis, A. (2005). *Convex analysis and nonlinear optimization*. CMS-Springer Books. 2nd edition.

Boyd, S., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press.

Bradley, D. M., & Bagnell, J. A. (2009a). *Convex coding* (Technical Report CMU-TR-09-22). Robotics Institute, Carnegie Mellon University.

Bradley, D. M., & Bagnell, J. A. (2009b). Convex coding. *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*.

Bradley, D. M., & Bagnell, J. A. (2009c). Differentiable sparse coding. *Neural Information Processing Systems 22*.

Cesa-Bianchi, N., & Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge University Press.

Ding, C., He, X., & Simon, H. (2005). On the equivalence of nonnegative matrix factorization and spectral clustering. *Proceedings of the SIAM Data Mining Conference*.

Elad, M., & Aharon, M. (2006). Image denoising via learned dictionaries and sparse representation. *CVPR '06: Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (pp. 895–900). Washington, DC, USA: IEEE Computer Society.

- Friedman, J. (2001). Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, 29, 1189–1232.
- Karklin, Y., & Lewicki, M. S. (2005). A hierarchical bayesian model for learning non-linear statistical regularities in non-stationary natural signals. *Neural Computation*, 17, 397–423.
- Lashkari, D., & Golland, P. (2008). Convex clustering with exemplar-based models. *Advances in Neural Information Processing Systems 20* (pp. 825–832). Cambridge, MA: MIT Press.
- Lee, H., Battle, A., Raina, R., & Ng, A. Y. (2007). Efficient sparse coding algorithms. *Advances in Neural Information Processing Systems 19*. Cambridge, MA: MIT Press.
- Mairal, J., Bach, F., Ponce, J., Sapiro, G., & Zisserman, A. (2009). Supervised dictionary learning. In D. Koller, D. Schuurmans, Y. Bengio and L. Bottou (Eds.), *Advances in neural information processing systems 21*.
- Mairal, J., Sapiro, G., & Elad, M. (2008). Learning multiscale sparse representations for image and video restoration. *Multiscale Modeling & Simulation*, 7, 214–241.
- Mason, L., Baxter, J., Bartlett, P., & Frean, M. (2000). Boosting algorithms as gradient descent. In *Advances in Neural Information Processing Systems 12* (pp. 512–518).
- Nowozin, S., & Bakir, G. (2008). A decoupled approach to exemplar-based unsupervised learning. *ICML '08: Proceedings of the 25th international conference on Machine learning* (pp. 704–711). New York, NY, USA: ACM.
- Obozinski, G., Taskar, B., & Jordan, M. (2006). Multi-task feature selection. *ICML-06 Workshop on Structural Knowledge Transfer for Machine Learning*.
- Olshausen, B., & Field, D. (1997). Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision Research*.
- Raina, R., Battle, A., Lee, H., Packer, B., & Ng, A. Y. (2007). Self-taught learning: Transfer learning from unlabeled data. *ICML '07: Proceedings of the 24th international conference on Machine learning*.
- Rifkin, R., & Lippert, R. (2007). Value regularization and fenchel duality. *The Journal of Machine Learning Research*, 8, 441–479.
- Singh, A. P., & Gordon, G. J. (2008). A unified view of matrix factorization models. *Machine Learning and Knowledge Discovery in Databases, European Conference (ECML/PKDD)*. ECML/PKDD-2008.
- Tropp, J. A. (2006). Algorithms for simultaneous sparse approximation: part ii: Convex relaxation. *Signal Process.*, 86, 589–602.
- Tropp, J. A., Gilbert, A. C., & Strauss, M. J. (2006). Algorithms for simultaneous sparse approximation: part i: Greedy pursuit. *Signal Process.*, 86, 572–588.

- Wright, J., Yang, A. Y., Ganesh, A., Sastry, S. S., & Ma, Y. (2009). Robust face recognition via sparse representation. *IEEE Trans. Pattern Anal. Mach. Intell.*, *31*, 210–227.
- Zinkevich, M. (2003). Online convex programming and generalized infinitesimal gradient ascent. *Twentieth International Conference on Machine Learning*.