Path Orthogonal Matching Pursuit for Sparse Reconstruction and Denoising of SWIR Maritime Imagery

Timothy Doster
timothy.doster@nrl.navy.mil

Tegan Emerson
tegan.emerson@nrl.navy.mil

Colin Olson
colin.olson@nrl.navy.mil

U.S. Naval Research Laboratory
Optical Sciences Division
4555 Overlook Ave., SW, Washington, DC 20375

Abstract

We introduce an extension that may be used to augment algorithms used for the sparse decomposition of signals into a linear combination of atoms drawn from a dictionary such as those used in support of, for example, compressive sensing, k-sparse representation, and denoising. Our augmentation may be applied to any reconstruction algorithm that relies on the selection and sorting of high-correlation atoms during an analysis or identification phase by generating a "path" between the two highest-correlation atoms. Here we investigate two types of path: a linear combination (Euclidean geodesic) and a construction relying on an optimal transport map (2-Wasserstein geodesic). We test our extension by performing image denoising and k-sparse representation using atoms from a learned overcomplete kSVD dictionary. We study the application of our techniques on SWIR imagery of maritime vessels and show that our methods outperform orthogonal matching pursuit. We conclude that these methods, having shown success in our two tested problem domains, will also be useful for reducing "basis mismatch" error that arises in the recovery of compressively sampled images.

1. Introduction

Shortwave infrared (SWIR) imaging is the preferred imaging modality for daylight and low-light adverse imaging as found, for example, in the presence of haze, rain and fog. Adequate SWIR imaging performance remains hampered, however, by certain challenges. In particular, SWIR-sensing technology does not have a strong consumer demand (as opposed to military) and therefore focal plane array (FPA) design is still evolving–resulting in high costs for the densely packed arrays required for adequate long-distance imaging performance. In addition, FPA designs for SWIR sensors are more likely to yield images with low signal-to-noise ratios (SNRs). A potential solution to these problems is the use of k-sparse denoising algorithms and compressive sensing (CS) to bridge the gap in performance between SWIR and visible FPA performance.

Compressive sensing (CS) is a methodology that enables higher-resolution digital sampling of natural phenomena by leveraging good signal models to reconstruct signals that are undersampled according to classical Nyquist sampling theory [2, 10, 16, 36]. In this case, "good" models are those that can sparsely represent signals as linear combinations of relatively few atoms drawn from a dictionary. Signals that can be represented to within some acceptable error tolerance using at most k atoms are defined as k-sparse relative to that dictionary. CS theory predicts, to a level of probabilistic certainty, successful reconstruction of an undersampled signal when the underlying true signal satisfies upper limits on sparsity relative to a given dictionary [3, 8, 9, 23]. As a result, significant effort has been spent on designing dictionaries or developing algorithms that are capable of learning dictionaries that are highly representative of the expected signal class [1, 13, 21].

A persistent problem exists, however, because even if the underlying signal model could perfectly represent the signal with a single atom, the atoms must be discretely sampled from the underlying model and therefore, with high probability, will fail to represent any given signal component exactly. For example, a 1-D sinusoidal signal composed of a single tone is well-represented by a sinusoidal signal model, but if the frequency of the signal falls between the discrete

Fourier frequencies of a given Fourier basis then the number of non-zero Fourier coefficients can actually be quite large [30]. Guarantees on successful reconstruction of undersampled signals begin to fail when the assumption of sparsity is violated which can lead to the introduction of artifacts at best or complete failure to reconstruct at worst.

This "basis mismatch" problem has been considered in the literature [6, 12, 18, 35, 37] with some success in the case of 1-D sinusoidal signals where a search over the frequencies residing between the two Fourier atoms with the largest correlation can find the exact representative atom [33]. But the task is more difficult when reconstructing undersampled images because relatively few local image regions can be exactly represented by an atom from a single structured (parameterized) dictionary (e.g., Fourier, wavelet, ridgelet, etc.). Unstructured, learned dictionaries constructed from ensembles of expected images produce higher-sparsity representations but still suffer from basis mismatch. In such unstructured cases there is no parameter over which one can search to find some better atom residing between the two atoms with which the image has the largest correlation.

We seek to improve sparse reconstructions by proposing two methods for constructing better image exemplars from an underlying learned dictionary: a linear combination of the two most-correlated atoms (Euclidean geodesic) and a construction relying on the optimal transport map between said atoms (2-Wasserstein geodesic) [24]. This "path-based" augmentation can be applied to any reconstruction algorithm that relies on the selection and sorting of high-correlation atoms during an analysis or identification phase. In particular, we consider the matching pursuit (MP) family of algorithms [5, 29] which contains a number of algorithm variations predicated on the selection of high-correlation atoms.

MP algorithms are used extensively for CS recovery and k-sparse denoising so we illustrate here that augmenting MP with our path-based modification leads to sparser representations [11, 14, 15, 19]. We demonstrate these improvements by constructing sparse image representations and performing image denoising from a learned kSVD dictionary. We compare our method to a traditional MP algorithm variation and show improved results at the expense of a modest increase in computational complexity.

We begin with a summary and discussion of related algorithms to better understand the landscape of the problem in Section 2. Additionally, we define necessary terms and metrics critical to the design of the proposed algorithm. In Section 3 we present the proposed algorithm Path Orthogonal Matching Pursuit (POMP). Next, in Section 4 we describe our set of SWIR maritime imagery and construction of our learned kSVD dictionary. A description and results of the experiments (sparse representation and denoising) are presented in Section 5. Finally, we conclude with a discussion of the results as well as directions of future work in Section 6.

2. Background and Related Work

Given an image, or generic signal, questions often surround the ability of the user to reconstruct or compress the information contained in the signal. A common assumption is that the signal may be written as a linear combination of a set of reference images/signals. The reference set is called the *dictionary* and the elements of the dictionary are called *atoms*. There is extensive literature considering when linearity holds and development of tools for intelligently building overcomplete dictionaries through augmentation when the base dictionary is insufficient for sparse linear reconstructions. Dictionary design and dictionary augmentation are, however, outside the scope of this paper.

2.1. Matching Pursuit and Existing Variations

Reconstruction algorithms can terminate based on an error threshold or aim to find the best $k-$sparse representation. For a fixed degree of sparsity, consideration of all possible atom combinations of that order is computationally intractable other than for a limited set of uninteresting problems. A popular and successful work-around of this combinatorial optimization problem is a greedy algorithm called *Matching Pursuit (MP)* [5, 29]. Standard MP begins by greedily searching for the best reconstruction produced from a single atom. "Best" is determined by the magnitude of the inner product between the signal and the dictionary atoms. In other words, optimality is determined by the atom with the smallest angle between the 1-dimensional signal space and the 1-dimensional space spanned by the atom. The optimal atom is scaled by the length of the projection of the signal onto the space spanned by the optimal atom and is then subtracted from the original signal. The residual image is then fit in the same greedy way, updated, and the process repeats. k iterations of MP yields a k-sparse representation with some associated final error/residual $\mathbf{R}_k$.

Several variations of MP exist: Orthogonal MP (OMP) [7, 34, 38], Generalized OMP (GOMP) [40], Regularized OMP (ROMP) [32], Stagewise OMP (StOMP) [17], Subspace Pursuit (SP) [14], Multipath MP (MMP) [28], and Compressive Sampling MP (CoSaMP) [31]. The *identification* step in an iteration of an MP-based algorithm refers to determining which atom(s) is(are) closest to the current residual. *Augmentation* is used to describe the step of adding the atom(s) identified to the support of the reconstruction. Finally, each pursuit-type algorithm is concluded by a *residual update*. The fundamental difference between MP and OMP (as well as OMP derivatives) is that in OMP the residual is updated by projecting the image onto the orthogonal complement of the span of the current support.

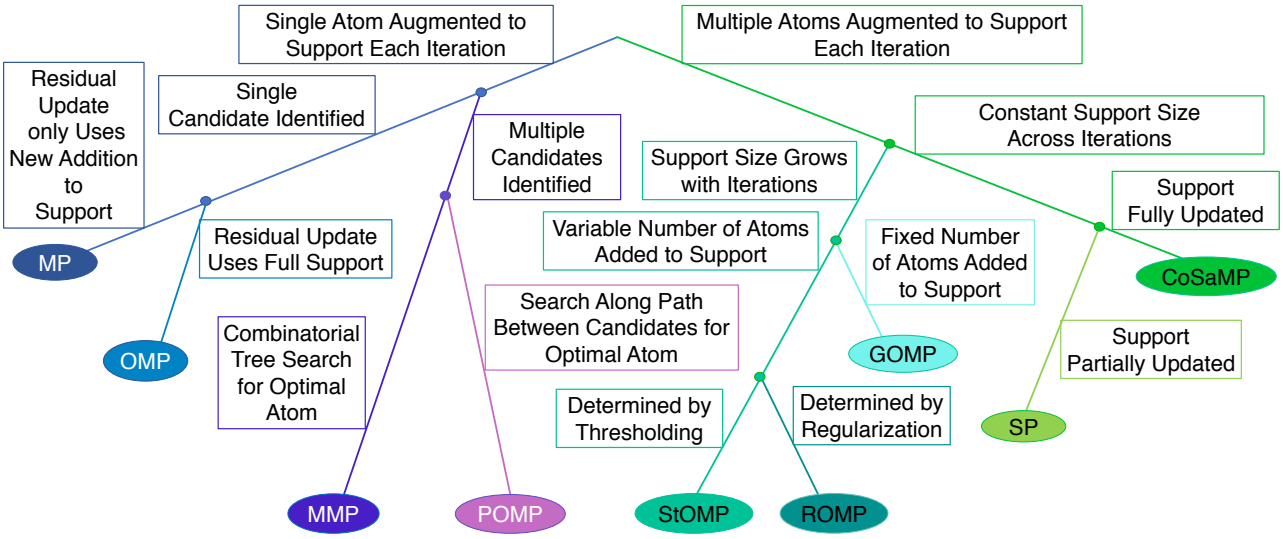

Figure 1. Tree highlighting key differences between well-known variations of the MP algorithm. Each algorithm is a leaf in the tree and the branching criterion is summarized in a color-coded box for each branch.

Note that when the dictionary consists of pairwise orthogonal atoms, MP and OMP are equivalent. For the readers' benefit, a tree describing the relationships between the MP-based algorithms listed above can be found in Figure 1.

ROMP and StOMP involve nomination of multiple candidates in the first stage of a two-stage identification step. The second stage then prunes the number of candidates based on an additional criterion (regularization or thresholding, respectively) in the identification stage. Residuals are updated in the same manner as in OMP. A key distinction between OMP and StOMP and ROMP is that multiple atoms, not just a single atom, are added to support in the augmentation steps of the latter. Based on regularization/thresholding a different number of atoms may be identified and added to the support at each iteration. Alternatively, GOMP identifies and augments the support with a fixed number of atoms at each iteration. The standard OMP algorithm is a special case of GOMP when the number of atoms identified and augmented is one at each step. In StOMP, ROMP, and GOMP the cardinality of the support grows at each iteration. CoSaMP augments the support with (potentially) multiple atoms in a given iteration but then prunes the support to the desired sparsity and updates the residual by solving a reduced-rank, least-squares problem.

Taking a more combinatorial approach, MMP searches paths through trees resulting from a fixed number of identified candidates and their children. The most senior relative of the path producing the smallest residual is then augmented to the support. GOMP and MMP aim to minimize the effect of a poor choice near the middle of the search, which has been shown to have large ramifications on reconstruction accuracy. Finally, SP views the reconstruction challenge from a more geometric perspective. An initial set of candidates, whose cardinality equals the desired level of sparsity, are identified and form the initial support. New candidates are identified based on distances between the residual and the subspace spanned by the support and traded out for existing support atoms until a stopping criterion is met.

The number of variants of the standard algorithm lends credence to the usefulness of the fundamental greedy approach. Guarantees of monotonic convergence when the signal is contained in the span of the dictionary also contribute to the popularity of the tool. Algorithms on the right branch of the tree in Figure 1 are characterized by the selection of multiple atoms during the identification and augmentation stages of the algorithm. Algorithms on the left branch only update a single atom during each step of the process.

We propose another modification to MP based on leveraging additional information contained in each greedy search by identifying the two most-representative atoms and searching along a path between the two atoms to reveal a novel single atom that better matches the current residual. From this perspective the proposed algorithm is most similar to GOMP when the number of atoms identified in each iteration is two. Given two promising atoms we consider here two different means of constructing a path between them as detailed next.

2.2. Paths Between Atoms

A *path* is a smooth map from the closest dictionary atom to the second-closest dictionary atom that is parameterized

by a variable $t \in [0,1]$. Explicitly, a path p is defined as

$$p(\mathbf{D}_1, \mathbf{D}_2, t) : \mathbf{D}_1 \rightarrow \mathbf{D}_2, \tag{1}$$

$$\text{s.t.} \quad p(\mathbf{D}_1, \mathbf{D}_2, 0) = \mathbf{D}_1 \quad \text{and} \tag{2}$$

$$p(\mathbf{D}_1, \mathbf{D}_2, 1) = \mathbf{D}_2. \tag{3}$$

One possible path is the line segment with endpoints $\mathbf{D}_1$ and $\mathbf{D}_2$ given by

$$\mathbf{D} = (1-t)\mathbf{D}_1 + (t)\mathbf{D}_2, \tag{4}$$

for $t \in [0,1]$. This path is called the Euclidean geodesic between $\mathbf{D}_1$ and $\mathbf{D}_2$, that is, the shortest path between the two points in Euclidean space.

Another possible path is that constructed from the Optimal Transport (OT) map between two images wherein one image is mapped into the other while minimizing the energy associated with the map [39]. As discussed below, a geodesic can be constructed from this map and in recent years such OT geodesics between images have yielded impressive results and insights in image registration and warping [22], super-resolution of low-resolution face images [27], and cell morphology [4]. Two main versions of OT exist: (1) the Monge OT problem in which all the intensity located at a pixel in $\mathbf{D}_1$ must be mapped to a single pixel in $\mathbf{D}_2$, and (2) the Kantorovich OT problem which allows for intensities at starting pixels to be split among multiple destination pixels [39].

Within both Monge and Kantorovich OT formulations there are additional sub-versions for different combinations of objective function and constraints. Specifically, different formulations of the OT problem produce constant-speed geodesics with respect to different distance measures. As such, a solution to the OT problem between images yields a vector field of direction vectors that implicitly indicate the terminal location (in $\mathbf{D}_2$) of intensity from a given pixel in $\mathbf{D}_1$. Let

$$\mathbf{V} = \begin{bmatrix} \vec{v}_{1,1} & \vec{v}_{1,2} & \cdots & \vec{v}_{1,m} \\ \vec{v}_{2,1} & \vec{v}_{2,2} & \cdots & \vec{v}_{2,m} \\ \vdots & \vdots & \ddots & \vdots \\ \vec{v}_{n,1} & \vec{v}_{n,2} & \cdots & \vec{v}_{n,m} \end{bmatrix} \tag{5}$$

where $\vec{v}_{j,k}$ is the velocity vector for the intensity of the pixel indexed by (j,k) in $\mathbf{D}_1$. Let $p_V(\mathbf{D}_1, \mathbf{D}_2, t)$ be the path induced by $\mathbf{V}$. We define

$$p_V(\mathbf{D}_1, \mathbf{D}_2, t) : \mathbf{D}_1 \rightarrow \mathbf{D}_2 \text{ s.t.} \tag{6}$$

$$p_V(\mathbf{D}_1, \mathbf{D}_2, t) = t\mathbf{V}(\mathbf{D}_1, \mathbf{D}_2) \tag{7}$$

where $t\mathbf{V}(\mathbf{D}_1)$ indicates movement of the intensity in $\mathbf{D}_1$ a partial step (of size t) in the directions given by OT. This OT path can be thought of as a set of linear path approximations to a globally nonlinear path between images (the 2-Wasserstein geodesic). Recently, a computationally-efficient approximation to the solution of Monge OT based on the Radon Cumulative Distribution Transform (RCDT) has been developed and has shown reduced computational time and increased performance on multiple tasks in machine learning, signal processing, and image classification [25, 26]. Due to the success of this method, we use it to estimate the path induced by the approximate solution to Monge OT.

Other viable paths between the atoms exist and will be considered in later work. At present we narrow the focus to a globally linear path (standard line segment or Euclidean geodesic) and a globally nonlinear path (2-Wasserstein geodesic) approximated by pixel-wise linear trajectories. Given a path between two dictionary atoms we search along the path for a novel atom which is closer to the test signal than either of the path end-points. If such a novel atom exists it takes the place of the single best atom in OMP. Details of the proposed algorithm are presented in the subsequent section. An example of samples along linear and OT paths between two images, as well as the angle formed between the path samples and the test image, are shown in Figure 2.

3. The Algorithm: POMP

Path Orthogonal Matching Pursuit (POMP) is our modification to the well-known OMP algorithm for signal reconstruction described in Section 2. Instead of finding a single nearest dictionary element and removing its contribution, two closest dictionary elements are chosen at each iteration and a path is formed which moves between the two atoms. We may then search for an optimal atom along this path, where optimal is defined as having the largest-magnitude, positive inner product with the test image, that is, the smallest angle between the pair of images when considered as vectors. A positive inner product can be interpreted as two images sharing more in-phase (same sign) intensities than out-of-phase intensities. When considered from the phase perspective, it is reasonable to pick the two nearest neighbors whose inner products share the same sign.

Let $\mathbf{D}_1$ be the dictionary atom having the largest-magnitude inner product with the test image. Define $\mathbf{D}_2$ to be the second-closest dictionary atom to our test image $\mathbf{T}$ whose inner product shares the same sign with the inner product of $\mathbf{D}_1$ with $\mathbf{T}$. The simplest, and perhaps most natural, form of a path is linear. In Section 2 we describe a globally linear path (Equation 4) and a pixel-wise linear path. When the globally linear path is used to choose an optimal atom within an iteration of POMP we refer to it as L-POMP. The pixel-wise linear path resulting from solving the OT problem between $\mathbf{D}_1$ and $\mathbf{D}_2$ selected within an iteration of POMP is denoted by OT-POMP. For samples along the paths between $\mathbf{D}_1$ and $\mathbf{D}_2$, the angle between $\mathbf{T}$ may be computed by

$$\theta_t = \frac{\langle \mathbf{D}_t, \mathbf{T} \rangle}{||\mathbf{D}_t||_F ||\mathbf{T}||_F}, \tag{8}$$

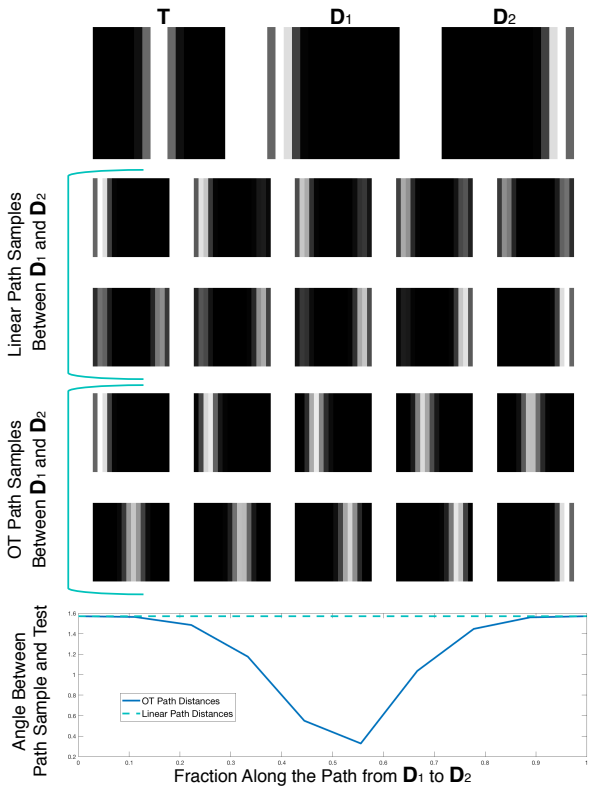

Figure 2. Examples of samples along the paths between two image atoms and the angle between the labeled test image and the path samples. Note that the images along the linear path (Euclidean geodesic) are characterized by the presence of intensity that matches the combined support from $\mathbf{D}_1$ and $\mathbf{D}_2$ and simply shifts intensity magnitude from the support represented by $\mathbf{D}_1$ to the support represented by $\mathbf{D}_2$ as t increases. As such, all samples along the Euclidean geodesic are equidistant from the test image. Conversely, the intensity support shifts between images along the OT path (2-Wasserstein geodesic) as t increases which yields an atom on the geodesic that is clearly closest to the test image.

where $t \in [0,1]$ parameterizes the distance along the path from $\mathbf{D}_1$ to $\mathbf{D}_2$ and $\|A\|_F = \sqrt{\sum_i \sum_j |a_{i,j}|^2}$ is the Frobenius norm. Let $p(\mathbf{D}_1, \mathbf{D}_2, t)$ be the path from $\mathbf{D}_1$ to $\mathbf{D}_2$. At each iteration the optimal atom is

$$\mathbf{D}_{t^*} = p(\mathbf{D}_1, \mathbf{D}_2, t^*) \quad \text{where} \quad t^* = \arg\min_{t \in [0,1]} \theta_t. \tag{9}$$

Lines 2-4 initialize variables. In line 6 we find the optimal atom and in line 7 we add its index to the list of support atoms. The same procedure is performed for the D_2 in lines 8-9. The signs of the inner products of the first and second closest (identified in line 8) atoms are identified and matched, in lines 10. The optimal atom along the path between the two nearest neighbors is selected in lines 11-12 and is then appended to the support in line 13. Residual updates and updating of indexing variables are performed in

Algorithm 1: Path Orthogonal Matching Pursuit

Input: $\mathbf{T}$, the test image and $\mathcal{D}$, the dictionary, and k the number of iterations/sparsity level.

Output: $\mathbf{X}$, the image estimate, $\mathbf{S}$, the support of the reconstruction, $\{d_k^1\}_{k=1}^K$, $\{d_k^2\}_{k=1}^K$ the vectors of first and second closest atom indices, and $\{t_k\}_{k=1}^K$, the vector containing the path parameter values.

1 **begin**
2 $\quad \mathbf{R}_1 \leftarrow \mathbf{T}$;
3 $\quad \mathbf{S} = []$;
4 $\quad k \leftarrow 1$;
5 $\quad$ **while** $k \leq K$ **do**
6 $\quad\quad \mathbf{D}_1 \leftarrow \arg\max_{\mathbf{D} \in \mathcal{D}} |\langle \mathbf{D}, \mathbf{T} \rangle|$;
7 $\quad\quad d_k^1 \leftarrow \text{index}(\mathbf{D}_1)$;
8 $\quad\quad \mathbf{D}_2 \leftarrow \arg\max_{\mathbf{D} \in \mathcal{D} \setminus \mathbf{D}_1} |\langle \mathbf{D}, \mathbf{T} \rangle|$;
9 $\quad\quad d_k^2 \leftarrow \text{index}(\mathbf{D}_2)$;
10 $\quad\quad \mathbf{D}_2 \leftarrow \text{sgn}(\langle D_1, T \rangle)\text{sgn}(\langle D_2, T \rangle)\mathbf{D}_2$;
11 $\quad\quad t_k \leftarrow \arg\min_{t \in [0,1]} \theta_t$;
12 $\quad\quad \mathbf{D}^* \leftarrow \text{path}(\mathbf{D}_1, \mathbf{D}_2, t_n)$;
13 $\quad\quad \mathbf{S} \leftarrow \text{augment}(\mathbf{S}, \mathbf{D}^*)$;
14 $\quad\quad \mathbf{P}_S \leftarrow \mathbf{S}(\mathbf{S}^\top \mathbf{S})^{-1} \mathbf{S}^\top$;
15 $\quad\quad \mathbf{X} \leftarrow \mathbf{P}_S \mathbf{T}$;
16 $\quad\quad \mathbf{R}_{k+1} \leftarrow (\mathbf{I} - \mathbf{P}_S)\mathbf{T}$;
17 $\quad\quad k \leftarrow k + 1$;
18 $\quad$ **end**
19 **end**

lines 14-17. It should be noted that standard OMP is simply lines 1-6 and 13-19 where $D^* = D_1$.

It is important to note that without further constraints, there is currently no guarantee of a novel/nontrivial minimum angle being found along the path, i.e. $\mathbf{D}^* = \mathbf{D}_t$ for $t \in (0,1)$. As long as the path is continuous for $t \in [0,1]$ a minimum will exist since the composition of continuous functions is also continuous (inner product composed with the path). Uniqueness, however, is also not guaranteed. Consider $|\langle \mathbf{T}, \mathbf{D}_t \rangle|$ for some $t \in [0,1]$. The desired inner product can be written as

$$\langle \mathbf{T}, \mathbf{D}_t \rangle = \frac{1}{2} \begin{bmatrix} \mathbf{T} \\ \mathbf{D}_t \end{bmatrix}^\top \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ \mathbf{I} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{T} \\ \mathbf{D}_t \end{bmatrix}. \tag{10}$$

This is an indefinite quadratic form, i.e. the characteristic matrix has eigenvalues $\{-1, 1\}$ with each one having the same multiplicity [20]. Resultantly, the inner product is not generally convex. However, with added constraints on the equations governing $\mathbf{D}_t$ it may be possible to prove convexity. With convexity comes uniqueness. A rigorous study of these necessary and sufficient conditions is a focus of ongoing and future work. The existence of explicit formulas

Figure 3. A subsample of the training data set used to create the learned dictionary.

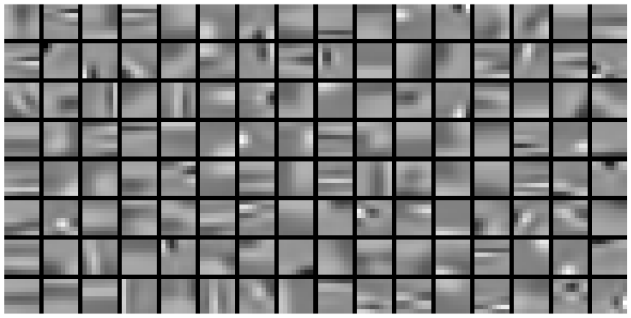

Figure 4. The kSVD dictionary.

for determining the optimal path parameter could also reduce the computational complexity of POMP. An empirical study of the computational costs of the proposed algorithm and a comparison to two existing MP-based algorithms are included in Section 5.

The proposed algorithm can be seamlessly combined with OMP as well as many of its derivatives. When the dictionary consists of pairwise orthogonal atoms, MP and OMP are equivalent. If the dictionary atoms are orthogonal then the linear combination of two of the atoms will also be orthogonal to all other atoms. For this reason when an orthogonal dictionary is used, a linear path-based MP algorithm will be equivalent to an OMP algorithm. When guarantees about orthogonality along a path cannot be made, the reconstruction resulting from an MP implementation can differ from an OMP implementation.

4. Data and Dictionary Construction

To construct a learned dictionary one must identify a collection of a training images that roughly represent the expected statistical qualities of the imagery under test. We choose a set of 85 broadband SWIR images of marine vessels including Naval (military), fishing, cargo, and sailing. Each 16-bit image is of size 1024x1280–a subset of which can be seen in Figure 3. This corpus of images was divided randomly such that 80 images were designated to train the

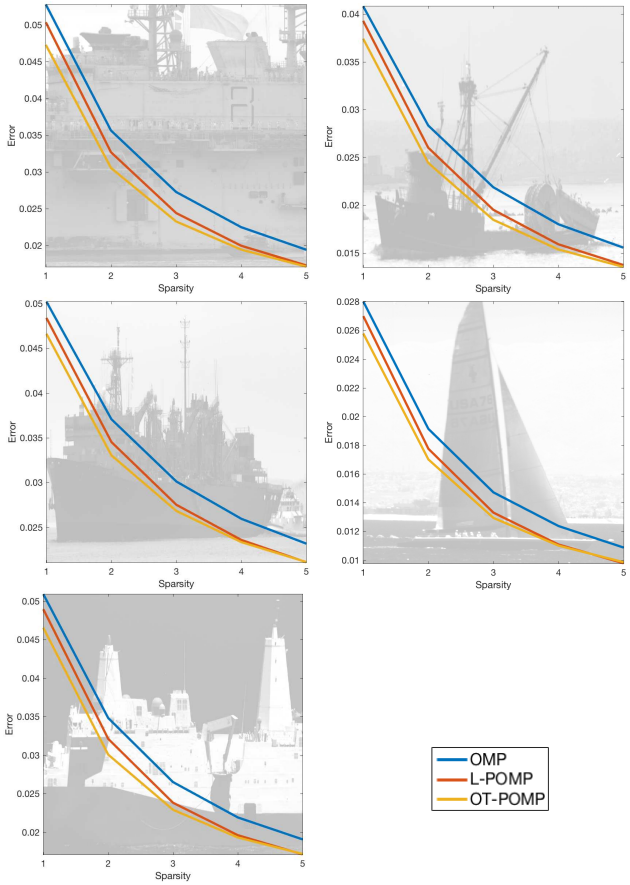

Figure 5. Relative reconstruction error as a function of iterations for the three algorithms considered. Results over 5 iterations are shown. Background of each plot is the test image under consideration (note that the alpha value has been lowered for graph readability). Images are numbered 1-5 left-to-right top-to-bottom to match with the numbering in Table 1.

dictionary and 5 images were withheld for testing. We randomly selected 5,000 patches of size 8×8 from each training image to generate the kSVD dictionary. Over 100 iterations of the kSVD algorithm 256 atoms were selected to minimize the representation error using OMP and a target sparsity of 6 atoms to yield the learned dictionary shown in Figure 4.

5. Experiments and Results

Computational results are shown for both k-sparse signal reconstruction (see Section 5.1) and image denoising (see Section 5.2 using a fixed number of OMP iterations). Both applications are tested on the five test images shown as backgrounds in Figure 5. For both algorithms we compare the current residual to five discrete samples along the path to check for a closer atom. This implementation is suboptimal but provides sufficient proof of concept.

Image	Original	OMP	L-POMP	OT-POMP
1	28.378	33.831	34.239	**34.786**
2	28.352	31.996	32.205	**33.232**
3	28.369	30.541	30.767	**31.525**
4	28.374	33.246	33.397	**34.645**
5	28.364	32.314	32.506	**33.629**

Table 1. Starting PSNR values for five images and denoised PSNR values for three methods. See Figure 5 for a representation of the imagery.

5.1. $k-$Sparse Reconstructions

We begin by considering the reconstruction error between a pristine original image and a patch-based k-sparse reconstruction using the dictionary. The dictionary is used to reconstruct 8x8 patches overlapped with a stride of one where the error between the original and the reconstructed image is given by $||\mathbf{R}_k||_F/||\mathbf{T}||_F$, where k is the number of atoms included in the reconstruction. Reconstruction errors as a function of k for all algorithms and images are shown in Figure 5; the background of each plot is the image being sparsely reconstructed.

Figure 5 shows that the path-augmented approach yields lower reconstruction error for all tested sparsity levels and images. In each case the reconstructions obtained using the proposed heuristic show L-POMP and OT-POMP performing better than OMP. It is also clear that OT-POMP outperforms L-POMP over the first few iterations (sparsity level) but after approximately 5 the two versions converge in performance. We have noticed–though the results are omitted here–that both L-POMP and OT-OMP continue to outperform OMP for decreasing sparsity levels but L-POMP begins to overtake OT-POMP in performance. This is believed to occur due to additional possibility of errors introduced during the OT geodesic construction which are not a problem in the linear case.

5.2. Denoising

Denoising experiments were also performed where additive white Gaussian noise (AWGN) with $\sigma = 2,500$ (images are 16-bit) was added to each pristine image. Patches of the noisy image, of size 8×8, are then estimated using several iterations of the indicated algorithm. A denoised image is then constructed by stitching together the OMP estimated patches. Performance is measured using output PSNR as is standard practice. Results of denoising are shown in Table 1 for a sparsity of 1. Path augmentation of OMP improves PSNR for each image after patch fitting for a fixed number of OMP iterations. Average improvement in PSNR using L-POMP is $\sim$ 0.2dB and using OT-POMP is $\sim$ 1.1dB.

6. Discussion

In this paper we have shown for sparse image representation and image denoising using a patch-based learned dictionary that the proposed POMP algorithm outperforms the baseline OMP algorithm. In this preliminary step we have shown that this algorithm is capable of addressing the basis mismatch problem seen in compressive sensing which is of particular concern when dealing with SWIR imagers.

Application of POMP to data sets expressing different signal-to-noise ratios will allow for a more comprehensive comparison to existing MP-based algorithms in the settings where they have been shown to excel. Furthermore, we intend to perform a detailed computational complexity analysis of POMP so that the gains of implementation may be appropriately qualified.

In the future we aim to develop a characterization of the tasks for which novel minimums along the path exist. If guarantees of such minimums exist it would further reduce computational expense and alleviate the need to compare test images to a discrete sampling of the paths between nearest atoms. Additional consideration of the necessary and/or sufficient conditions under which there exist closed-form solutions for determination of both the optimal path parameter value and its corresponding atom are underway. If such a characterization proves too difficult we proposed using a search-based methodology (e.g., binary search or golden-section search) to find a better approximation of the optimal atom contained along the path in the POMP algorithm.

Results presented in this work inspire confidence that there are further improvements to be discovered through the integration of POMP with dictionary augmentation approaches.

References

[1] M. Aharon, M. Elad, and A. Bruckstein. rmk-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Trans. Signal Process.*, 54(11):4311–4322, 2006. 1

[2] R. G. Baraniuk. Compressive sensing. *IEEE Signal Process. Mag.*, 24(4):118–121, 2007. 1

[3] R. G. Baraniuk, M. Davenport, R. A. DeVore, and M. B. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approx.*, 28(3):253–263, 2008. 1

[4] S. Basu, S. Kolouri, and G. K. Rohde. Detecting and visualizing cell phenotype differences from microscopy images using transport-based morphometry. *Proc. Natl. Acad. Sci.*, 111(9):3448–3453, 2014. 4

[5] F. Bergeaud and S. Mallat. Matching pursuit of images. In *Image Process., 1995. Proceedings., International Conference on*, volume 1, pages 53–56. IEEE, 1995. 2

[6] P. Boufounos, V. Cevher, A. C. Gilbert, Y. Li, and M. J. Strauss. What's the frequency, Kenneth?: Sublinear Fourier

sampling off the grid. *Lecture Notes in Computer Science*, 7408:61–72, 2012. 2

[7] T. T. Cai and L. Wang. Orthogonal matching pursuit for sparse signal recovery with noise. *IEEE Trans. Inf. theory*, 57(7):4680–4688, 2011. 2

[8] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency inf. *IEEE Trans. Inf. Theory*, 52(2):489–509, 2006. 1

[9] E. J. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Commun. Pure Appl. Math.*, 59(8):1207–1223, 2006. 1

[10] E. J. Candès and T. Tao. Decoding by linear programming. *IEEE Trans. Inf. Theory*, 51(12):4203–4215, 2005. 1

[11] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM Review*, 43(1):129–159, 2001. 2

[12] Y. Chi, L. Scharf, A. Pezeshki, and A. R. Calderbank. Sensitivity to basis mismatch in compressed sensing. *IEEE Trans. Signal Process.*, 59(5):2182–2195, 2011. 2

[13] R. R. Coifman and M. V. Wickerhauser. Entropy-based algorithms for best basis selection. *IEEE Trans. Inf. Theory*, 38(2):713–718, 1992. 1

[14] W. Dai and O. Milenkovic. Subspace pursuit for compressive sensing signal reconstruction. *IEEE Trans. Inf. Theory*, 55(5):2230–2249, 2009. 2

[15] T. T. Do, L. Gan, N. Nguyen, and T. D. Tran. Sparsity adaptive matching pursuit algorithm for practical compressed sensing. In *Signals, Systems and Computers, 2008 42nd Asilomar Conference on*, pages 581–587. IEEE, 2008. 2

[16] D. L. Donoho. Compressed sensing. *IEEE Trans. Inf. Theory*, 52(4):1289–1306, 2006. 1

[17] D. L. Donoho, Y. Tsaig, I. Drori, and J.-L. Starck. Sparse solution of underdetermined systems of linear equations by stagewise orthogonal matching pursuit. *IEEE Trans. Inf. Theory*, 58(2):1094–1121, 2012. 2

[18] C. Ekanadham, D. Tranchina, and E. P. Simoncelli. Recovery of sparse translation-invariant signals with continuous basis pursuit. *IEEE Trans. Signal Process.*, 59(10):4735–4744, 2011. 2

[19] M. Elad and M. Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Trans. Image Process.*, 15(12):3736–3745, 2006. 2

[20] G. H. Golub and C. F. Van Loan. *Matrix Computations*, volume 3. JHU Press, 2012. 5

[21] R. Gribonval and M. Nielsen. Sparse representations in unions of bases. *IEEE Trans. Inf. Theory*, 49(12):3320–3325, 2003. 1

[22] S. Haker, L. Zhu, A. Tannenbaum, and S. Angenent. Optimal mass transport for registration and warping. *International Journal of Computer Vision*, 60(3):225–240, 2004. 4

[23] J. Haupt and R. Nowak. Signal reconstruction from noisy random projections. *IEEE Trans. Inf. Theory*, 52(9):4036–4048, 2006. 1

[24] S. Kolouri, S. Park, M. Thorpe, D. Slepčev, and G. K. Rohde. Transport-based analysis, modeling, and learning from signal and data distributions. *preprint*, arXiv:1609.04767, 2017. 2

[25] S. Kolouri, S. R. Park, and G. K. Rohde. The Radon Cumulative Distribution Transform and its application to image classification. *IEEE Trans. Image Process.*, 25(2):920–934, 2016. 4

[26] S. Kolouri, S. R. Park, M. Thorpe, D. Slepčev, and G. K. Rohde. Optimal mass transport: Signal processing and machine-learning applications. *IEEE Signal Process. Mag.*, 34(4):43–59, 2017. 4

[27] S. Kolouri and G. K. Rohde. Transport-based single frame super resolution of very low resolution face images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4876–4884, 2015. 4

[28] S. Kwon, J. Wang, and B. Shim. Multipath matching pursuit. *IEEE Trans. Inf. Theory*, 60(5):2986–3001, 2014. 2

[29] S. G. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Trans. Signal Process.*, 41(12):3397–3415, 1993. 2

[30] C. V. McLaughlin, J. M. Nichols, and F. Bucholtz. Basis mismatch in a compressively sampled photonic link. *IEEE Photonics Technology Letters*, 25(23):2297–2300, 2013. 2

[31] D. Needell and J. A. Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. *Applied and Computational Harmonic Analysis*, 26(3):301–321, 2009. 2

[32] D. Needell and R. Vershynin. Uniform uncertainty principle and signal recovery via regularized orthogonal matching pursuit. *Foundations of Computational Mathematics*, 9(3):317–334, 2009. 2

[33] J. M. Nichols, A. K. Oh, and R. M. Willett. Reducing basis mismatch in harmonic signal recovery via alternating convex search. *IEEE Signal Process. Letters*, 21(8):1007–1011, 2014. 2

[34] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad. Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition. In *Proceedings of the 27th Asilomar Conference on Signals, Systems and Computers*, pages 40–44. IEEE, 1993. 2

[35] N. Rao, P. Shah, S. Wright, and R. Nowak. A greedy forward-backward algorithm for atomic norm constrained minimization. In *International Conference on Acoustics, Speech and Signal Process. (ICASSP)*, pages 5885–5889. IEEE, 2013. 2

[36] J. Romberg. Imaging via compressive sampling. *IEEE Signal Process. Mag.*, 25(2):14–20, 2008. 1

[37] G. Tang, B. N. Bhaskar, P. Shah, and B. Recht. Compressive sensing off the grid. In *50th Annual Allerton Conference on Communication, Control, and Computing*, pages 778–785. IEEE, 2012. 2

[38] J. A. Tropp and A. C. Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Trans. Inf. Theory*, 53(12):4655–4666, 2007. 2

[39] C. Villani. *Optimal Transport: Old and New*, volume 338. Springer Science & Business Media, 2008. 4

[40] J. Wang, S. Kwon, and B. Shim. Generalized orthogonal matching pursuit. *IEEE Trans. Signal Process.*, 60(12):6202–6216, 2012. 2