

ONLINE CONVOLUTIONAL DICTIONARY LEARNING FOR MULTIMODAL IMAGING

Kévin Degraux^{}, Ulugbek S. Kamilov[†], Petros T. Boufounos[†], and Dehong Liu[†]*

^{*} ISPGROUP/ICTEAM, FNRS, Université catholique de Louvain, 1348, Louvain-la-Neuve, Belgium
email: kevin.degraux@uclouvain.be

[†] Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA 02139, USA
email: {kamilov, petrosb, liudh}@merl.com

ABSTRACT

Computational imaging methods that can exploit multiple modalities have the potential to enhance the capabilities of traditional sensing systems. In this paper, we propose a new method that reconstructs multimodal images from their linear measurements by exploiting redundancies across different modalities. Our method combines a convolutional group-sparse representation of images with total variation (TV) regularization for high-quality multimodal imaging. We develop an online algorithm that enables the unsupervised learning of convolutional dictionaries on large-scale datasets that are typical in such applications. We illustrate the benefit of our approach in the context of joint intensity-depth imaging.

Index Terms— multimodal imaging, convolutional dictionary learning, online learning, sparse regularization

1. INTRODUCTION

Multimodal imaging systems acquire several measurements of an object using multiple distinct sensing modalities. Often, the data acquired from the sensors is jointly processed to improve the imaging quality in one or more of the acquired modalities. Such imaging methods have the potential to enable new capabilities in traditional sensing systems, providing complementary sources of information about the object. Some of the most common applications of multimodal imaging include remote sensing [1], biomedical imaging [2], and high-resolution depth sensing [3].

We consider a joint imaging inverse problem with multiple noisy linear measurements

$$\mathbf{y}_\ell = \Phi_\ell \mathbf{x}_\ell + \mathbf{e}_\ell, \quad (1)$$

where for each modality $\ell \in [1, \dots, L]$, $\mathbf{y}_\ell \in \mathbb{R}^{M_\ell}$ denotes the corresponding measurement vector, $\mathbf{x}_\ell \in \mathbb{R}^N$ denotes the unknown image, $\Phi_\ell \in \mathbb{R}^{M_\ell \times N}$ denotes the sensing matrix, and $\mathbf{e}_\ell \in \mathbb{R}^{M_\ell}$ denotes the noise in the measurements. The images $\{\mathbf{x}_\ell\}_{\ell \in [1 \dots L]}$ correspond to the same physical object viewed from different modalities. For example, each \mathbf{x}_ℓ may represent a different color channel, spectral band, or a type of sensor. For simplicity, we assume that the desired dimension of the images is the same across all modalities and that acquisition devices are perfectly registered. The key insight used in our paper is that information about a single modality exists, in some form, in other modalities. This information can be exploited to improve the quality of multimodal imaging, as long as it can be extracted from the measurements.

This work was completed while K. Degraux was with MERL.

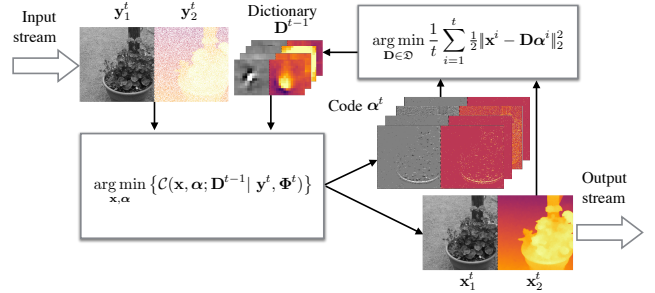


Fig. 1. Illustration of the proposed multimodal imaging method.

1.1. Main Contributions

In this work, we propose a novel approach based on jointly sparse representation of multimodal images. Specifically, we are interested in learning data-adaptive convolutional dictionaries for both reconstructing and representing the signals given their linear measurements. The main benefit of a convolutional approach is that it is translation invariant and leads to a sparse representation over the entire image. This, however, comes with the increase in the computational cost, which we address by developing a new online convolutional dictionary learning method suitable for working with large-scale datasets. Our key contributions are summarized as follows:

- We provide a new formulation for multimodal computational imaging, incorporating a convolutional joint sparsity prior and a total variation (TV) regularizer. In this formulation, the high resolution images are determined by solving an optimization problem, where the regularizer exploits the redundancies across different modalities.
- We develop an online convolutional dictionary learning algorithm, illustrated in Figure 1. By accommodating an additional TV regularizer in the cost, the algorithm is able to learn the convolutional dictionary in an unsupervised fashion, directly from the noisy measurements. We validate our approach for joint intensity-depth imaging.

1.2. Related Work

Starting from early work by Olshausen and Field [4, 5], dictionary learning has become a standard tool for various tasks in image processing [6–10]. Our approach builds upon two prior lines of research, one on convolutional sparse representations [11–13] and one on online dictionary learning [14–16]. Since our method relies on TV regularization, it is also related to TV-based imaging algorithms [17–20]. Specifically, our method is based on the popular

fast iterative shrinkage/thresholding algorithm (FISTA) for reconstructing images from measurements. Our method is validated on the problem of joint intensity–depth imaging, also considered in [21–26]. In particular, [23, 24] use traditional sparse coding for combining depth and intensity, and [27] uses convolutional dictionaries for representing multiple modalities. Our approach extends earlier work by performing multi-modal image reconstruction with convolutional dictionaries and developing a dedicated online learning algorithm for large-scale settings.

2. PROPOSED METHOD

2.1. Problem Formulation

The underlying assumption in our approach is that a jointly sparse convolutional model can accurately approximate the images $\{\mathbf{x}_\ell\}$ as

$$\mathbf{x}_\ell \approx \mathbf{D}_\ell \boldsymbol{\alpha}_\ell \triangleq \sum_{k=1}^K \mathbf{d}_{\ell k} * \boldsymbol{\alpha}_{\ell k}, \quad (2)$$

where $\{\mathbf{d}_{\ell k}\}$ is the set of LK convolutional filters in \mathbb{R}^P , $*$ denotes convolution, and $\{\boldsymbol{\alpha}_{\ell k}\}$ is the set of coefficient maps in $\mathbb{R}^{\tilde{N}}$. Note that \mathbf{D}_ℓ and $\boldsymbol{\alpha}_\ell$ denote the concatenation of all K dictionaries and coefficient maps, respectively. Given the complete dictionary $\mathbf{D} = (\mathbf{D}_1, \dots, \mathbf{D}_L)$, we can define our imaging problem as the following joint optimization

$$(\hat{\mathbf{x}}, \hat{\boldsymbol{\alpha}}) = \arg \min_{\mathbf{x}, \boldsymbol{\alpha}} \{\mathcal{C}(\mathbf{x}, \boldsymbol{\alpha}; \mathbf{D} | \mathbf{y}, \boldsymbol{\Phi})\}, \quad (3)$$

where the cost function \mathcal{C} is given by

$$\mathcal{C}(\mathbf{x}, \boldsymbol{\alpha}; \mathbf{D} | \mathbf{y}, \boldsymbol{\Phi}) \triangleq \frac{1}{2} \|\mathbf{y} - \boldsymbol{\Phi} \mathbf{x}\|_2^2 + \frac{\rho}{2} \|\mathbf{x} - \mathbf{D} \boldsymbol{\alpha}\|_2^2 + \lambda \|\boldsymbol{\alpha}\|_{2,1} + \tau \mathcal{R}(\mathbf{x}), \quad (4)$$

with $\mathbf{y} \triangleq \text{vc}(\mathbf{y}_1, \dots, \mathbf{y}_L)$, $\mathbf{x} \triangleq \text{vc}(\mathbf{x}_1, \dots, \mathbf{x}_L)$, and $\boldsymbol{\alpha} \triangleq \text{vc}(\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_L)$, denoting the vertical concatenation (vc) of corresponding signals and $\boldsymbol{\Phi} \triangleq \text{diag}(\boldsymbol{\Phi}_1, \dots, \boldsymbol{\Phi}_L)$ denoting the block diagonal sensing matrix. The first quadratic term in (4) measures the data-fidelity, while the second controls the approximation quality of the dictionaries. The first regularization term

$$\|\boldsymbol{\alpha}\|_{2,1} \triangleq \sum_{k=1}^K \sum_{n=1}^{\tilde{N}} \|\boldsymbol{\alpha}_{\cdot kn}\|_2 \quad (5)$$

imposes group- or joint-sparsity of coefficients across L modalities. Here, $\boldsymbol{\alpha}_{\cdot kn} \in \mathbb{R}^L$ denotes the vector formed by the aligned entries of the coefficient maps associated with kernel k for every modality ℓ . Specifically, this regularizer promotes the co-occurrence of image features, encoded by the dictionary \mathbf{D} , in all the modalities. The second regularizer in (4) corresponds to the isotropic TV penalty [17]

$$\mathcal{R}(\mathbf{x}) \triangleq \sum_{\ell=1}^L \sum_{n=1}^N \|[\mathbf{L} \mathbf{x}_\ell]_n\|_2, \quad (6)$$

where \mathbf{L} denotes the discrete gradient operator. Unsupervised learning of dictionaries from \mathbf{y} is complicated when the imaging problem is ill-posed. The goal of including the TV regularizer is to assist this learning. In practice, we observed significant improvement in quality when TV was included, both during learning and reconstruction. Finally, the positive constants ρ , λ , and τ are parameters controlling

the tradeoff between the data fidelity and regularization.

The joint optimization program in (3) is a convex problem. To solve it, we use the monotonic variant of FISTA [18]. In particular, we split $\mathcal{C}(\mathbf{x}, \boldsymbol{\alpha}; \mathbf{D} | \boldsymbol{\Phi}, \mathbf{y})$ into a smooth quadratic term

$$\frac{1}{2} \|\mathbf{y} - \boldsymbol{\Phi} \mathbf{x}\|_2^2 + \frac{\rho}{2} \|\mathbf{x} - \mathbf{D} \boldsymbol{\alpha}\|_2^2 \quad (7)$$

and a non-smooth term that is separable in \mathbf{x} and $\boldsymbol{\alpha}$

$$\lambda \|\boldsymbol{\alpha}\|_{2,1} + \tau \mathcal{R}(\mathbf{x}), \quad (8)$$

The proximal operator associated with $\lambda \|\boldsymbol{\alpha}\|_{2,1}$ is equal to

$$[\text{prox}_{\lambda \|\cdot\|_{2,1}}(\boldsymbol{\alpha})]_{\cdot kn} = \left(\|\boldsymbol{\alpha}_{\cdot kn}\|_2 - \lambda \right)_+ \frac{\boldsymbol{\alpha}_{\cdot kn}}{\|\boldsymbol{\alpha}_{\cdot kn}\|_2}, \quad (9)$$

where the operator $(\cdot)_+$ extracts the positive part of its argument. While the proximal of TV does not have a closed-form solution, it can be efficiently implemented [18].

2.2. Learning Algorithm

Suppose the input data is streamed so that at every time step $t \in \mathbb{N}$ we get a pair $(\mathbf{y}^t, \boldsymbol{\Phi}^t)$. The learning procedure attempts to minimize (4) for all t , jointly for \mathbf{x} , $\boldsymbol{\alpha}$ and \mathbf{D} . Specifically, let $\mathcal{J}^t(\mathbf{D}) \triangleq \min_{\mathbf{x}, \boldsymbol{\alpha}} \{\mathcal{C}(\mathbf{x}, \boldsymbol{\alpha}; \mathbf{D} | \mathbf{y}^t, \boldsymbol{\Phi}^t)\}$, then this amounts to solving

$$\min_{\mathbf{D} \in \mathfrak{D}} \{\mathbb{E}[\mathcal{J}^t(\mathbf{D})]\}, \quad (10)$$

with respect to \mathbf{D} , where the expectation is taken over t . Note that, to compensate for scaling ambiguities, we restrict the optimization of \mathbf{D} to a closed convex set \mathfrak{D} . Specifically, \mathfrak{D} is the set of convolutional dictionaries that have kernels in the ℓ_2 ball, i.e., $\|\mathbf{d}_{\ell k}\|_2 \leq 1$.

The joint optimization program in (10) is difficult to solve directly. Thus, we use an alternating minimization procedure. In particular, at iteration t , given the current dictionary \mathbf{D}^{t-1} , and a new pair of data $(\mathbf{y}^t, \boldsymbol{\Phi}^t)$, we first solve

$$(\mathbf{x}^t, \boldsymbol{\alpha}^t) \leftarrow \arg \min_{\mathbf{x}, \boldsymbol{\alpha}} \{\mathcal{C}(\mathbf{x}, \boldsymbol{\alpha}; \mathbf{D}^{t-1} | \mathbf{y}^t, \boldsymbol{\Phi}^t)\}, \quad (11)$$

using the method presented in Section 2.1. Then, following the principle of [14], we use all the previous iterates and chose \mathbf{D} to minimize a surrogate of $\mathbb{E}[\mathcal{J}^t(\mathbf{D})]$ given by

$$\frac{1}{t} \sum_{i=1}^t \mathcal{C}(\mathbf{x}^i, \boldsymbol{\alpha}^i; \mathbf{D} | \mathbf{y}^i, \boldsymbol{\Phi}^i). \quad (12)$$

This second step, performed using a block gradient descent on the kernels $\{\mathbf{d}_{\ell k}\}$, is described in Section 2.3. The complete learning algorithm is summarized in Algorithm 1.

2.3. Dictionary update

Keeping \mathbf{x}^i and $\boldsymbol{\alpha}^i$ fixed, the only term in \mathcal{C} that depends on \mathbf{D} is the quadratic coupling penalty $\frac{\rho}{2} \|\mathbf{x}^i - \mathbf{D} \boldsymbol{\alpha}^i\|_2^2$. Therefore, we can equivalently minimize $\frac{1}{2t} \sum_{i=1}^t \|\mathbf{x}_\ell^i - \mathbf{D}_\ell \boldsymbol{\alpha}_\ell^i\|_2^2$, for each modality ℓ . Since everything is separable in ℓ , in the remainder we drop the subscript for notational clarity. Note that, since the convolution operation is commutative and the $\boldsymbol{\alpha}^i$ are fixed, we can rewrite

$$\mathbf{D} \boldsymbol{\alpha}^i = \sum_{k=1}^K \mathbf{d}_k * \boldsymbol{\alpha}_k^i = \sum_{k=1}^K \boldsymbol{\alpha}_k^i * \mathbf{d}_k = \mathbf{A}^i \mathbf{d}, \quad (13)$$

Algorithm 1 Online Convolutional Dictionary Learning

```

1: procedure ONLINECDL
2: Input: Stream of data  $t \mapsto (\mathbf{y}^t, \Phi^t)$ , initial dictionary  $\mathbf{D}^0$ .
3:    $\mathbf{C}^0 \leftarrow \mathbf{0}$ ;    $\mathbf{b}^0 \leftarrow \mathbf{0}$ ;
4:   while streaming data, do
5:     Draw a pair  $(\mathbf{y}^t, \Phi^t)$ ;
6:     Sparse coding step:
7:      $(\mathbf{x}^t, \alpha^t) \leftarrow \arg \min_{\mathbf{x}, \alpha} \{\mathcal{C}(\mathbf{x}, \alpha; \mathbf{D}^{t-1} | \mathbf{y}^t, \Phi^t)\}$ ;
8:     Update memory:
9:      $\mathbf{b}^t \leftarrow (1 - \frac{1}{t})\mathbf{b}^{t-1} + \frac{1}{t} \sum_{i=1}^t \mathbf{A}^{iT} \mathbf{x}^i$ ;
10:     $\mathbf{C}^t \leftarrow (1 - \frac{1}{t})\mathbf{C}^{t-1} + \frac{1}{t} \sum_{i=1}^t \mathbf{A}^{iT} \mathbf{A}^i$ ;
11:    Dictionary update (14), (15) initialized with  $\mathbf{D}^{t-1}$ :
12:     $\mathbf{D}^t \leftarrow \arg \min_{\mathbf{D} \in \mathcal{D}} \frac{1}{2t} \sum_{i=1}^t \|\mathbf{x}^i - \mathbf{D}\alpha^i\|_2^2$ ;

```

where $\mathbf{A}^i \triangleq (\mathbf{A}_1^i, \dots, \mathbf{A}_K^i) \in \mathbb{R}^{N \times KP}$ is the sum-of-convolutions linear operator and $\mathbf{d} \triangleq \text{vc}(\mathbf{d}_1, \dots, \mathbf{d}_K)$. In order to minimize $\mathcal{G}^t(\mathbf{d}) \triangleq \frac{1}{2t} \sum_{i=1}^t \|\mathbf{x}^i - \mathbf{A}^i \mathbf{d}\|_2^2$, subject to $\|\mathbf{d}_k\|_2 \leq 1$, as in [14], we apply a projected block-coordinate gradient descent. The algorithm starts for $s = 0$ with $\mathbf{d}^{t,0} \leftarrow \mathbf{D}^{t-1}$ and iteratively applies the following two steps for all $k \in [1, \dots, K]$,

$$\tilde{\mathbf{d}}_k^{t,s} \leftarrow \mathbf{d}_k^{t,s-1} - \frac{1}{L_k^t} \nabla_{\mathbf{d}_k} \mathcal{G}^t \left(\text{vc}(\dots, \tilde{\mathbf{d}}_{k-1}^{t,s}, \mathbf{d}_k^{t,s-1}, \dots) \right) \quad (14)$$

$$\mathbf{d}_k^{t,s} \leftarrow \tilde{\mathbf{d}}_k^{t,s} / \max\{\|\tilde{\mathbf{d}}_k^{t,s}\|, 1\} \quad (15)$$

until convergence or until a maximum number of iterations is reached. Note that $\nabla_{\mathbf{d}_k}$ denotes the partial gradient

$$\nabla_{\mathbf{d}_k} \mathcal{G}^t(\mathbf{d}) = \frac{1}{t} \sum_{i=1}^t \mathbf{A}_k^{iT} (\mathbf{A}^i \mathbf{d} - \mathbf{x}^i), \quad (16)$$

and L_k^t is the Lipschitz constant of $\nabla_{\mathbf{d}_k} \mathcal{G}^t(\mathbf{d})$. Importantly, we can take advantage of all the previous iterates to compute this gradient. Indeed, we can write it as

$$\nabla_{\mathbf{d}_k} \mathcal{G}^t(\mathbf{d}) = \mathbf{C}_k^t \mathbf{d} - \mathbf{b}_k^t, \quad (17)$$

where the *memory vector* and the symmetric *memory matrix*,

$$\mathbf{b}^t = \text{vc}(\mathbf{b}_1^t, \dots, \mathbf{b}_K^t) \triangleq \frac{1}{t} \sum_{i=1}^t \mathbf{A}^{iT} \mathbf{x}^i, \quad (18)$$

$$\mathbf{C}^t = \text{vc}(\mathbf{C}_1^t, \dots, \mathbf{C}_K^t) \triangleq \frac{1}{t} \sum_{i=1}^t \mathbf{A}^{iT} \mathbf{A}^i, \quad (19)$$

with $\mathbf{C}_k^t \triangleq (\mathbf{C}_{k,1}^t, \dots, \mathbf{C}_{k,K}^t) = (\mathbf{A}_k^{iT} \mathbf{A}_1^i, \dots, \mathbf{A}_k^{iT} \mathbf{A}_K^i)$, are computed recursively from the previous iterates as

$$\mathbf{b}^t \leftarrow \frac{t-1}{t} \mathbf{b}^{t-1} + \frac{1}{t} \mathbf{A}^{tT} \mathbf{x}^t, \quad (20)$$

$$\mathbf{C}^t \leftarrow \frac{t-1}{t} \mathbf{C}^{t-1} + \frac{1}{t} \mathbf{A}^{tT} \mathbf{A}^t. \quad (21)$$

Note that the aforementioned Lipschitz constant is $L_k^t = \|\mathbf{C}_{kk}^t\|_2$.

2.4. Implementation details

Convolutional implementation. A naive implementation of the online convolutional learning algorithm would require to store \mathbf{C}^t , which is a dense symmetric $KP \times KP$ matrix. By definition, $\mathbf{A}_k^i \in \mathbb{R}^{N \times P}$ is a convolution operator whose columns and rows are restricted to match the size of its input \mathbf{d}_k and output \mathbf{x} , respectively. Similarly, $\mathbf{A}_k^{iT} \in \mathbb{R}^{P \times N}$ is a restricted convolution operator whose kernel $\check{\alpha}_k^i$ is the flipped version of α_k^i . Therefore, if the size of α_k^i is chosen such that it implements a full convolution, then the operator $\mathbf{A}_k^{iT} \mathbf{A}_{k'}^i$ corresponds to a restricted convolution with a kernel $\check{\alpha}_k^i * \alpha_{k'}^i$. Importantly, the restrictions to a $P \times P$ operator imply that only part of the kernel, of size proportional to P , is actually used. We denote this effective part by $\mathbf{S}_P(\check{\alpha}_k^i * \alpha_{k'}^i)$ where \mathbf{S}_P is the selection operator. This implies that $\mathbf{C}_{kk'}^t$ is itself a convolution operator with kernel

$$\mathbf{c}_{kk'}^t \triangleq \frac{1}{t} \sum_{i=1}^t \mathbf{S}_P(\check{\alpha}_k^i * \alpha_{k'}^i). \quad (22)$$

Therefore, an efficient implementation of \mathbf{C}^t only requires to store or convolve with those K^2 kernels. Note that this computational trick is a big argument in favor of decoupling $\frac{1}{2} \|\mathbf{y} - \Phi \mathbf{x}\|_2^2$ from $\frac{1}{2} \|\mathbf{x} - \mathbf{D} \alpha\|_2^2$. By contrast, using the same technique to minimize $\frac{1}{2} \|\mathbf{y} - \Phi \mathbf{D} \alpha\|_2^2$ instead, requires the explicit storage of a dense symmetric $KP \times KP$ matrix. When Φ is a mask and the dictionary is not convolutional, Mensch *et al.* [16] adopt a different strategy which consists in approximating the surrogate function.

Data centering. Patch-based dictionary learning is known to be more effective when the input data is centered to have zero mean. While this is not required, it is common to pre-process the data by removing the means of the training patches [28]. Accordingly, we first estimate a local mean, *i.e.*, a low-pass component, \mathbf{x}^{lo} of the data \mathbf{x} . The remaining component is the high-pass image $\mathbf{x}^{\text{hi}} \triangleq \mathbf{x} - \mathbf{x}^{\text{lo}}$. We aim to learn the sparse synthesis model $\mathbf{x}^{\text{hi}} \approx \mathbf{D} \alpha$. To do so, we adapt the method above, replacing the coupling term in (4) by $\frac{\rho}{2} \|\mathbf{x} - \mathbf{x}^{\text{lo}} - \mathbf{D} \alpha\|_2^2$. Specifically, we use a mask-aware low-pass filter to estimate \mathbf{x}^{lo} from (\mathbf{y}, Φ) where Φ is a mask operator. Let $\mathcal{L}(\cdot)$ be a low-pass filter and $\varphi \in \{0, 1\}^N$ the mask (*i.e.*, the diagonal of Φ). Then, for every pixel $[\mathcal{L}(\varphi)]_n > 0$, we compute $[\mathbf{x}^{\text{lo}}]_n = \frac{[\mathcal{L}(\mathbf{y})]_n}{[\mathcal{L}(\varphi)]_n}$. We use the nearest neighbor interpolation to fill the remaining pixels where $[\mathcal{L}(\varphi)]_n = 0$.

Mini-batch extension. Similarly to [14], we enhance the algorithm by performing the sparse coding step on a few samples between every dictionary update. This is particularly appealing when we can process several samples in parallel. It is worth noting that there is a trade off between the number of input samples per mini-batch and their size. A big input sample contains a lot of redundant information and leads to a slower coding step. Conversely, a mini-batch of a few small but diverse samples is faster and can mitigate the effect of a single iteration biasing towards a specific scene.

Forgetting factor. In the first few iterations of the algorithm, the initial dictionary may not be informative for effective sparse representation. Thus, the corresponding coefficient maps α^t are inaccurate, compared to coefficient maps computed with later iterates. Consequently, we also introduce a forgetting factor $\gamma \geq 0$, which allocates more weight to newer samples than to older ones. In practice, during the update of the memory vector (20) and matrix (21), we weight the old and new ones respectively by $\theta^t \triangleq (1 - \frac{1}{t})^{1+\gamma}$ and $1 - \theta^t$.

Table 1. Average PSNR for various subsampling rates on 23 images from the Middlebury 2014 dataset.

| Method | 2× | 3× | 4× |
|---------------|-----------------|-----------------|-----------------|
| Linear | 30.72 dB | 30.39 dB | 29.97 dB |
| Guided Filter | 34.18 dB | 33.46 dB | 32.76 dB |
| Weighted TV | 36.85 dB | 35.58 dB | 34.77 dB |
| Proposed | 37.13 dB | 35.73 dB | 34.88 dB |

3. EXPERIMENTAL EVALUATION

To evaluate our multimodal imaging method, we focus on joint intensity-depth reconstruction. We consider two modalities ($L = 2$), where $\Phi_1 = \mathbf{I}$ is the sensing matrix associated to an intensity image and $\Phi_2 = \text{diag}(\varphi)$ is a random mask selecting a fraction of the depth map pixels. To generate \mathbf{y}_ℓ , Gaussian noise of 30dB PSNR is added to $\Phi_\ell \mathbf{x}_\ell$ on both modalities. Our quality metric is the prediction PSNR over the missing pixels of \mathbf{x}_2 . Since the reconstruction method (3) optimizes over two distinct variables ($\hat{\mathbf{x}}, \hat{\alpha}$), two options are available for the prediction: either using $\hat{\mathbf{x}}$ or using $\mathbf{D}\hat{\alpha} + \mathbf{x}^{10}$. In our experiments the second solution provides better performance. We rely only on subsampled data for both training and reconstruction, and set the number of convolutional kernels to $K = 32$ and size $P = 15 \times 15$.

Comparison with other methods

In this experiment, we first train a global dictionary using 160 mini-batches of 8 randomly selected patches of size 45×46 from the Middlebury dataset [29]. Then, to process each specific frame, the dictionary is specialized with 120 mini-batches of 8 patches sampled from that frame. Finally, the full (480×672) frame is reconstructed by solving (3) with the specialized dictionary. Note that, in contrast to patch-based dictionary learning, the input patches are not necessarily of the same size as the dictionary kernels. In fact, the full frame could be used to train the convolutional dictionary. However, using smaller patches instead, helps to accelerate the training.

Table 1 compares the average performance of our reconstruction procedure with three alternative approaches, listed in increasing complexity: linear interpolation, guided filtering [22], and weighted TV [25]. Note that, similarly to our method, the guided filter and the weighted TV both use intensity information as a guide for depth estimation. The parameters were hand-tuned using heuristics for every method in order to achieve the best average performance.

Figure 2 shows the specialized dictionary and the reconstruction results for the *Backpack* image. One can clearly recognize the typical image features manifested in the learned kernels. Some kernels present sharp edges or corners. Others show more elaborate grid-like features. Most paired kernels have striking similarities between depth and intensity in terms of shape, orientation, and alignment. These results highlight the ability of our method to learn multimodal convolutional dictionaries directly from noisy and subsampled data.

Online training on a video

In order to demonstrate the *online* capability of our learning algorithm, we use the *Road* intensity-depth video sequence from [30]. We use a mask with $2\times$ subsampling and add a 30 dB Gaussian noise. We start with a dictionary filled with Dirac deltas. On each 512×512 frame of the video, we extract 8 randomly chosen 50×50 patches and perform one mini-batch iteration of the learning algorithm. Then, we reconstruct the full frame using the current dictio-




Fig. 2. The top-left and top-right images are the intensity and depth modalities from the *Backpack* image. White pixels correspond to missing pixels due, for example, to occlusions. Bottom-left image shows the trained dictionary with each pair of corresponding intensity-depth kernels grouped. Bottom-right shows the reconstruction of the region highlighted in red for $2\times$ subsampling.

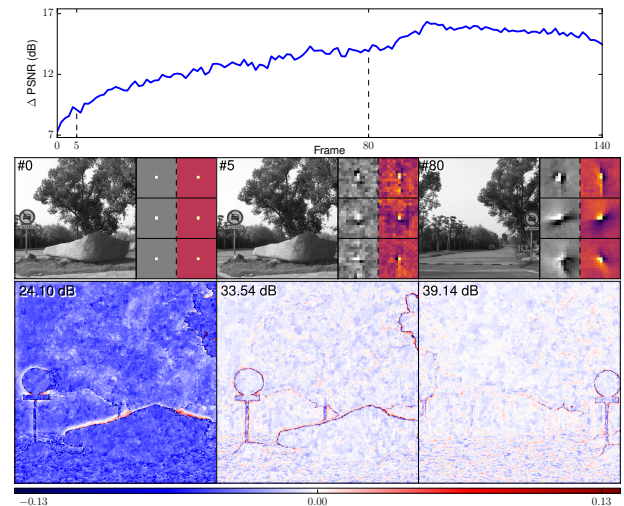


Fig. 3. Top row plots the evolution of the improvement in dB using the learned dictionary instead of deltas. Middle row shows the intensity frames 0, 5 and 80 next to three kernels from the corresponding dictionary. Bottom row shows the residual of the reconstructed depth.

nary and compare with the performances obtained with the initial dictionary of deltas. Figure 3 presents the evolution, as the video is streamed, of the PSNR improvement. As evident, the energy of the residual decreases as the dictionary improves, especially near object edges. Note that a temporary drop of quality might be observed when an unexpected feature appears in the scene. This decrease is then compensated when the dictionary adapts to the new feature.

4. REFERENCES

- [1] W.-K. Ma, J. M. Bioucas-Dias, J. Chanussot, and P. Gader, "Signal and image processing in hyperspectral remote sensing," *IEEE Signal Process. Mag.*, vol. 22, pp. 22–23, January 2014.
- [2] H. Fatakdawala and *et al.*, "Multimodal in vivo imaging of oral cancer using fluorescence lifetime, photoacoustic and ultrasound techniques," *Biomed. Opt. Express*, vol. 4, no. 9, pp. 1724–1741, September 2013.
- [3] J. Diebel and S. Thrun, "An application of Markov random field to range sensing," in *Proc. Advances in Neural Information Processing Systems*, Vancouver, BC, December 5-8, 2005, vol. 18, pp. 291–298.
- [4] B. A. Olshausen and D. J. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature*, vol. 381, pp. 607–609, June 1996.
- [5] B. A. Olshausen and D. J. Field, "Sparse coding with an overcomplete basis set: A strategy employed by V1?," *Vision Research*, vol. 37, no. 23, pp. 3311–3325, 1997.
- [6] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4311–4322, November 2006.
- [7] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, "Non-local sparse models for restoration," in *Proc. IEEE Int. Conf. Comp. Vis. (ICCV)*, Kyoto, Japan, September 29-October 2, 2009, pp. 2272–2279.
- [8] J. Yang, Z. Wang, Z. Lin, S. Cohen, and T. Huang, "Coupled dictionary training for image super-resolution," *IEEE Trans. Image Process.*, vol. 21, no. 8, pp. 3467–3478, August 2012.
- [9] Q. Wei, J. Bioucas-Dias, N. Dobigeon, and J.-Y. Tourneret, "Hyperspectral and Multispectral Image Fusion based on a Sparse Representation," pp. 1–27, sep 2014.
- [10] A. M. Tillmann, Y. C. Eldar, and J. Mairal, "DOLPHIn - Dictionary learning for phase retrieval," *IEEE Trans. Signal Process.*, 2016.
- [11] M. D. Zeiler, D. Krishnan, G. W. Taylor, and R. Fergus, "Deconvolutional networks," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, San Francisco, CA, USA, June 13-18, 2010, pp. 2528–2535.
- [12] J. Yang, K. Yu, and T. Huang, "Supervised translation-invariant sparse coding," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, San Francisco, CA, USA, June 13-18, 2010, pp. 3517–3524.
- [13] B. Wohlberg, "Efficient algorithms for convolutional sparse representations," *IEEE Trans. Image Process.*, vol. 25, no. 1, pp. 301–315, January 2016.
- [14] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online Learning for Matrix Factorization and Sparse Coding," *Journal of Machine Learning Research*, vol. 11, pp. 19–60, aug 2010.
- [15] A. Mensch, J. Mairal, B. Thirion, and G. Varoquaux, "Dictionary Learning for Massive Matrix Factorization," *International Conference on Machine Learning*, 2016.
- [16] A. Mensch, J. Mairal, G. Varoquaux, and B. Thirion, "Subsampled online matrix factorization with convergence guarantees," in *Proc. NIPS Workshop on Optimization for Machine Learning*, Barcelona, Spain, 2016, number 2, pp. 7–11.
- [17] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear Phenomena*, vol. 60, no. 1-4, pp. 259–268, nov 1992.
- [18] A. Beck and M. Teboulle, "Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems," *IEEE Trans. Image Process.*, 2009.
- [19] M.V. Afonso, J.M. Bioucas-Dias, and M.A.T. Figueiredo, "Fast Image Recovery Using Variable Splitting and Constrained Optimization," *IEEE Transactions on Image Processing*, vol. 19, no. 9, pp. 1–11, 2010.
- [20] U. S. Kamilov, "A parallel proximal algorithm for anisotropic total variation minimization," *IEEE Trans. Image Process.*, vol. 26, no. 2, pp. 539–548, February 2017.
- [21] J. Kopf, M. F. Cohen, D. Lischinski, and M. Uyttendaele, "Joint bilateral upsampling," in *ACM Trans. Graph. (Proc. SIGGRAPH 2007)*, San Diego, CA, USA, August 5-9, 2007, vol. 26, p. 96.
- [22] K. He, J. Sun, and X. Tang, "Guided image filtering," *IEEE Trans. Patt. Anal. and Machine Intell.*, vol. 35, no. 6, pp. 1397–1409, June 2013.
- [23] I. Tosic and P. Frossard, "Dictionary learning for stereo image representation," *IEEE Trans. Image Process.*, vol. 20, no. 4, pp. 921–934, 2011.
- [24] I. Tosic and S. Drewes, "Learning joint intensity-depth sparse representations," *IEEE Transactions on Image Processing*, vol. 23, no. 5, pp. 2122–2132, 2014.
- [25] J. Castorena, U. S. Kamilov, and P. T. Boufounos, "Auto-calibration of lidar and optical cameras via edge alignment," in *IEEE Int. Conf. Acoustics, Speech and Signal Process. (ICASSP 2016)*, may 2016, vol. 2016-May, pp. 2862–2866, Institute of Electrical and Electronics Engineers Inc.
- [26] U. S. Kamilov and P. T. Boufounos, "Depth superresolution using motion adaptive regularization," *Proc. 2015 IEEE Int. Conf. Multimedia & Expo Workshops (ICMEW 2016)*, pp. 1–13, 2016.
- [27] Y. Liu, X. Chen, R. K. Ward, and Z. Jane Wang, "Image fusion with convolutional sparse representation," *IEEE Signal Processing Letters*, vol. 23, no. 12, pp. 1882–1886, Dec 2016.
- [28] J. Mairal, F. Bach, and J. Ponce, "Sparse modeling for image and vision processing," *Foundations and Trends in Machine Learning*, vol. 8, no. 2-3, pp. 1–199, 2014.
- [29] D. Scharstein, H. Hirschmüller, Y. Kitajima, G. Krathwohl, N. Nešić, X. Wang, and P. Westling, "High-resolution stereo datasets with subpixel-accurate ground truth," *Lecture Notes in Computer Science*, vol. 8753, no. 1, pp. 31–42, 2014.
- [30] G. Zhang, J. Jia, T. T. Wong, and H. Bao, "Consistent depth maps recovery from a video sequence," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 6, pp. 974–988, 2009.