Optimization Methods for Convolutional Sparse Coding

Hilton Bristow · **Simon Lucey**

Abstract Sparse and convolutional constraints form a natural prior for many optimization problems that arise from physical processes. Detecting motifs in speech and musical passages, super-resolving images, compressing videos, and reconstructing harmonic motions can all leverage redundancies introduced by convolution. Solving problems involving sparse and convolutional constraints remains a difficult computational problem, however. In this paper we present an overview of *convolutional sparse coding* in a consistent framework. The objective involves iteratively optimizing a convolutional least-squares term for the basis functions, followed by an L_1-regularized least squares term for the sparse coefficients. We discuss a range of optimization methods for solving the convolutional sparse coding objective, and the properties that make each method suitable for different applications. In particular, we concentrate on computational complexity, speed to ϵ convergence, memory usage, and the effect of implied boundary conditions. We present a broad suite of examples covering different signal and application domains to illustrate the general applicability of convolutional sparse coding, and the efficacy of the available optimization methods.

Keywords convolution · sparse coding · machine learning · SISC · Fourier · L_1 · ADMM · FISTA

Hilton Bristow
The Queensland University of Technology
Brisbane, Australia
E-mail: hilton.bristow@gmail.com

Simon Lucey
Carnegie Mellon University
Pittsburgh, USA
E-mail: slucey@cs.cmu.edu
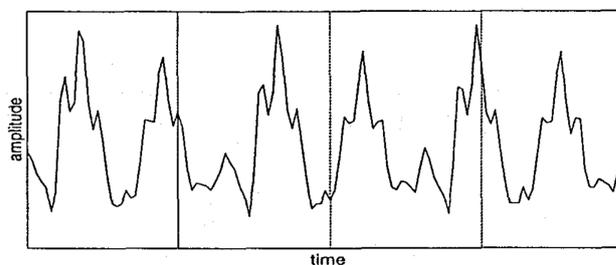
Fig. 1 Blocking results in arbitrary alignment of the underlying signal structure, artificially inflating the rank of the basis required to reconstruct the signal. Convolutional sparse coding relaxes this constraint, allowing shiftable basis functions to discover a lower rank structure.

1 Introduction

Convolutional constraints arise in response to many natural phenomena. They provide a means of expressing an intuition that the relationships between variables should be both spatially or temporally localized, and redundant across the support of the signal.

Coupled with sparsity, one can also enforce selectivity, *i.e.* that the patterns observed in the signal stem from some structured underlying basis whose phase alignment is arbitrary.

In this paper, we discuss a number of optimization methods designed to solve problems involving convolutional constraints and sparse regularization in the form of the convolutional sparse coding algorithm,[1] and show through examples that the algorithm is well suited to a wide variety of problems that arise in computer vision.

---

[1] Convolutional Sparse Coding has also been coined *Shift Invariant* Sparse Coding (SISC), however the authors believe the term "convolutional" is more representative of the algorithm's properties.

Consider the signal in Fig. 1 from (Lewicki and Sejnowski, 1999). The signal is composed of two distinct modes, appearing at multiple intervals (you could consider the modes to be expressions that are repeated over the course of a conversation, features that appear multiple times within an image, or particular motifs in a musical passage or speech). We wish to discover these modes and their occurrences in the signal in an unsupervised manner.

If the signal is segmented into blocks, the latent structure becomes obfuscated, and the basis learned must have the capacity to reconstruct each block in isolation. In effect, we learn a basis that is higher rank than the true basis due to the artificial constraints placed on the temporal alignment of the bases.

Convolutional sparse coding makes no such assumptions, allowing shiftable basis functions to discover a lower rank structure, and should therefore be preferred in situations where the basis alignment is not known *a priori*. As we aim to demonstrate in this paper, this arises in a large number of practical applications.

The notion of translation invariant optimization stems from the thought experiments of Simoncelli *et al*. (Simoncelli et al., 1992). His motivation for considering translation invariance came from the context of wavelet transforms for signal processing. Amongst others, he had observed that block-based wavelet algorithms were sensitive to translation and scaling of the input signal.

As an example, he chose an input signal to be one of the wavelet basis functions (yielding a single reconstruction coefficient), then perturbed that signal slightly to produce a completely dense set of coefficients. The abrupt change in representation in the wavelet domain due to a small change in the input illustrated the wavelet transform's unsuitability for higher-level summarization.

Olshausen and Field showed that sparsity alone is a sufficient driver for learning structured overcomplete representations from signals, and used it to learn a basis for natural image patches (Olshausen and Field, 1996). The resulting basis, featuring edges at different scales and orientations, was similar to the receptive fields observed in the primary visual cortex.

The strategy of sampling patches from natural images has come under fire however, since many of the learned basis elements are simple translations of each other - an artefact of having to reconstruct individual patches, rather than entire image scenes (Kavukcuoglu et al., 2010). Removing the artificial assumption that image patches are independent - by modelling interactions in a convolutional objective - results in more

expressive basis elements that better explain the underlying mechanics of the signal.

Lewicki and Sejnowski (Lewicki and Sejnowski, 1999) made the first steps towards this realization, by finding a set of sparse coefficients (value and temporal position) that reconstructed the signal with a fixed basis. They remarked at the spike-like responses observed, and the small number of coefficients needed to achieve satisfactory reconstructions.

Introducing sparsity brought with it a set of computational challenges that made the resulting objectives difficult to optimize. (Olshausen and Field, 1996) explored coefficients drawn from a Cauchy distribution (a smooth heavy-tailed distribution) and a Laplacian distribution (a non-smooth heavy-tailed distribution), citing that in both cases *they favour among activity states with equal variance, those with fewest non-zero coefficients*. (Olshausen and Field, 1996) inferred the coefficients as the equilibrium solution to a differential equation. (Lewicki and Sejnowski, 1999) assumed Laplacian distributed coefficients, and noted that due to the high sparsity of the desired response, it would be sufficient to replace exact inference with a procedure for guessing the values and temporal locations of the non-zero coefficients, then refining the results through a modified conjugate gradient local search.

Tibshirani (Tibshirani, 1996) introduced a convex form of the sparse inference problem - estimating Laplacian distributed coefficients which minimize a least-squares reconstruction error - using L_1-norm regularization and presented a method for solving it with existing quadratic programming techniques.

The full convolutional sparse coding algorithm culminated in the work of (Grosse et al., 2007). Fundamentally, Grosse extended Olshausen and Field's sparse coding algorithm to include convolutional constraints and generalized Lewicki and Sejnowski's convolutional sparse inference to 2D. Algorithmically, Grosse drew on the work of Tibshirani (Tibshirani, 1996) in expressing Laplacian distributed coefficients as L_1-norm regularization, and used the feature sign search minimization algorithm proposed by his colleague in the same year (Lee et al., 2007) to solve it efficiently. The form he introduced is the now canonical bilinear convolutional sparse coding algorithm.

Convolutional sparse coding has found application in learning Gabor-like bases that reflect the receptive fields of the primary visual cortex (Olshausen and Field, 1997), elemental motifs of visual (Zeiler et al., 2010), speech (Grosse et al., 2007; Lewicki and Sejnowski, 1999) and musical (Grosse et al., 2007; Mø rup et al., 2008) perception, a basis for human motion and articulation (Zhu and Lucey, 2014), mid-level discriminative
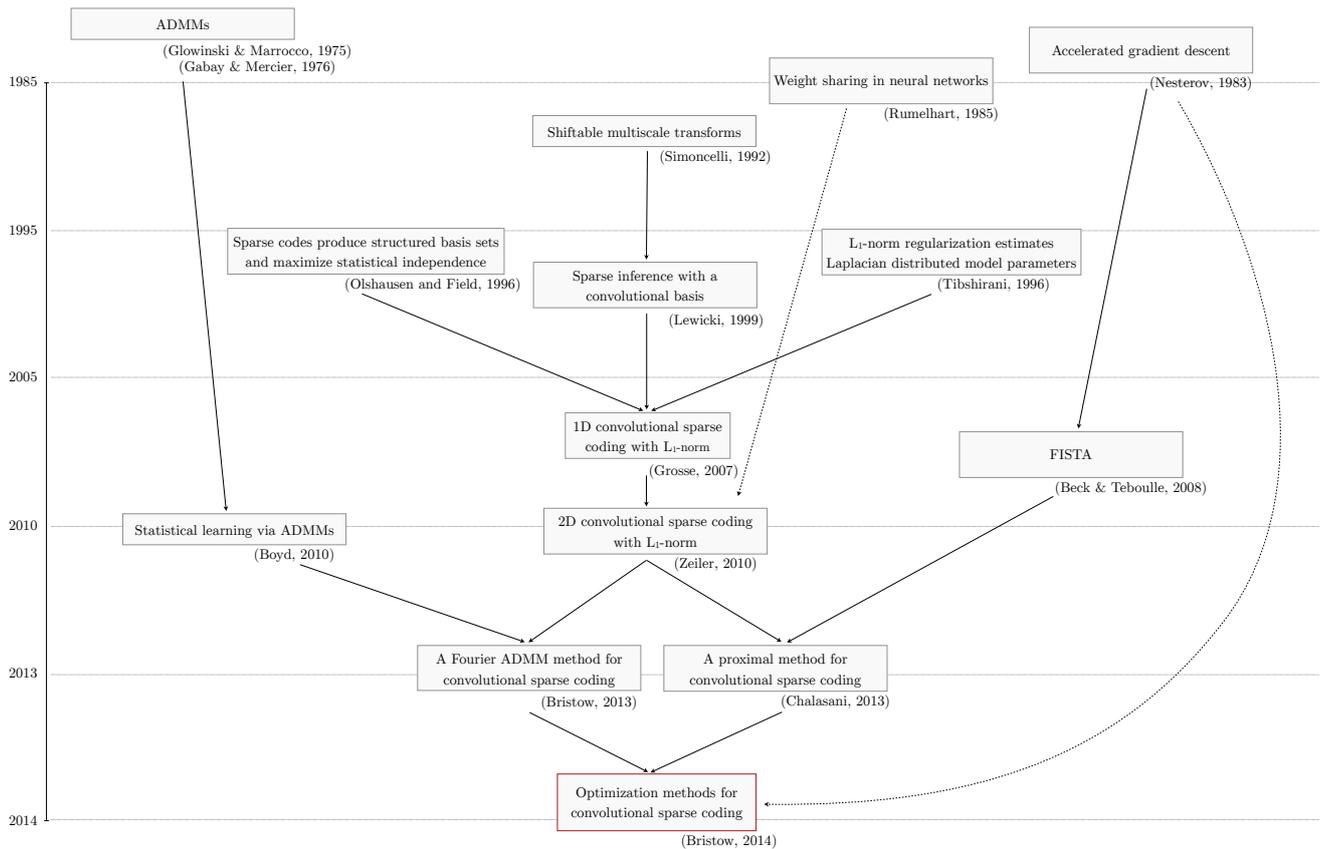
Fig. 2 A brief history of the works that influenced the direction of the convolutional sparse coding problem, and how it is optimized. The text within each box indicates the theme or idea that the paper introduced.

patches (Sohn and Jung, 2011) and unsupervised learning of hierarchical generative models (Lee et al., 2009; Zeiler et al., 2010).

The large-scale nature of the latter applications have placed great demands on the computational efficiency of the underlying algorithms. Coupled with the steady advances in machine learning and computing, this has given rise to a range of optimization approaches for convolutional sparse coding. (Chalasani and Principe, 2012) introduced a convolutional extension to the FISTA algorithm for sparse inference (Beck and Teboulle, 2009), and (Bristow et al., 2013) introduced a Fourier method based on the closely-related ADMM (Boyd, 2010). Convolutional basis learning has largely been relegated to gradient descent, however (Grosse et al., 2007) introduced a Fourier domain adaptation that minimized the inequality constrained problem via the Lagrange dual.

In this paper, we make the following contributions:

1. We provide a brief discussion detailing the works that influenced the development of the convolutional sparse coding algorithm and how it stems from natural processes.

2. We argue that *convolutional* sparse coding makes a set of assumptions that are more appropriate in many tasks where block or sampled sparse coding is currently used.

3. We discuss practical optimization of convolutional sparse coding, including speed, memory usage and assumptions on boundary conditions.

4. We show through a number of examples the applicability of convolutional sparse coding to a wide range of problems that arise in computer vision, and when particular problems benefit from different optimization methods.

It is worth emphasising that the algorithms presented in this work are not new. Detailed treatments of the algorithms can be found in their respective works (Bristow et al., 2013; Chalasani and Principe, 2012; Grosse et al., 2007). Rather, this paper aims to present the available algorithms in a unified framework and act as a resource for practitioners wishing to understand their properties.

(a)

$$\min \; f(\boldsymbol{x}) \qquad\qquad \min \; f(\boldsymbol{x})$$
$$\text{s.t.} \;\; ||x||_2 \le \epsilon \qquad \text{s.t.} \;\; |x|_1 \le \epsilon$$

(b)
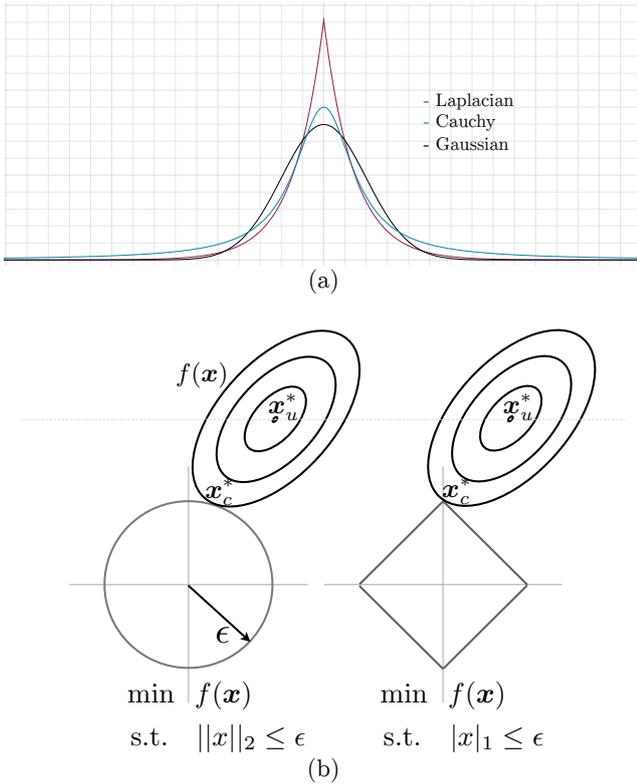
Fig. 3 Sparsifying distributions and their effect as regularizers. (a) For a given variance, the Laplacian and Cauchy distributions have more probability mass centered around zero. (b) As a result, $L1$ constraints favour axis-aligned solutions (only the second dimension is non-zero).

1.1 Outline

We begin in §2 with an introduction of the convolutional sparse coding formulation. We discuss optimization methods for convolutional sparse inference first in §3, as they are of most algorithmic interest, followed by methods for basis/filter updates in §4. In §6 we present a number of applications of convolutional sparse coding, each with a discussion of the preferred optimization method and implication of boundary effects. Proofs and implementation details are deferred until the appendices.

1.2 Notation

Matrices are written in bold uppercase (*e.g.* $\mathbf{A}$), vectors in bold lowercase (*e.g.* $\mathbf{x}$) and scalars in regular typeface (*e.g.* K). Convolution is represented by the $*$ operator, and correlation by the $\star$ operator. We treat all signals as vectors - higher-dimensional spatial or spatio-temporal signals are implicitly vectorized - however operations on these signals are performed in their original domain. We adopt this notation because (i) the domain of the signal

is largely irrelevant in our analysis and the algorithms we present, and (ii) for any linear transform in the original domain, there exists an equivalent transform in the vectorized domain.

A ˆ applied applied to any vector denotes the Discrete Fourier Transform (DFT) of a vectorized signal $\mathbf{a}$ such that $\hat{\mathbf{a}} \leftarrow \mathcal{F}(\mathbf{a}) = \mathbf{F}\mathbf{a}$, where $\mathcal{F}()$ is the Fourier transform operator and $\mathbf{F}$ is the $D \times D$ matrix of complex Fourier bases.

2 Convolutional Sparse Coding

The convolutional sparse coding problem consists of minimizing a convolutional model-fitting term f and a sparse regularizer g,

$$\arg\min_{\mathbf{d},\mathbf{z}} \;\; f(\mathbf{d},\mathbf{z}) + \beta g(\mathbf{z}) \tag{1}$$

where $\mathbf{d}$ is the convolutional kernel, $\mathbf{z}$ are the set of sparse coefficients, and β controls the tradeoff between reconstruction error and sparsity of representation. The input can be reconstructed via the convolution,

$$\mathbf{x} = \mathbf{d} * \mathbf{z} . \tag{2}$$

Assuming Gaussian distributed noise and Laplacian distributed coefficients, Eqn. 1 can be written more formally as,

$$\arg\min_{\mathbf{d},\mathbf{z}} \;\; \frac{1}{2}||\mathbf{x} - \mathbf{d} * \mathbf{z}||_2^2 + \beta ||\mathbf{z}||_1$$
$$\text{s.t.} \;\;\;\; ||\mathbf{d}||_2^2 \le 1 \tag{3}$$

The remainder of our analysis is based around efficient methods of optimizing this objective. The objective naturally extends to multiple images and filters,

$$\arg\min_{\mathbf{d},\mathbf{z}} \;\; \frac{1}{2}\sum_{i=1}^{M}||\mathbf{x}_i - \sum_{j=1}^{N}(\mathbf{d}_j * \mathbf{z}_{i,j})||_2^2 + \beta \sum_{i=1}^{M}\sum_{j=1}^{N}||\mathbf{z}_{i,j}||_1$$
$$\text{s.t.} \;\;\;\; ||\mathbf{d}_j||_2^2 \le 1 \;\; \forall j \in 1 \dots N \tag{4}$$

however this form quickly becomes unwieldy, so we only use it when the number of filters or images needs to be emphasized.

Contrast the objective of Eqn. 3 with that of conventional sparse coding,

$$\arg\min_{\mathbf{B},\mathbf{Z}} \;\; ||\mathbf{X} - \mathbf{B}\mathbf{Z}||_2^2 + \beta ||\mathbf{Z}||_1$$
$$\text{s.t.} \;\;\;\; ||\mathbf{B}_i||_2^2 \le 1 \;\; \forall i \tag{5}$$

Here, we are solving for a set of basis vectors $\mathbf{B}$ and sparse coefficients $\mathbf{Z}$ in alternation that reconstruct patches or samples of the signal $\mathbf{X}$ *in isolation*. This

is an important distinction to make, and forms the fundamental difference between convolutional and patch-based sparse coding.

In the limit when $\mathbf{X}$ contains *every patch* from the full image, the two sparse coding algorithms behave equivalently, however the patch-based algorithm must store a redundant amount of data, and cannot take advantage of fast methods for evaluating the inner product of the basis with each patch (*i.e.* convolution).

The objective of Eqn. 3 is bilinear – solving for each variable whilst holding the other fixed yields a convex subproblem, however the objective is not jointly convex in both. We optimize the objective in alternation, iterating until convergence. There are no guarantees that the final minima reached is the global minima, however in practice multiple trials reach minima of comparable quality, even if the exact bases and distribution of coefficients learned are slightly different.

In the sections that follow, we introduce a range of algorithms that can solve for the filters and coefficients. Since the alternation strategy treats each independently, we largely consider the algorithms in isolation, however some care must be taken in matching appropriate boundary condition assumptions.

2.1 Other Formulations

It should be noted that the algorithm of Eqn. 3 is not the only conceivable formulation of convolutional sparse coding. In particular, there has been growing interest in non-convex sparse coding with hyper-Laplacian, L_{0+} and other exotic priors (Zhang, 2010; Wipf et al., 2011). Whilst these methods have not yet been extended to *convolutional* sparse coding, there is no fundamental barrier to doing so.

3 Solving for Coefficients

It is natural to begin with methods for convolutional sparse inference, since this is where most of the research effort has been focussed. Solving for the coefficients $\mathbf{z}$ involves optimizing the unconstrained objective,

$$\arg\min_{\mathbf{z}} ||\mathbf{x} - \mathbf{d} * \mathbf{z}||_2^2 + \beta||\mathbf{z}||_1 \tag{6}$$

where β controls the tradeoff between sparsity and reconstruction error. This objective is difficult to solve because (i) it involves a non-smooth regularizer, (ii) the least-squares system cannot be solved directly (forming explicit convolutional matrices for large inputs is infeasible), and (iii) the system involves a large number of variables, especially if working with megapixel imagery, *etc*.

In the case of multiple images and filters,

$$\arg\min_{\mathbf{z}} \quad \frac{1}{2}\sum_{i=1}^{M}||\big(\mathbf{x}_i - \sum_{j=1}^{N}(\mathbf{d}_j * \mathbf{z}_{i,j})\big)||_2^2$$
$$+ \beta\sum_{i=1}^{M}\sum_{j=1}^{N}||\mathbf{z}_{i,j}||_1 \tag{7}$$

and given that the minima of the sum of convex functions is the sum of their minima,

$$\min_{\mathbf{z}}\sum_{i=1}^{M}f_i(\mathbf{z}) = \sum_{i=1}^{M}\min_{\mathbf{z}}f_i(\mathbf{z}) \tag{8}$$

the coefficients for each image can be inferred separately,

$$\mathbf{z}_i^{\star} = \arg\min_{\mathbf{z}} \quad \frac{1}{2}||\mathbf{x}_i - \sum_{j=1}^{N}(\mathbf{d}_j * \mathbf{z}_{i,j})||_2^2 + \beta\sum_{j=1}^{N}||\mathbf{z}_{i,j}||_1 \tag{9}$$

The two methods that we introduce are both based around partitioning the objective into the smooth model-fitting term and non-smooth regularizer that can then be handled separately.

3.1 ADMM Partitioning

The alternating direction method of minimizers (ADMM), was proposed jointly by (Glowinski and Marroco, 1975; Gabay and Mercier, 1976), though the idea can be traced back as early as Douglas-Rachford splitting in the mid-1950s. (Boyd, 2010) presents a thorough overview of ADMMs and their properties. ADMMs solve problems of the form,

$$\arg\min_{\mathbf{z},\mathbf{t}} \quad g(\mathbf{z}) + h(\mathbf{t})$$
$$\text{s.t.} \quad \mathbf{A}\mathbf{z} + \mathbf{B}\mathbf{t} = \mathbf{c} \tag{10}$$

To express convolutional sparse inference in this form, we perform the (somewhat unintuitive) substitution,

$$g(\mathbf{z}) = \frac{1}{2}||\mathbf{x} - \mathbf{d} * \mathbf{z}||_2^2 \qquad h(\mathbf{z}) = \beta||\mathbf{z}||_1 \tag{11}$$
$$\mathbf{A} = \mathbf{I} \qquad \mathbf{B} = -\mathbf{I} \qquad \mathbf{c} = 0 \tag{12}$$

to obtain,

$$\arg\min_{\mathbf{z},\mathbf{t}} \quad \frac{1}{2}||\mathbf{x} - \mathbf{d} * \mathbf{z}||_2^2 + \beta||\mathbf{t}||_1$$
$$\text{s.t.} \quad \mathbf{z} = \mathbf{t} \tag{13}$$

By introducing a proxy $\mathbf{t}$, the loss function can be treated as a sum of functions of two independent variables, and with the addition of equality constraints, the

minima of the new constrained objective is the same as the original unconstrained objective.

Whilst the individual functions may be easier to optimize, there is added complexity in enforcing the equality constraints. Taking the Lagrangian of the augmented objective,

$$\mathcal{L}(\mathbf{z}, \mathbf{t}, \boldsymbol{\lambda}) = g(\mathbf{z}) + h(\mathbf{t}) + \frac{\rho}{2}||\mathbf{z} - \mathbf{t} + \boldsymbol{\lambda}||_2^2 \qquad (14)$$

the solution is to minimize the primal variables, and maximize the dual variables,

$$\mathbf{z}^\star, \mathbf{t}^\star, \boldsymbol{\lambda}^\star = \arg\max_{\boldsymbol{\lambda}} \left(\arg\min_{\mathbf{z},\mathbf{t}} \mathcal{L}(\mathbf{z}, \mathbf{t}, \boldsymbol{\lambda}) \right) \qquad (15)$$

Optimizing in alternation (which accounts for the term *alternating direction*) yields a strategy for updating the variables involving a function of a single variable plus a proximal term binding all variables,

$$\mathbf{z}^{k+1} = \arg\min_{\mathbf{z}} \quad g(\mathbf{z}^k) + \frac{\rho}{2}||\mathbf{z}^k - (\mathbf{t} - \boldsymbol{\lambda})||_2^2 \qquad (16)$$

$$\mathbf{t}^{k+1} = \arg\min_{\mathbf{t}} \quad h(\mathbf{t}^k) + \frac{\rho}{2}||\mathbf{t}^k - (\mathbf{z} + \boldsymbol{\lambda})||_2^2 \qquad (17)$$

$$\boldsymbol{\lambda}^{k+1} = \boldsymbol{\lambda}^k + \rho(\mathbf{z} - \mathbf{t}) \qquad (18)$$

The intuition behind this strategy is to find a set of model parameters which don't deviate too far from the regularization parameters, then visa versa, to find a set of regularization parameters which are close to the model parameters. The Lagrange variables impose a linear descent direction that force the two primal variables to equality over time.

For this strategy to be effective, the sum of the function g or h and an isotropic least-squares must be easy to solve.

Substituting the convolutional sparse inference objective into Eqn. 16,

$$\mathbf{z}^{k+1} = \arg\min_{\mathbf{z}} \frac{1}{2}||\mathbf{x} - \mathbf{d} * \mathbf{z}^k||_2^2 + \frac{\rho}{2}||\mathbf{z}^k - (\mathbf{t} - \boldsymbol{\lambda})||_2^2 (19)$$

$$\mathbf{t}^{k+1} = \arg\min_{\mathbf{t}} \beta||\mathbf{t}^k||_1 + \frac{\rho}{2}||\mathbf{t}^k - (\mathbf{z} + \boldsymbol{\lambda})||_2^2 \qquad (20)$$

The $\mathbf{z}$ update takes the form of generalized Tikhonov regularization. In the $\mathbf{t}$ update, the L_1-regularizer can now be treated independently of the model-fitting term, and importantly the added proximal term is an isotropic Gaussian, so each pixel of $\mathbf{t}$ can be updated independently,

$$t^{k+1} = \arg\min_{t} \quad \beta|t^k| + \frac{\rho}{2}(t^k - z + \lambda)^2 \qquad (21)$$

the solution to which is the soft thresholding operator,

$$t^\star = \operatorname{sgn}(z + \lambda) \cdot \max\left\{ |z + \lambda| - \frac{\beta}{\rho}, 0 \right\} \qquad (22)$$

In the ADMM, the minimizer $\mathbf{t}^\star$ is exactly sparse, whilst the minimizer $\mathbf{z}^\star$ is only close to sparse. If exact sparsity is a concern in the problem domain, $\mathbf{t}^\star$ should be retained at the point of convergence.

3.2 FISTA/Proximal Gradient

Proximal gradient methods are a close parallel to ADMMs. They generalize the problem of projecting a point onto a convex set, and often admit closed-form solutions. (Parikh and Boyd, 2013) present a thorough review of proximal algorithms and their relation to ADMMs. The convolutional Lasso problem introduces the splitting,

$$g(\mathbf{z}) = \frac{1}{2}||\mathbf{x} - \mathbf{d} * \mathbf{z}||_2^2 \qquad h(\mathbf{z}) = \beta||\mathbf{z}||_1 \qquad (23)$$

with gradient and proximal operator,

$$\nabla g(\mathbf{z}) = \mathbf{D}^T(\mathbf{D}\mathbf{z} - \mathbf{x}) \qquad \operatorname{prox}_h(\mathbf{z}) = \mathbf{S}(\mathbf{z}) \qquad (24)$$

where $\mathbf{S}$ is the soft thresholding operator of Eqn. 22. In this case, the proximal algorithm is finding the closest minimizer to the convolutional least squares model fitting term that projects onto the L_1-ball (with radius proportional to β).

The FISTA algorithm of (Beck and Teboulle, 2009) presents an efficient update strategy for solving this problem by incorporating an optimal first-order method in the gradient computation (discussed in §4.2).

FISTA and ADMMs both have similar computational complexity, each requiring updates to a least-squares problem and evaluation of a soft thresholding operator. One disadvantage of FISTA is the requirement of gradient-based updates to the functional term. ADMMs, on the other hand, make a more general set of assumptions, requiring only that the objective value in each iteration is reduced. In cases where solving the objective is similar in complexity to evaluating the gradient, ADMMs may converge faster.

3.3 Iterative Optimziation

So far we have neglected to show how convolution in Eqn. 11 and Eqn. 23 is actually performed.

The classical approach to convolution is to assume Dirichlet boundary conditions, *i.e.* that values outside the domain of consideration are zero. In such a case $\mathbf{x} \in \mathbb{R}^{P \times Q}$, $\mathbf{d} \in \mathbb{R}^{R \times S}$ and $\mathbf{z} \in \mathbb{R}^{P \times Q}$. Another approach is to take only the convolution between the fully overlapping portions of the signals – 'valid' convolution – which results in $\mathbf{z} \in \mathbb{R}^{(P+R-1) \times (Q+S-1)}$.

In both approaches, the signals *must* be convolved in the spatial domain, which is an $\mathcal{O}(PQRS)$ operation.

One way to alleviate the computational cost is to assume periodic extension of the signal, where convolution is then diagonalized by the Fourier transform,

$$\mathbf{d} * \mathbf{z} = \mathcal{F}^{-1}\{\mathcal{F}(\mathbf{d}) \cdot \mathcal{F}(\mathbf{z})\} \qquad (25)$$

where $\mathbf{d} \in \mathbb{R}^{P \times Q}$ and $\mathbf{z} \in \mathbb{R}^{P \times Q}$ (see §4.5 for the correct method of padding $\mathbf{d}$ to size). This method of convolution has cost $\mathcal{O}(PQ \log(PQ))$.

Both ADMMs and FISTA can take advantage of iterative methods, by taking accelerated (proximal) gradient steps. In the next section we show how solving the entire system of equations in the Fourier domain is only slightly more complex than performing a single gradient step, and can lead to faster overall convergence of the ADMM. An important distinction between FISTA and ADMMs is that FISTA is constrained to use gradient updates – it cannot take advantage of direct solvers.

3.4 Direct Optimization

Given Dirichlet boundary assumptions, a filter kernel $\mathbf{d}$ can be represented (in 2D) as a block-Toeplitz-Toeplitz matrix such that $g(\mathbf{z})$ of Eqn. 23 becomes,

$$g(\mathbf{z}) = \frac{1}{2}||\mathbf{x} - \mathbf{Dz}||_2^2 \tag{26}$$

where,

$$\mathbf{D} = [\mathbf{D}_1 \dots \mathbf{D}_N] \qquad \mathbf{z} = \begin{bmatrix} \mathbf{z}_1 \\ \vdots \\ \mathbf{z}_N \end{bmatrix} \tag{27}$$

which is the canonical least-squares problem. There are a plethora of direct solvers for this problem, however unlike Toeplitz matrices, there are no known fast ($< \mathcal{O}(n^2)$) methods for inverting block-Toeplitz-Toeplitz matrices, so general purpose solvers must be used. This requires constructing the full $\mathbf{D}$. For a $P \times Q$ input, and N filters each of support $R \times S$, the matrix $\mathbf{D}$ will have $\mathcal{O}(NPQRS)$ non-zero values.

One way to alleviate the computational cost is to assume periodic extension of the signal, where convolution is then diagonalized by the Fourier transform.

This provides an efficient strategy for inverting the system,

$$g(\hat{\mathbf{z}}) = \frac{1}{2}||\hat{\mathbf{x}} - \hat{\mathbf{D}}\hat{\mathbf{z}}||_2^2 \tag{28}$$

where,

$$\hat{\mathbf{D}} = \left[\text{diag}(\hat{\mathbf{d}}_1) \dots \text{diag}(\hat{\mathbf{d}}_N) \right] \qquad \hat{\mathbf{z}} = \begin{bmatrix} \hat{\mathbf{z}}_1 \\ \vdots \\ \hat{\mathbf{z}}_N \end{bmatrix} \tag{29}$$

The equivalence between Eqn. 23 and Eqn. 28 relies upon Parseval's theorem,

$$\mathbf{x}^T \mathbf{x} = K \hat{\mathbf{x}}^T \hat{\mathbf{x}} \tag{30}$$

where K is a constant scaling factor between the domains. Since the L_2 norm is rotation invariant, the minimizer of Eqn. 28 is the Fourier transform of the minimize of Eqn. 23.

The system of Eqn. 19 can be solved directly in an efficient manner in the Fourier domain by observing that each frequency band in a single image $\mathbf{x}$ is related only by the same frequency band in each of the filters and coefficients, *e.g.*,

$$\hat{\mathbf{x}}_{1,1} = \hat{\mathbf{d}}_{1,n}\hat{\mathbf{z}}_{n,1} \quad \forall \, n \in 1 \dots N \tag{31}$$

Finding the optima of $\hat{\mathbf{z}}$ across all n channels for frequency band 1 now involves

$$\hat{\mathbf{z}}_{n,1} = \left(\hat{\mathbf{d}}_{1,n}^T \hat{\mathbf{d}}_{1,n} + \rho \mathbf{I}_{n,1} \right)^{-1} \left(\hat{\mathbf{x}}_{1,1} \hat{\mathbf{d}}_{1,n}^T + \rho(\hat{\mathbf{t}}_{n,1} - \hat{\boldsymbol{\lambda}}_{n,1}) \right) \tag{32}$$

Note that $\hat{\mathbf{d}}_{1,n}^T \hat{\mathbf{d}}_{1,n}$ is a rank-1 matrix. Thus from the matrix inversion lemma (Appendix B),

$$\left(\hat{\mathbf{d}}_{1,n}^T \hat{\mathbf{d}}_{1,n} + \rho \mathbf{I}_{n,1} \right) = \frac{1}{\rho} \left(\mathbf{I}_{n,1} - \frac{1}{\rho + \hat{\mathbf{d}}_{1,n}\hat{\mathbf{d}}_{1,n}^T} \right) \hat{\mathbf{d}}_{1,n}^T \hat{\mathbf{d}}_{1,n}$$

$$= \frac{1}{\rho}(\mathbf{I}_{n,1} - K)\hat{\mathbf{d}}_{1,n}^T \hat{\mathbf{d}}_{1,n} \tag{33}$$

where K is a scalar, since $\mathbf{d}_{1,n}\mathbf{d}_{1,n}^T$ is a scalar. The block of $\mathbf{z}_{n,1}$ can thus be solved by,

$$\hat{\mathbf{z}}_{n,1} = \frac{1}{\rho}(\mathbf{I}_{n,1} - K)\hat{\mathbf{d}}_{1,n}^T \hat{\mathbf{d}}_{1,n} \left(\hat{\mathbf{x}}_{1,1}\hat{\mathbf{d}}_{1,n}^T + \rho(\hat{\mathbf{t}}_{n,1} - \hat{\boldsymbol{\lambda}}_{n,1}) \right) \tag{34}$$

which is just a series of multiplications, so a solution can be found in $\mathcal{O}(n^2)$ time – no inversion is required.

The assumption of periodic extension is sometimes not indicative of the structure observed in the signal of interest, and can cause artefacts along the boundaries. In such a case, a more sensible assumption is to assume Neumann boundary conditions, or *symmetric* reflection across the boundary. This assumption tends to minimize boundary distortion for small displacements. The resulting matrix can be diagonalised by the discrete cosine transform (DCT). There exists a generalization of Parseval's theorem that extends to the DCT, however some care must be taken with understanding the nuances between the 16 types of DCT. (Martucci, 1993) provides a comprehensive exposé on the topic.

4 Solving for Filters

Solving for the filters $\mathbf{d}$ involves optimizing,

$$\arg\min_{\mathbf{d}} \quad \frac{1}{2}||\mathbf{x} - \mathbf{d} * \mathbf{z}||_2^2$$
$$\text{s.t.} \quad ||\mathbf{d}||_2^2 \leq 1 \tag{35}$$

This is a classical least squares objective with norm inequality constraints. Norm constraints on the bases are necessary since there always exists a linear transformation of $\mathbf{d}$ and $\mathbf{z}$ which keeps $(\mathbf{d} * \mathbf{z})$ unchanged whilst making $\mathbf{z}$ approach zero. Inequality constraints are sufficient since deflation of $\mathbf{z}$ will always cause $\mathbf{d}$ to lie on the constraint boundary $||\mathbf{d}||_2^2 = 1$ (whilst forming a convex set).

Eqn. 35 is a quadratically constrained quadratic program (QCQP), which is difficult to optimize in general. Each of the following methods relax this form in one way or another to make it more tractable to solve.

The convolutional form of the least squares term also poses some challenges for optimization, since (i) the filter has smaller support than the image it is being convolved with, and (ii) forming an explicit multiplication between the filters and a convolutional matrix form of the images is prohibitively expensive (as per §3.4).

In the case of multiple images and filters,

$$\arg\min_{\mathbf{d}} \quad \frac{1}{2}||\sum_{i=1}^{M}\big(\mathbf{x}_i - \sum_{j=1}^{N}(\mathbf{d}_j * \mathbf{z}_{i,j})\big)||_2^2$$

$$\text{s.t.} \qquad ||\mathbf{d}_j||_2^2 \leq 1 \ \ \forall \, j \in 1 \ldots N \tag{36}$$

one can see that all filters are jointly involved in reconstructing the inputs, and so must be updated jointly.

4.1 Gradient Descent

The gradient of the objective is given by,

$$\nabla f_i = \left(\sum_{j=1}^{N}\mathbf{d}_j * \sum_{k=1}^{M}(\mathbf{z}_{i,j} \star \mathbf{z}_{i,k})\right) - \sum_{k=1}^{M}(\mathbf{x}_i \star \mathbf{z}_{i,j}) \tag{37}$$

This involves collecting the correlation statistics across the coefficient maps and images. Since the filters are of smaller support than the coefficient maps and images, we collect only "valid" statistics, or regions that don't incur boundary effects. In the autocorrelation of $\mathbf{z}$, one of the arguments must be zero padded to the appropriate size.

Given the gradient direction, updating the the filters involves,

$$\mathbf{d}^{k+1} = \mathbf{d}^k - t\nabla f \tag{38}$$

Computing the step size, t, can be done either via line search or by solving a 1D optimization problem which minimizes the reconstruction error in the gradient direction,

$$\arg\min_{t} \ \ ||\mathbf{x} - \underbrace{\big(\mathbf{d}^k - t\nabla f\big)}_{\mathbf{d}^{k+1}} * z||_2^2 \tag{39}$$

the closed-form solution to which is,

$$t = \frac{(\mathbf{x} - \mathbf{d} * \mathbf{z})^T(\nabla f * \mathbf{z})}{(\nabla f * \mathbf{z})^T(\nabla f * \mathbf{z})} \tag{40}$$

and involves evaluating only $2N$ convolutions (N if the $\mathbf{x} - \mathbf{d} * \mathbf{z}$ term has been previously computed as part of a stopping criteria, *etc.*). In the case of a large number of inputs M, stochastic gradient descent is typically used (Mairal et al., 2009).

Typically after each gradient step, if the new iterate exists outside the L_2 unit ball, the result is projected back onto the ball. This is not strictly the correct way to enforce the norm-constraints, however it works without side-effects in practice. The ADMM (§4.3) and Lagrange dual (§4.4) methods on the other hand, both solve for the norm constraints exactly.

4.2 Nesterov's Accelerated Gradient Descent

Prolific mathematician Yurii Nesterov introduced an optimal[2] first-order method for solving smooth convex functions (Nesterov, 1983). Convolutional least squares objectives require a straightforward application of accelerated proximal gradient (APG).

Without laboring on the details introduced by Nesterov, the method involves iteratively placing an isotropic quadratic tangent to the current gradient direction, then shifting the iterate to the minima of the quadratic. The curvature of the quadratic is computed by estimating the Lipschitz smoothness of the objective.

Checking for Lipschitz smoothness feasibility using backtracking requires two projections per iteration. This can be costly since each projection involves multiple convolutions. In practice we find this method no faster than regular gradient descent with optimal step-size calculation.

4.3 ADMM Partitioning

As per §3.1, treating convolution in the Fourier domain can lead to efficient direct optimization. Unlike solving for the coefficients, however, the learned filters must be constrained to be small support, and there is no way to do this explicitly via Fourier convolution.

---

[2] Optimal in the sense that it has a worst-case convergence rate that cannot be improved whilst remaining first-order.

An approach to handling this is via ADMMs again,

$$\arg\min_{\mathbf{d},\mathbf{s}} \quad \frac{1}{2}||\hat{\mathbf{x}} - \hat{\mathbf{Z}}\hat{\mathbf{s}}||_2^2$$

$$\text{s.t.} \qquad \mathbf{d}_j = \boldsymbol{\Phi}^T\hat{\mathbf{s}}_j \quad \forall\, j \in 1\ldots N$$

$$\qquad\qquad ||\mathbf{d}_j||_2^2 \leq 1 \quad \forall\, j \in 1\ldots N \tag{41}$$

where,

$$\hat{\mathbf{x}} = \begin{bmatrix} \hat{\mathbf{x}}_1 \\ \vdots \\ \hat{\mathbf{x}}_M \end{bmatrix} \qquad \hat{\mathbf{d}} = \begin{bmatrix} \hat{\mathbf{d}}_1 \\ \vdots \\ \hat{\mathbf{d}}_N \end{bmatrix}$$

$$\hat{\mathbf{Z}} = \begin{bmatrix} \text{diag}(\hat{\mathbf{z}}_{1,1}) & \ldots & \text{diag}(\hat{\mathbf{z}}_{1,N}) \\ \vdots & \ddots & \\ \text{diag}(\hat{\mathbf{z}}_{M,1}) & & \text{diag}(\hat{\mathbf{z}}_{M,N}) \end{bmatrix} \tag{42}$$

and $\boldsymbol{\Phi}$ is a submatrix of the Fourier matrix that corresponds to a small spatial support transform.

Intuitively, we are trying to learn a set of filters $\hat{\mathbf{s}}$ that minimize reconstruction error in the Fourier domain and are small support in the spatial domain.

Taking the augmented Lagrangian of the objective and optimizing over $\mathbf{d}$ and $\mathbf{s}$ in alternation yields the update strategy,

$$\hat{\mathbf{s}}^{k+1} = \arg\min_{\hat{\mathbf{s}}} ||\hat{\mathbf{x}} - \hat{\mathbf{Z}}\hat{\mathbf{s}}^k||_2^2 + \frac{\rho}{2}||\hat{\mathbf{s}}^k - (\boldsymbol{\Phi}\mathbf{d} - \hat{\boldsymbol{\lambda}})||_2^2 \tag{43}$$

$$\mathbf{d}^{k+1} = \arg\min_{\mathbf{d}} ||\boldsymbol{\Phi}\mathbf{d}^k - (\hat{\mathbf{s}} + \hat{\boldsymbol{\lambda}})||_2^2$$

$$\text{s.t.} \qquad ||\mathbf{d}_j||_2^2 \leq 1 \quad \forall\, j \in 1\ldots N \tag{44}$$

$$\boldsymbol{\lambda}^{k+1} = \boldsymbol{\lambda}^k + \rho(\boldsymbol{\Phi}\mathbf{d} - \hat{\mathbf{s}}) \tag{45}$$

In a similar fashion to Eqn. 32, each frequency component in $\hat{\mathbf{s}}$ jointly across all filters can be solved for independently via a variable reordering to produce $P \times Q$ dense systems of equations.

Solving for $\mathbf{d}$ appears more involved, however the unconstrained loss function can be minimized in closed-form,

$$\mathbf{d}^\star = \mathbf{d}^T\boldsymbol{\Phi}^T\boldsymbol{\Phi}\mathbf{d} - 2\mathbf{d}^T\boldsymbol{\Phi}^T(\hat{\mathbf{s}} - \hat{\boldsymbol{\lambda}}) + c$$

$$= \boldsymbol{\Phi}^T(\hat{\mathbf{s}} - \hat{\boldsymbol{\lambda}}) \tag{46}$$

since $\boldsymbol{\Phi}$ is an orthonormal matrix, and thus $\boldsymbol{\Phi}^T\boldsymbol{\Phi} = \mathbf{I}$. Further, the matrix multiplication can instead be replaced by the inverse Fourier transform, followed by a selection operator $\mathcal{M}$ which keeps only the small support region,

$$\mathbf{d}^\star = \mathcal{M}\left(\mathcal{F}^{-1}\{\hat{\mathbf{s}} - \hat{\boldsymbol{\lambda}}\}\right) \tag{47}$$

Handling the inequality constraints is now trivial. Since $\boldsymbol{\Phi}$ is orthonormal (implying an isotropic regression problem), projecting the optimal solution to the

unconstrained problem onto the L_2 ball,

$$\mathbf{d}^\star = \begin{cases} ||\mathbf{d}_k^\star||_2^{-2}\mathbf{d}_k^\star, & \text{if } ||\mathbf{d}_k^\star||_2^2 \geq 1 \\ \mathbf{d}_k^\star, & \text{otherwise} \end{cases} \tag{48}$$

is *equivalent* to solving the constrained problem.

4.4 Lagrange Dual

Rather than using gradient descent with iterative projection, Eqn. 35 can be solved by taking the Lagrange dual in the Fourier domain,

$$\mathcal{L}(\hat{\mathbf{d}}, \hat{\boldsymbol{\lambda}}) = ||\hat{\mathbf{x}} - \hat{\mathbf{Z}}\mathbf{d}||_2^2 + \hat{\mathbf{d}}^T\boldsymbol{\Lambda}\hat{\mathbf{d}} - 1^T\boldsymbol{\lambda} \tag{49}$$

where $\boldsymbol{\lambda} \geq 0$ are the dual variables, and $\boldsymbol{\Lambda} = \text{diag}(\boldsymbol{\lambda})$. Finding the optimum involves minimizing the primal variables, and maximizing the dual variables,

$$\hat{\mathbf{d}}^\star, \hat{\boldsymbol{\lambda}}^\star = \arg\max_{\hat{\boldsymbol{\lambda}}}\left(\arg\min_{\hat{\mathbf{d}}} \mathcal{L}(\hat{\mathbf{d}}, \hat{\boldsymbol{\lambda}})\right) \tag{50}$$

The closed-form solution to the primal variables is,

$$\hat{\mathbf{d}}^\star = (\hat{\mathbf{Z}}^T\hat{\mathbf{Z}} + \boldsymbol{\Lambda})^{-1}(\hat{\mathbf{Z}}^T\hat{\mathbf{x}}) \tag{51}$$

Substituting this expression for $\hat{\mathbf{d}}$ into the Lagrangian, we analytically derive the dual optimization problem.

4.5 Convolving with Small Support Filters in the Fourier Domain

In order to convolve two signals in the Fourier domain, their lengths must commute. This involves padding the shorter signal to the length of the longer. Some care must be taken to avoid introducing phase shifts into the response, however. Given a 2D filter $\mathbf{z} \in \mathbb{R}^{P,Q}$, we can partition it into 4 blocks,

$$\mathbf{z} = \begin{vmatrix} \mathbf{z}_{1,1} & \mathbf{z}_{1,2} \\ \mathbf{z}_{2,1} & \mathbf{z}_{2,2} \end{vmatrix} \tag{52}$$

where, in the case of odd-sized filters, the blocks are partitioned *above and to the left* of the central point. Given a 2D image $\mathbf{x} \in \mathbb{R}^{M,N}$, the padded representation $\mathbf{z}^\star \in \mathbb{R}^{M,N}$ can thus be formed as,

$$\mathbf{z} = \begin{vmatrix} \mathbf{z}_{2,2} & & & \mathbf{z}_{2,1} \\ & \ddots & & \\ & \ldots & \mathbf{0} & \ldots & \\ & & \ddots & \\ \mathbf{z}_{1,2} & & & \mathbf{z}_{1,1} \end{vmatrix} \tag{53}$$

This transform is illustrated in Fig. 4. For a comprehensive guide to Fourier domain transforms and identities, including appropriate handling of boundary effects and padding, see (Oppenheim et al., 1996).
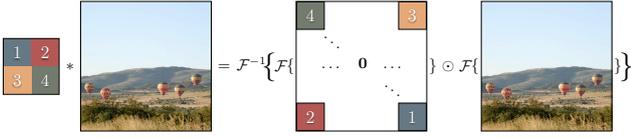
Fig. 4 Swapping quadrants and padding the filter to the size of the image permits convolution in the Fourier Domain

5 Stopping Criteria

For the gradient-based algorithms – gradient descent, APG and FISTA – a sufficient stopping criteria is to threshold the residual between two iterates,

$$||\mathbf{x}^{k+1} - \mathbf{x}^k||_2^2 \leq \epsilon \tag{54}$$

where $\mathbf{x}$ is the variable being minimized.

Estimating convergence of the ADMM-based methods is more involved, including deviation from primal feasibility,

$$||\mathbf{d} - \mathbf{s}||_2^2 \leq \epsilon \, , \qquad ||\mathbf{z} - \mathbf{t}||_2^2 \leq \epsilon \tag{55}$$

and dual feasibility,

$$||\boldsymbol{\lambda}^{k+1} - \boldsymbol{\lambda}^k||_2^2 \leq \epsilon \tag{56}$$

However, in practice it is usually sufficient to measure only primal feasibility, since if the iterates have reached primal feasibility, dual feasibility is unlikely to improve (this is doubly true when using a strategy for increasing ρ).

6 Applications

6.1 Example 1 - Image and Video Compression

Many image and video coding algorithms such as JPEG (Wallace, 1991) and H.264 (Sullivan, 2004) discretize each image into blocks which are transformed, quantized and coded independently.

Using convolutional sparse coding, an entire image $\mathbf{x}$ can be coded onto a basis $\mathbf{d}$ with sparse reconstruction coefficients $\mathbf{z}$. Quality and size can be controlled via the β parameter. The basis can either be specific to the image, or a generic basis (such as Gabor) which is part of the coding spec and need not be transferred with the image data.

Since the coefficients of $\mathbf{z}$ are exactly sparse by virtue of the soft-thresholding operator, the representation can make effective use of run-length and Huffman entropy coding techniques. [3]

---

[3] JPEG also uses run-length coding but its efficiency is a function of the quantization artefacts.

To reconstruct the image $\mathbf{x}_r$, the decoder simply convolves the bases with the transmitted coefficients,

$$\mathbf{x}_r = \sum_{j=1}^{N} \mathbf{d}_j * \mathbf{z}_j \tag{57}$$

This matches the media model well: media is consumed more frequently than it is created, so encoding can be costly (in this case an inverse inference problem) but decoding should be fast – convolutional primitives are hardware-accelerated on almost all modern chipsets.

6.2 Example 2 - A basis for Natural Images

The receptive fields of simple cells in the mammalian primary visual cortex can be characterized as being spatially localized, oriented and bandpass. (Olshausen and Field, 1996) hypothesized that such fields could arise spontaneously in an unsupervised strategy for maximizing the sparsity of the representation. Sparsity can be interpreted biologically as a metabolic constraint - firing only a few neurons in response to a stimulus is clearly more energy efficient than firing a large number.

(Olshausen and Field, 1996) use traditional patch-based sparse coding to solve for a set of basis functions. The famous result is that the basis functions resemble Gabor filters. However, a large number of the bases learned are translations of others – an artefact of sampling and treating each image patch independently.

It is well-understood that the statistics of natural images are translation invariant (Hyvärinen et al., 2009), *i.e.* the covariance of natural images depends only on the distance,

$$\Sigma \left(\mathbf{I}(x,y), \, \mathbf{I}(x',y') \right) = f \left((x-x')^2 + (y-y')^2 \right) \tag{58}$$

Thus it makes sense to code natural images in a manner that does not depend on exact position of the stimulus. Convolutional sparse coding permits this, and as a result produces a more varied range of basis elements than simple Gabor filters when coding natural images. The sparsity pattern of convolutional coefficients also has a mapping onto the receptive fields of active neurons in V1.[4]

6.3 Example 3 - Structure from Motion

Trajectory basis Non-Rigid Structure from Motion (NRSfM) refers to the process of reconstructing the mo-

---

[4] Unlike many higher regions within the visual cortex, V1 is retinotopic – the spatial location of stimulus in the visual world is highly correlated with the spatial location of active neurons within V1.

tion of 3D points of a non-rigid object from only their 2D projected trajectories.

Reconstruction relies on two inherently conflicting factors: (i) the condition of the composed camera and trajectory basis matrix, and (ii) whether the trajectory basis has enough degrees of freedom to model the 3D point trajectory. Typically, (i) is improved with a low-rank basis, and (ii) is improved with a higher-rank basis.

The Discrete Cosine Transform (DCT) basis has traditionally been used as a generic basis for encoding motion trajectories, however choosing the correct rank has been a difficult problem. (Zhu and Lucey, 2014) proposed the use of convolutional sparse coding to learn a compact basis that could model the trajectories taken from a corpus of training data.

Learning the basis proceeds as per usual,

$$\arg\min_{\mathbf{d},\mathbf{z}} \quad \frac{1}{2}\sum_{i=1}^{M}||\big(\mathbf{x}_i - \sum_{j=1}^{N}(\mathbf{d}_j * \mathbf{z}_{i,j})\big)||_2^2$$
$$+ \beta \sum_{i=1}^{M}\sum_{j=1}^{N}||\mathbf{z}_{i,j}||_1$$
$$\text{s.t.} \qquad ||\mathbf{d}_j||_2^2 \leq 1 \;\; \forall\, j \in 1 \ldots N \qquad (59)$$

where each $\mathbf{x}_i$ is a 1D trajectory of arbitrary length, $\mathbf{d}$ the trajectory basis being learned, and $\mathbf{z}$ the sparse reconstruction coefficients.

Given the convolutional trajectory basis $\mathbf{d}$, reconstructing the sparse coefficients for the $3D$ trajectory from $2D$ observations involves,

$$\mathbf{z}^* = \arg\min_{\mathbf{z}} \;\; ||\mathbf{z}||_1$$
$$\text{s.t.} \qquad \underbrace{\mathbf{Q}\mathbf{x}}_{\mathbf{u}} = \mathbf{Q}\sum_{j=1}^{N}\mathbf{d}_j * \mathbf{z}_j \qquad (60)$$

where $\mathbf{u}$ are the $2D$ observations of the $3D$ points $\mathbf{x}$ that have been imaged by the camera matrices $\mathbf{Q}$, one for each frame in the trajectory,

$$\mathbf{Q} = \begin{bmatrix} \mathbf{Q}_1 & & \\ & \ddots & \\ & & \mathbf{Q}_F \end{bmatrix} \qquad (61)$$

In single view reconstruction, back-projection is typically enforced as a constraint, and the objective is to minimize the number of non-zero coefficients in the reconstructed $3D$ trajectory that satisfy this constraint.

A convolutional sparse coded basis produces less $3D$ reconstruction error than previously explored bases, including one learned from patch-based sparse coding, and a generic DCT basis. This illustrates convolutional sparse coding's ability to learn low rank structure from

misaligned trajectories stemming from the same underlying dynamics (*e.g.* articulated human motion).

6.4 Example 4 - Mid-level Generative Parts

Zeiler *et al.* show how a cascade of convolutional sparse coders can be used to build robust, unsupervised mid-level representations, beyond the edge primitives of §6.2.

The convolutional sparse coder at each level of the hierarchy can be defined as,

$$C_l(\mathbf{d}^l, \mathbf{z}^l) \;\; = \frac{1}{2}\sum_{i=1}^{M}||\underbrace{f_s(\mathbf{z}_i^{l-1})}_{\mathbf{x}_i} - \sum_{j=1}^{M}(\mathbf{d}_j^l * \mathbf{z}_{i,j}^l)||_2^2$$
$$+ \beta \sum_{i=1}^{M}\sum_{j=1}^{N}||\mathbf{z}_{i,j}^l||_1 \qquad (62)$$

where the extra superscripts on each element indicate the layer to which they are native.

The inputs $\mathbf{x}$ to each layer are the sparse coefficients of the previous layer, $\mathbf{z}^{l-1}$ after being passed through a pooling/subsampling operation $f_s(\cdot)$. For the first layer, $\mathbf{z}^{l-1} = \mathbf{x}$, *i.e.* the input image.

The idea behind this coding strategy is that structure within the signal is progressively gathered at a higher and higher level, initially with edge primitives, then mergers between these primitives into line-segments, and eventually into recurrent object parts.

Layers of convolutional sparse coders have also been used to produce high quality latent representations for convolutional neural networks (Lee et al., 2009; Kavukcuoglu et al., 2010; Chen et al., 2013), though fully-supervised back-propagation across layers has become popular more recently (Krizhevsky et al., 2012).

6.5 Example 5 - Single Image Super-Resolution

Single-Image Super Resolution (SISR) is the process of reconstructing a high-resolution image from an observed low-resolution image. SISR can be cast as the inverse problem,

$$\mathbf{y} = \mathbf{D}\mathbf{B}\mathbf{x} \qquad (63)$$

where $\mathbf{x}$ is the latent high-resolution image that we wish to recover, $\mathbf{B}$ is an anti-aliasing filter, $\mathbf{D}$ is a downsampling matrix and $\mathbf{y}$ is the observed low-resolution image. The system is underdetermined, so there exist infinitely many solutions to $\mathbf{x}$. A strategy for performing SISR is

via a straightforward convolutional extension of (Yang et al., 2010),

$$\underset{\mathbf{d}_L,\mathbf{d}_H,\mathbf{z}}{\arg\min} \sum_{i=1}^{M} ||\mathbf{DB}\mathbf{x}_i - \sum_{j=1}^{N}(\mathbf{d}_{L,j}*\mathbf{Dz}_{i,j})||_2^2$$

$$+ ||\mathbf{x}_i - \sum_{j=1}^{N}(\mathbf{d}_{H,j}*\mathbf{z}_{i,j})||_2^2$$

$$+ \beta \sum_{i=1}^{M}\sum_{j=1}^{M} ||\mathbf{z}_{i,j}||_1$$

$$\text{s.t.} \quad ||\mathbf{d}_{L,j}||_2^2 \le 1 \quad \forall\, j \in 1\dots M$$
$$||\mathbf{d}_{H,j}||_2^2 \le 1 \quad \forall\, j \in 1\dots M \qquad (64)$$

where $\mathbf{x}_i$ and $\mathbf{DB}\mathbf{x}_i$ are a high-resolution and derived low-resolution training pair, $\mathbf{D}$ is the downsampling filter as before and $\mathbf{z}$ are a common set of coefficients that tie the two representations together. The dictionaries $\mathbf{d}_L$ and $\mathbf{d}_H$ learn a mapping between low- and high-resolution image features.

Given a new low-resolution input image $\mathbf{x}_L$, the sparse coefficients are first inferred with respect to the low-resolution basis,

$$\mathbf{z}^\star = \arg\min_{\mathbf{z}} ||\mathbf{x}_L - \sum_{j=1}^{N}(\mathbf{d}_{L,j}*\mathbf{Dz}_j)||_2^2 + \beta||\mathbf{z}||_1 \qquad (65)$$

and then convolved with the high-resolution basis to reconstruct the high-resolution image,

$$\mathbf{x}_H = \sum_{j=1}^{N}(\mathbf{d}_{H,j}*\mathbf{z}_j) \qquad (66)$$

6.6 Example 6 - Visualizing Object Detection Features

(Vondrick et al., 2013) presented a method for visualizing HOG features via the inverse mapping,

$$\phi^{-1}(\mathbf{y}) = \arg\min_{\mathbf{x}} ||\phi(\mathbf{x}) - \mathbf{y}||_2^2 \qquad (67)$$

where $\mathbf{x}$ is the image to recover, and $\mathbf{y} = \phi(\mathbf{x})$ is the mapping of the image into HOG space. Direct optimization of this objective is difficult, since it is highly nonlinear through the HOG operator $\phi()$, *i.e.* multiple distinct images can map to the same HOG representation.

One possible approach to approximating this objective is through paired dictionary learning, in a similar manner to §6.5.

Given an image $\mathbf{x}$ and its representation $\mathbf{y} = \phi(\mathbf{x})$ in the HOG domain, we wish to find two basis sets, $\mathbf{d}_\mathbf{I}$ in the image domain and $\mathbf{d}_\phi$ in the HOG domain, and a common set of sparse reconstruction coefficients $\mathbf{z}$, such that,

$$\mathbf{x} = \sum_{j=1}^{N}(\mathbf{d}_{\mathbf{I},j}*\mathbf{z}_j) \qquad \mathbf{y} = \sum_{j=1}^{N}(\mathbf{d}_{\phi,j}*\mathbf{z}_j) \qquad (68)$$

Intuitively, the common reconstruction coefficients force the basis sets to represent the same appearance information albeit in different domains. The basis pairs thus provide *a* mapping between the domains.

Optimizing this objective is a straightforward extension of the patch based sparse coding used by (Vondrick et al., 2013),

$$\underset{\mathbf{d}_\mathbf{I},\mathbf{d}_\phi,\mathbf{z}}{\arg\min} \sum_{i=1}^{M} ||\mathbf{x}_i - \sum_{j=1}^{N}(\mathbf{d}_{\mathbf{I},j}*\mathbf{z}_{i,j})||_2^2$$

$$+ ||\phi(\mathbf{x}_i) - \sum_{j=1}^{N}(\mathbf{d}_{\phi,j}*\mathbf{z}_{i,j})||_2^2$$

$$+ \beta \sum_{i=1}^{M}\sum_{j=1}^{M} ||\mathbf{z}_{i,j}||_1$$

$$\text{s.t.} \quad ||\mathbf{d}_{\mathbf{I},j}||_2^2 \le 1 \quad \forall\, j \in 1\dots M$$
$$||\mathbf{d}_{\phi,j}||_2^2 \le 1 \quad \forall\, j \in 1\dots M \qquad (69)$$

Because we are optimizing over entire images rather than independently sampled patches, the bases learned will (i) produce a more unique mapping between pixel features and HOG features (since translations of features are not represented), and (ii) be more expressive for any given basis set size as a direct result of (i).

Image-scale optimization also reduces blocking artefacts in the image reconstructions, leading to more faithful/plausible representations, with potentially finer-grained detail.

References

Amir Beck and Marc Teboulle. A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems. *SIAM Journal on Imaging Sciences*, 2(1): 183–202, 2009. 3, 6

Stephen Boyd. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2010. 3, 5

Hilton Bristow, Anders Eriksson, and Simon Lucey. Fast Convolutional Sparse Coding. *Computer Vision and Pattern Recognition (CVPR)*, 1(2), 2013. 3

Rakesh Chalasani and JC Principe. A fast proximal method for convolutional sparse coding. *International Joint Conference on Neural Networks (IJCNN)*, 2012. 3

B Chen, Gungor Polatkan, Guillermo Sapiro, and David Blei. Deep Learning with Hierarchical Convolutional Factor Analysis. *Pattern Analysis and Machine Intelligence (PAMI)*, pages 1–30, 2013. 11

D Gabay and B Mercier. A Dual Algorithm for the Solution of Nonlinear Variational Problems via Finite Element Approximations. *Computers and Mathematics with Applications*, 1976. 5

R Glowinski and A Marroco. Sur L'Approximation, par Elements Finis d'Ordre Un, et la Resolution, par Penalisation-Dualite, d'une Classe de Problemes de Dirichlet non Lineares. *Revue Francaise d'Automatique*, 1975. 5

Roger Grosse, Rajat Raina, Helen Kwong, and Andrew Y Ng. Shift-invariant sparse coding for audio classification. *UAI*, 2007. 2, 3

Aapo Hyvärinen, Jarmo Hurri, and Patrik O. Hoyer. Natural Image Statistics A probabilistic approach to early computational vision. *Computational Imaging and Vision*, 39, 2009. 10

Koray Kavukcuoglu, Pierre Sermanet, Y-lan Boureau, K. Gregor, M. Mathieu, and Y. LeCun. Learning convolutional feature hierarchies for visual recognition. *Advances in Neural Information Processing Systems (NIPS)*, (1):1–9, 2010. 2, 11

Alex Krizhevsky, I Sutskever, and G Hinton. Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems (NIPS)*, pages 1–9, 2012. 11

Honglak Lee, A. Battle, R. Raina, and A.Y. Ng. Efficient sparse coding algorithms. *Advances in Neural Information Processing Systems*, 19:801, 2007. 2

Honglak Lee, Roger Grosse, Rajesh Ranganath, and Andrew Y. Ng. Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. *International Conference on Machine Learning (ICML)*, pages 1–8, 2009. 3, 11

MS Lewicki and TJ Sejnowski. Coding time-varying signals using sparse, shift-invariant representations. *NIPS*, 1999. 2

Julien Mairal, Francis Bach, and J Ponce. Online dictionary learning for sparse coding. *Conference on Machine Learning*, 2009. 8

SA Martucci. *Symmetric convolution and the discrete sine and cosine transforms: Principles and applications*. PhD thesis, Georgia Institute of Technology, 1993. 7

Morten Mø rup, Mikkel N Schmidt, and Lars Kai Hansen. Shift invariant sparse coding of image and music data. *Journal of Machine Learning*, pages 1–14, 2008. 2

Y Nesterov. A method of solving a convex programming problem with convergence rate O(1/kˆ2). *Soviet Mathematics Doklady*, 1983. 8

B.A. Olshausen and D.J. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, 1996. 2, 10

Bruno A Olshausen and David J Field. Sparse Coding with an Overcomplete Basis Set: A Strategy Employed by V1? *Science*, 37(23):3311–3325, 1997. 2

Alan V Oppenheim, Alan S Willsky, and with S. Hamid. *Signals and Systems (2nd Edition)*. Prentice Hall, 1996. 9

Neal Parikh and Stephen Boyd. Proximal Algorithms. 1(3):1–108, 2013. 6

Eero Simoncelli, William Freeman, Edward Adelson, and David Heeger. Shiftable multiscale transforms. *Information Theory*, 1992. 2

K Sohn and DY Jung. Efficient learning of sparse, distributed, convolutional feature representations for object recognition. *International Conference on Computer Vision (ICCV)*, 2011. 3

GJ Sullivan. The H. 264/AVC advanced video coding standard: Overview and introduction to the fidelity range extensions. In *Conference on Applications of Digital Image Processing*, pages 1–22, 2004. 10

R Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society*, 58 (1):267–288, 1996. 2

Carl Vondrick, Aditya Khosla, Tomasz Malisiewicz, and Antonio Torralba. HOGgles: Visualizing Object Detection Features. *International Conference on Computer Vision (ICCV)*, 2013. 12

GK Wallace. The JPEG Still Picture Compression Standard. *Communications of the ACM*, pages 1–17, 1991. 10

DP Wipf, BD Rao, and Srikantan Nagarajan. Latent Variable Bayesian Models for Promoting Sparsity. *Information Theory*, 57(9):6236–6255, 2011. 5

Jianchao Yang, John Wright, Thomas Huang, and Yi Ma. Image Super-Resolution via Sparse Representation. *IEEE Transactions on Image Processing*, 19(11):1–13, November 2010. 12

MD Zeiler, Dilip Krishnan, and GW Taylor. Deconvolutional networks. *Computer Vision and Pattern Recognition (CVPR)*, 2010. 2, 3

Cun-Hui Zhang. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics (arXiv)*, 38(2):894–942, April 2010. 5

Yingying Zhu and Simon Lucey. Convolutional Sparse Coded Filters Nonrigid Structure From Motion. *Pattern Analysis and Machine Intelligence (PAMI)*, 2014. 2, 11

A Log Likelihood Interpretation

Minimizing the sparse convolutional objective has an equivalent maximum likelihood interpretation,

$$\mathbf{z}^{\star} = \arg\max_{\mathbf{z}} \quad P(\mathbf{z}|\mathbf{x}, \mathbf{d})$$
$$= \arg\max_{\mathbf{x}} \quad P(\mathbf{x}|\mathbf{d}, \mathbf{z})P(\mathbf{z}) \tag{70}$$

Replacing the product of two probability distributions with the log of their sums, and negating the expression yields,

$$\mathbf{z}^{\star} = \arg\min_{\mathbf{z}} \quad -\log(P(\mathbf{x}|\mathbf{z}, \mathbf{d})) - P(\mathbf{z}) . \tag{71}$$

The corresponding unbiased estimators are,

$$-\log(P(\mathbf{x}|\mathbf{z}, \mathbf{d})) \propto \frac{1}{2\sigma^2} ||\mathbf{x} - \mathbf{d} * \mathbf{z}||_2^2 \tag{72}$$

$$-\log(P(\mathbf{z})) \qquad \propto \frac{1}{2b} ||\mathbf{z}||_1 \tag{73}$$

After choosing appropriate values for the noise variance,

$$\mathbf{z}^{\star} = \arg\min_{\mathbf{z}} \quad \frac{1}{2} ||\mathbf{x} - \mathbf{d} * \mathbf{z}||_2^2 + \beta ||\mathbf{z}||_1 \tag{74}$$

B Inversion of a Rank-1 + Scaled Identity Matrix

The Woodbury matrix identity (matrix inversion lemma) states,

$$(A + UCV)^{-1} = A^{-1} - A^{-1}U(C^{-1} + VA^{-1}U)^{-1}VA^{-1} \tag{75}$$

Substituting $A = \rho I$, $C = 1$ and $U = x^T$, $V = x$, where x is a rank-1 column matrix gives,

$$\begin{aligned}
(\rho I + xx^T)^{-1} &= (\rho I)^{-1} - (\rho I)^{-1}x(1^{-1} \\
&\qquad + x^T(\rho I)^{-1}x)^{-1}x^T(\rho I)^{-1} \\
&= (\rho I)^{-1} - (\rho I)^{-1}x(1 + \frac{1}{\rho}x^Tx)^{-1}x^T(\rho I)^{-1} \\
&= (\rho I)^{-1} - (\rho I)^{-1}x(\rho + x^Tx)^{-1}x^T \\
&= \frac{1}{\rho}\left(I - \frac{1}{\rho + x^Tx}\right)xx^T \tag{76}
\end{aligned}$$