

Compressed-Sensing Recovery of Images and Video Using Multihypothesis Predictions

Chen Chen, Eric W. Tramel, and James E. Fowler

Department of Electrical and Computer Engineering, Geosystems Research Institute (GRI)
Mississippi State University, MS 39762 USA

Abstract—Compressed-sensing reconstruction of still images and video sequences driven by multihypothesis predictions is considered. Specifically, for still images, multiple predictions drawn for an image block are made from spatially surrounding blocks within an initial non-predicted reconstruction. For video, multihypothesis predictions of the current frame are generated from one or more previously reconstructed reference frames. In each case, the predictions are used to generate a residual in the domain of the compressed-sensing random projections. This residual being typically more compressible than the original signal leads to improved reconstruction quality. To appropriately weight the hypothesis predictions, a Tikhonov regularization to an ill-posed least-squares optimization is proposed. Experimental results demonstrate that the proposed reconstructions outperform alternative strategies not employing multihypothesis predictions.

I. INTRODUCTION

The compressed sensing (CS) of images and video faces several challenges including a large computational cost associated with multidimensional signal reconstruction and a huge memory burden when the random sampling operator is represented as a dense matrix. To address these issues, structurally random matrices (SRMs) (e.g., [1]) can be used to provide a sampling process with little computation and memory. An alternative to SRMs is to limit CS sampling to relatively small blocks (e.g., [2, 3]). Block-based CS (BCS) with smoothed projected Landweber reconstruction (BCS-SPL) [3], as well as a multiscale variant (MS-BCS-SPL) [4] deployed in the domain of a discrete wavelet transform (DWT), typically provides much faster reconstruction than techniques based on full-image CS sampling. For video, a motion-compensated reconstruction (MC-BCS-SPL) [5] extends this advantage to the CS reconstruction of video in which one or more frames are used to make predictions of the current frame such that the resulting residual is more efficiently reconstructed.

In this paper, we extend this concept of prediction plus residual reconstruction to the use of multihypothesis (MH) predictions (e.g., [6]). That is, we couple BCS-SPL with MH predictions for both still-image as well as video reconstruction. For video, we cull the multiple hypothesis predictions from previously reconstructed frames. For still images, we first reconstruct the image with an initial BCS-SPL reconstruction, cull predictions for each image block from spatially surrounding blocks, and then finally reconstruct the resulting prediction-residual image. We also consider a similar multiscale variant in which the MH predictions occur in the wavelet

domain.

In all cases, we determine the MH predictions in the domain of CS random projections. Due to the ill-posed nature of the resulting prediction problem, we apply Tikhonov regularization [7] to arrive at a solution. Experimental results for video demonstrate that this Tikhonov-regularized reconstruction usually provides higher PSNR as compared to a similar ℓ_1 -regularized approach [8, 9], as well as compared to a straightforward intraframe BCS-SPL reconstruction. The same Tikhonov-regularized reconstruction is then applied to the MH-based still-image reconstruction; corresponding experimental results indicate a significant gain in PSNR as compared to the original BCS-SPL as well as to a popular still-image reconstruction based on total-variation (TV) minimization [10].

II. BACKGROUND

Suppose we want to recover real-valued signal $x \in \mathbb{R}^N$ from M measurements such that $M \ll N$; i.e., $y = \Phi x$, where $y \in \mathbb{R}^M$, and Φ is an $M \times N$ measurement matrix with subsampling rate, or *subrate*, being $S = M/N$. CS theory holds that, if x is sufficiently sparse in some transform basis Ψ , then x is recoverable from y by the optimization,

$$\hat{x} = \arg \min_{x \in \mathbb{R}^N} \|\Psi x\|_1, \quad \text{such that } y = \Phi x, \quad (1)$$

as long as Φ and Ψ are sufficiently incoherent, and M is sufficiently large. High-dimensional signals, such as images or video, impose a huge memory burden when explicitly storing the sampling operator Φ as a dense matrix. In addition, the reconstruction process will be time consuming if the dimensionality is large. To assuage the computation complexity, in [2, 3], an image is partitioned into smaller blocks while sampling is applied on a block-by-block basis. In such BCS, the global measurement matrix takes a block-diagonal structure, $\Phi = \text{diag}(\Phi_B, \dots, \Phi_B)$, wherein Φ_B independently samples blocks within the image. That is, $y_i = \Phi_B x_i$, where x_i is a column vector with length B^2 representing block i of the image, and Φ_B is a $M_B \times B^2$ measurement matrix such that the subrate of BCS is $S = M_B/B^2$. In [2, 3], reconstruction uses a procedure that couples projected Landweber (PL) iteration with a smoothing operation intended to reducing blocking artifacts. The overall technique was called BCS-SPL in [3].

BCS-based techniques such as BCS-SPL that rely on a block-based sampling operator can be at a disadvantage in terms of reconstruction quality since CS sampling generally

works better the more global it is. To improve reconstruction quality, in [4], BCS-SPL was deployed independently within each subband of each decomposition level of a wavelet transform of an image to provide multiscale sampling and reconstruction; the resulting algorithm for image reconstruction was called MS-BCS-SPL.

For video, one can simply apply a CS image sampling and reconstruction independently frame by frame (i.e., “intraframe” CS sampling and reconstruction). Alternatively, one can incorporate motion estimation (ME) and motion compensation (MC) into the CS reconstruction of video while maintaining the same frame-by-frame image sampling (e.g., [5, 8, 9, 11, 12]). In the latter approach, a motion-compensated prediction of the current frame is created during reconstruction such that some CS image-reconstruction algorithm is applied to the residual between the current frame and its ME/MC prediction. Specifically, if a ME/MC prediction is similar to the original frame, then the prediction residual will be more amenable to CS reconstruction since the residual is typically more compressible than the original frame itself. Suppose \tilde{x} is a prediction of original frame x which satisfies $\tilde{x} \approx x$; the residual of between the two signals is $r = x - \tilde{x}$. With a measurement basis Φ , the projection of r is $q = \Phi r = y - \Phi \tilde{x}$. The final reconstruction of y is calculated as $\hat{x} = \tilde{x} + \text{Reconstruct}(q, \Phi)$, where $\text{Reconstruct}(\cdot)$ is some suitable CS image reconstruction.

To produce a highly compressible residual, one should create a prediction that is as close as possible to x , which implies that the following optimization problem is desired,

$$\tilde{x} = \arg \min_{p \in \mathcal{P}(x_{\text{ref}})} \|x - p\|_2^2, \quad (2)$$

where $\mathcal{P}(x_{\text{ref}})$ is the set of all ME/MC predictions that can be made from a reference frame, x_{ref} . However, since x is unknown in CS reconstruction, solving (2) directly is infeasible. Instead, one approach is to reformulate (2) as

$$\tilde{x} = \arg \min_{p \in \mathcal{P}(x_{\text{ref}})} \|\hat{x} - p\|_2^2, \quad (3)$$

wherein some initial reconstruction, \hat{x} , is used as a proxy for x in (2); this is, in fact, the approach taken in [5, 12].

An alternative is to recast the optimization of (2) from the ambient signal domain of x into the measurement domain of y ; specifically,

$$\tilde{x} = \arg \min_{p \in \mathcal{P}(x_{\text{ref}})} \|y - \Phi p\|_2^2. \quad (4)$$

The Johnson-Lindenstrauss (JL) lemma [13] holds that L points in \mathbb{R}^N can be projected into a K -dimensional subspace while approximately maintaining pairwise distances as long as $K \geq O(\epsilon^{-2} \log L)$ for any $0 < \epsilon < 1$. This suggests that the solution of (4) will likely coincide with that of (2). The next section explores two general strategies for implementing (4) with MH prediction.

III. MULTIHYPOTHESIS PREDICTIONS FOR VIDEO

A. MH with Tikhonov Regularization

For a MH CS reconstruction of video, the goal is to reformulate (4) so that, instead of choosing a single prediction, or *hypothesis*, we find an optimal linear combination of all hypotheses contained in some search set; i.e., (4) becomes $\tilde{x}_{t,i} = H_{t,i} \hat{w}_{t,i}$ where

$$\hat{w}_{t,i} = \arg \min_w \|y_{t,i} - \Phi H_{t,i} w\|_2^2, \quad (5)$$

and we have also recast (4) for block-based prediction with i being the block index and t being the temporal frame index. Here, $H_{t,i}$ is a matrix of dimensionality $B^2 \times K$ whose columns are the rasterizations of the possible blocks within the search space of the reference frames. In this context, $\hat{w}_{t,i}$ is a column vector which represents a linear combination of the columns of $H_{t,i}$. However, because $M \ll K$, the ill-posed nature of the problem requires some kind of regularization in order to differentiate among the infinite number of possible linear combinations which lie in the solution space of (5).

The most common approach to regularizing a least-squares problem is Tikhonov regularization [7] which imposes an ℓ_2 penalty on the norm of $\hat{w}_{t,i}$,

$$\hat{w}_{t,i} = \arg \min_w \|y_{t,i} - \Phi H_{t,i} w\|_2^2 + \lambda \|\Gamma w\|_2^2, \quad (6)$$

where Γ is known as the Tikhonov matrix; this strategy for MH prediction was initially proposed in [11]. The Γ term allows the imposition of prior knowledge on the solution—we take the approach that hypotheses which are the most dissimilar from the target block should be given less weight than hypotheses which are most similar. Specifically, we propose a diagonal Γ in the form of $\Gamma_{j,j} = \|y_{t,i} - \Phi h_{t,j}\|_2^2$, where $h_{t,j}$ are the columns of $H_{t,i}$, $J = 1, \dots, K$. For each block then, $\hat{w}_{t,i}$ can be calculated directly by the closed form solution,

$$\hat{w}_{t,i} = \left((\Phi H_{t,i})^T (\Phi H_{t,i}) + \lambda^2 \Gamma^T \Gamma \right)^{-1} (\Phi H_{t,i})^T y_{t,i}. \quad (7)$$

B. MH with ℓ_1 Regularization

An alternate to the Tikhonov regularization used in (6) was suggested in [8, 9]. Specifically, it was assumed in [8, 9] that the MH weights $\hat{w}_{t,i}$ in (5) are sparse; i.e., only relatively few of the possible hypotheses in $H_{t,i}$ should contribute the prediction. As a consequence of this assumption, the reconstructions in [8, 9] essentially impose an ℓ_1 penalty term on $\hat{w}_{t,i}$; i.e.,

$$\hat{w}_{t,i} = \arg \min_w \|\Phi H_{t,i} w - y_{t,i}\|_2^2 + \lambda \|w\|_1. \quad (8)$$

The intuition here is that only a few blocks within the search space contribute significantly to the linear combination. However, in the context of CS reconstruction, a regularization enforcing sparsity is needlessly restrictive on the structure of $\hat{w}_{t,i}$, which can potentially result in lower prediction quality. Furthermore, Tikhonov regularization in the form of (6) is a much more amenable solution than ℓ_1 regularization in terms of scalability and computation time, as well—with the ℓ_1 penalty, the optimization in (8) is approached as a traditional

CS problem using some generic CS solver independently on each block, while Tikhonov regularization simply uses (7).

IV. MULTIHYPOTHESIS PREDICTIONS FOR IMAGES

A. MH-BCS-SPL

Above, we considered the use of MH prediction for the CS of video. We now consider applying MH prediction for the CS of a still image. In this case, (6) becomes

$$\hat{w}_i = \arg \min_w \|y_i - \Phi H_i w\|_2^2 + \lambda \|\Gamma w\|_2^2, \quad (9)$$

and the matrix of hypotheses, H_i , is assembled from an initial reconstruction, \bar{x} , of the image x using either BCS-SPL or MS-BCS-SPL. That is, for each block in \bar{x} , MH predictions are generated from blocks spatially surrounding \bar{x} in the initial reconstruction. Specifically, suppose image x is split into blocks of size $B \times B$ in BCS; each block is further divided into subblocks of size $b \times b$. MH predictions are created for each individual subblock of the block by sliding a $b \times b$ mask across the entire search window to create all candidate predictions for each subblock. Since the block size is $B \times B$, the region in the block outside of the $b \times b$ subblock is set to all zeros; the resulting $B \times B$ “zero-padded” block is then placed as a column in H_i ; H_i thus contains all the predictions for all of the subblocks of block i . This subblock-based MH-prediction process is illustrated in Fig. 1(a)–(c).

The parameter λ in (9) controls the regularization. Unfortunately, there does not appear to be a straightforward approach for finding an optimal value without foreknowledge of x . Some possible approaches to choose an appropriate λ include the L-curve, generalized cross validation, and the discrepancy principle. Through empirical analysis, we test a set of λ values and choose the one gives the best performance.

We incorporate the proposed MH-based prediction into BCS-SPL image reconstruction, resulting in a technique we call MH-BCS-SPL (see Algorithm 1). In MH-BCS-SPL, MH prediction and residual reconstruction are repeated with increasing subblock size in order to improve the quality of the recovered image.

Specifically, the original BCS-SPL reconstruction of [3] uses a block size of $B = 32$ and a dual-tree DWT (DDWT) [14] as the sparsity transform Ψ . In MH-BCS-SPL, we start with an initial subblock size of $b = 16$ and an initial search window of $w = 8$. The subblock size b and search window w are increased based on a criterion involving structural similarity (SSIM) [15]. As a stopping criterion, we apply cross validation [16] to predict the performance. Specifically, three measurements as a holdout set y_H are used for the performance test. For example, at subrate = 0.1 and block size $B = 32$, the measurement matrix $\Phi \in \mathbb{R}^{102 \times 1024}$ has three more rows than $\Phi_R \in \mathbb{R}^{99 \times 1024}$ which is used for reconstruction. $\Phi_H \in \mathbb{R}^{3 \times 1024}$ is the measurement matrix for the holdout set. In other words, $\Phi = [\Phi_R; \Phi_H]$. The residual calculated in the projected domain is

$$R = \|\Phi_H x - \Phi_H \hat{x}\|_2 = \|y_H - \Phi_H \hat{x}\|_2. \quad (10)$$

This means that, if \hat{x} is close to x , then R should be small.

Algorithm 1 MH-BCS-SPL

Input: $y, \Phi = [\Phi_R; \Phi_H], \Psi, \bar{x}, b$ (initial subblock size), w (initial search window size), $B = 32$ (block size), τ .

Output: \bar{x} .

Initialization: $i = 1, \tilde{x}_0 = \bar{x}, s_0 = 0, R_0 = +\infty$.

repeat

$\tilde{x}_i = \text{MH_Prediction}(\bar{x}, y, \Phi_R, b, w, B)$

$\hat{x}_i = \tilde{x}_i + \text{BCS-SPL}(y - \Phi_R \tilde{x}_i, \Phi_R, \Psi, B)$

Compute $s_i = \text{SSIM}(\hat{x}_i, \hat{x}_{i-1}), R_i = \|y_H - \Phi_H \hat{x}_i\|_2$

if $b < B$ **then**

if $(R_i < R_{i-1} \ \& \ |s_i - s_{i-1}| \leq \tau)$ **or** $R_i > R_{i-1}$ **then**

$b \leftarrow b \times 2, w \leftarrow w \times 2$

end if

end if

Update $\bar{x} \leftarrow \hat{x}_i$

$i = i + 1$

until $R_i > R_{i-1} \ \& \ b = B$

B. MH-MS-BCS-SPL

As described in [4], MS-BCS-SPL performs both CS measurement and reconstruction in the wavelet domain; i.e., the CS measurement process becomes $y = \Phi \Omega x$, where Ω is a 3-level DWT with the popular 9/7 biorthogonal wavelet. Block size depends on transform level, with $B_l = 16, 32,$ and 64 for levels $l = 1, 2,$ and 3 , respectively ($l = 3$ is the highest-resolution level). A DDWT is again used as the sparsity transform Ψ .

We now formulate a MH version of MS-BCS-SPL by performing MH predictions within the wavelet domain. That is, in MH-MS-BCS-SPL, multiple predictions for a block are made using the procedure of Fig. 1(a) applied in the wavelet domain of the MS-BCS-SPL reconstructed image. In the case that the subblocks are smaller than the block size (i.e., $b_l < B_l$), the MH predictions are carried out within a subband of a DWT. However, for the $b_l = B_l$ case, predictions are calculated using a redundant DWT (RDWT) for Ω . Such a RDWT is an overcomplete transform, effectively created by eliminating downsampling from the traditional DWT; see, e.g., [17]. In this latter case, the multiple hypotheses are culled from the various RDWT phases associated with the subband as illustrated in Fig. 1(d).

In our experiments, $\lambda = 0.035$ is used when $b_l < B_l$, and $\lambda = 0.5$ when $b_l = B_l$. The initial subblock size b_l is one eighth of the block size B_l at each decomposition level. The initial search window is set to $w = 1$. For the stopping criterion, we did not apply cross validation, as we found that the reduced number of measurements y due to the holdout set would adversely affect the recovery performance since CS sampling is deployed in the wavelet domain. In this case, a threshold τ_s calculated from the SSIM between two successive residual reconstructed images is used as a stopping criterion: $\tau_s = 0.9995$ when subrate = 0.1, 0.2, and 0.3; $\tau_s = 0.99995$ when subrate = 0.4 and 0.5. Algorithm 2 details the MH-MS-BCS-SPL procedure.

V. RESULTS

A. Video

We consider the first three consecutive frames, $x_1, x_2,$ and x_3 of a given video sequence—the first and third frames,

Algorithm 2 MH-MS-BCS-SPL

Input: $y, \Phi, \Psi, \bar{x}, L = 3, \{b_l\}, w, \{B_l\}, \tau_s$.
Output: \hat{x} .
Initialization: $i = 1, \tilde{x}_0 = \bar{x}, s_0 = 0$.
repeat
 if $b_l = B_l$ for each level l **then**
 $\tilde{x}_i = \Omega_{\text{RDWT}} \bar{x}$ {Perform redundant wavelet transform}
 else
 $\tilde{x}_i = \Omega_{\text{DWT}} \bar{x}$ {Perform discrete wavelet transform}
 end if
 for $1 \leq l \leq L$ **do**
 for each subband $\theta \in \{H, V, D\}$ **do**
 $\tilde{x}_i(\theta) = \text{MH_Prediction}(\tilde{x}_i(\theta), y(\theta), \Phi, B_l, b_l, w)$
 end for
 end for
 $\tilde{x}_i = \Omega^{-1} \tilde{x}_i$
 $\hat{x}_i = \tilde{x}_i + \text{BCS-SPL}(y - \Phi \tilde{x}_i, \Phi, \Psi, \{B_l\})$
 Compute $s_i = \text{SSIM}(\hat{x}_i, \hat{x}_{i-1})$
 if $|s_i - s_{i-1}| \leq \tau$ **and** $b_l < B_l$ for each level l **then**
 for $1 \leq l \leq L$ **do**
 $b_l \leftarrow b_l \times 2$
 end for
 $w \leftarrow w \times 2$
 end if
 Update $\bar{x} \leftarrow \hat{x}_i$
until $|s_i - s_{i-1}| \leq \tau_s$

x_1 and x_3 , are used as reference frames, while the second frame, x_2 , is the “test frame” used to measure reconstruction performance. In all cases, the reference frame is BCS sampled with a relatively high subrate of $S_1 = 0.5$ and reconstructed using BCS-SPL [3]¹. On the other hand, the test frame is BCS sampled using a range of subrates, $S_2 \leq S_1$. For video-frame reconstruction, we use a block size of $B = 16$ for BCS and a DWT with 4 levels of decomposition as the sparsity basis for BCS-SPL reconstruction. Block-based sampling operator Φ_B is a $B \times B$ dense Gaussian matrix.

The reconstructed reference frames are used to create a prediction of each block of the test frame; afterward, residual reconstruction of the test frame is conducted. We investigate the Tikhonov and ℓ_1 prediction strategies as discussed in Sec. III. For the Tikhonov approach, we use a regularization parameter of $\lambda = 0.25$. Additionally, for the ℓ_1 -regularized MH prediction, we use GPSR² [18] to find the weights. We also consider performance of the straightforward BCS-SPL reconstruction of the test frame independent of the reference frame. In all cases, a spatial window size of ± 15 pixels about the current block is used as the search space for finding the hypotheses.

The PSNR performance of the test-frame reconstruction as the subrate, S_2 , for the test frame varies is presented in Table I. As can be seen in Table I, the proposed Tikhonov-regularized MH prediction provides significantly superior reconstruction for x_2 at low subrates as compared to the ℓ_1 -regularized prediction of [8]. For higher subrates near $S_2 \approx 0.5$, the performance of the ℓ_1 regularization is generally more competitive, and even exceeds that of the proposed Tikhonov regularization for the *News* sequence at $S_2 = 0.5$. However,

such a high-subrate case is of less interest than low-subrate reconstructions due to the necessity of maintaining the subrate of non-key frames as low as possible to minimize the overall sampling rate of the system.

In terms of computation, MH prediction performs much more quickly than the ℓ_1 method, taking just a few minutes for a single frame reconstruction, while the ℓ_1 method can take exceedingly long to calculate, up to 4 or 5 hours for a single frame.

B. Still Images

The performance of the MH-BCS-SPL and MH-MS-BCS-SPL is evaluated on a number of grayscale images of size 512×512 (see Fig. 2) with $\tau = 0.0001$. We compare to the original BCS-SPL [3] and MS-BCS-SPL [4] as well as to the TV reconstruction described in [10] and a multiscale variant of GPSR as described in [19]. Block-based sampling operator Φ_B is a $B \times B$ dense Gaussian matrix; on the other hand, TV uses the scrambled block-Hadamard SRM of [1] to provide a fast whole-image CS sampling. The multiscale GPSR (MS-GPSR) uses the same Ω as MS-BCS-SPL in implementing GPSR reconstruction at each DWT level. We use our implementations³ of BCS-SPL and MS-BCS-SPL, and ℓ_1 -MAGIC⁴ for TV.

The reconstruction performance of the various algorithms under consideration is presented in Table III. In all cases except the “Barbara” and “Barbara2” images, MH-MS-BCS-SPL performs uniformly better than other algorithms. For “Barbara,” MH-BCS-SPL provides a substantial gain in reconstruction quality over TV, generally on order of a 5- to 7-dB increase in PSNR. A visual comparison of the various algorithms is shown in Fig. 3.

As can be seen in Table II, in terms of execution time, reconstruction with MH-BCS-SPL and MH-MS-BCS-SPL is, as expected, slower than BCS-SPL and MS-BCS-SPL due to iterated MH prediction. Both algorithms run for less than 3 minutes on a quadcore 2.67-GHz machine. On the other hand, the execution times of TV are much slower than MH-BCS-SPL and MH-MS-BCS-SPL, with TV requiring more than 20 minutes to reconstruct a single image even with fast SRM implementation of the sampling operator.

VI. CONCLUSIONS

In this paper, we considered how the high degree of spatial correlation in images and frame-to-frame temporal correlation in video signals can be exploited to enhance CS reconstruction. In essence, we formed MH predictions using a distance-weighted Tikhonov regularization to find the best linear combination of hypotheses. The MH predictions were used to create a measurement-domain residual of the signal to be recovered—such a residual is typically more compressible than the original signal making it more amenable to CS reconstruction. The proposed approach to MH prediction showed a significant improvement in reconstruction quality over several alternative

¹<http://www.ece.msstate.edu/~fowler/BCSSPL/>

²<http://www.lx.it.pt/~mtf/GPSR/>

³<http://www.ece.msstate.edu/~fowler/BCSSPL/>

⁴<http://www.l1magic.org>

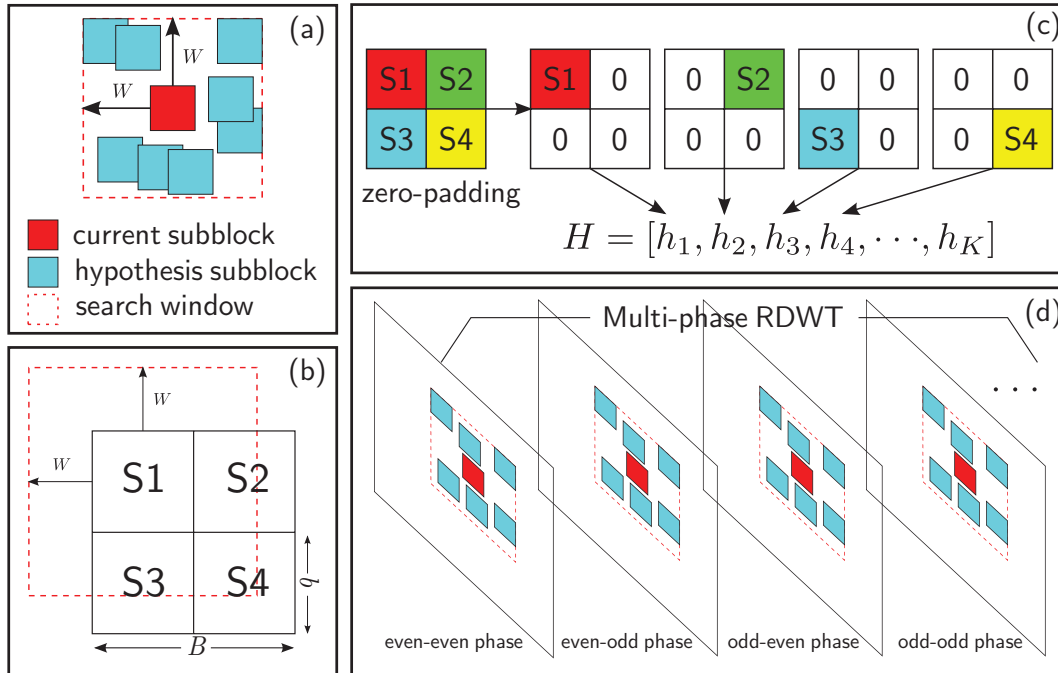


Fig. 1. (a) Generation of multiple hypotheses for a subblock in a search window; (b) generation of multiple hypotheses for each $b \times b$ subblock within a $B \times B$ block; (c) zero-padding of subblock predictions into blocks; (d) generation of multiple hypotheses within the multiple phases of an RDWT (four phases are shown—each phase has the same size as the corresponding DWT subband and results from various combinations of even and odd downsampling in the horizontal and vertical directions).

reconstructions for both image and video reconstruction, including, for video, a straightforward intraframe reconstruction as well as an alternative ℓ_1 -regularized prediction/residual reconstruction; and, for still images, a popular TV-based reconstruction.

REFERENCES

- [1] T. T. Do, T. D. Tran, and L. Gan, "Fast compressive sampling with structurally random matrices," in *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, Las Vegas, NV, March 2008, pp. 3369–3372.
- [2] L. Gan, "Block compressed sensing of natural images," in *Proceedings of the International Conference on Digital Signal Processing*, Cardiff, UK, July 2007, pp. 403–406.
- [3] S. Mun and J. E. Fowler, "Block compressed sensing of images using directional transforms," in *Proceedings of the International Conference on Image Processing*, Cairo, Egypt, November 2009, pp. 3021–3024.
- [4] J. E. Fowler, S. Mun, and E. W. Tramel, "Multiscale block compressed sensing with smoother projected Landweber reconstruction," in *Proceedings of the European Signal Processing Conference*, Barcelona, Spain, August 2011, pp. 564–568.
- [5] S. Mun and J. E. Fowler, "Residual reconstruction for block-based compressed sensing of video," in *Proceedings of the IEEE Data Compression Conference*, J. A. Storer and M. W. Marcellin, Eds., Snowbird, UT, March 2011, pp. 183–192.
- [6] G. J. Sullivan, "Multi-hypothesis motion compensation for low bitrate video coding," in *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, vol. 5, Minneapolis, MN, April 1993, pp. 437–440.
- [7] A. N. Tikhonov and V. Y. Arsenin, *Solutions of Ill-Posed Problems*. Washington, D.C.: V. H. Winston & Sons, 1977.
- [8] T. T. Do, Y. Chen, D. T. Nguyen, N. Nguyen, L. Gan, and T. D. Tran, "Distributed compressed video sensing," in *Proceedings of the International Conference on Image Processing*, Cairo, Egypt, November 2009, pp. 1393–1396.
- [9] J. Prades-Nebot, Y. Ma, and T. Huang, "Distributed video coding using compressive sampling," in *Proceedings of the Picture Coding Symposium*, Chicago, IL, May 2009.
- [10] E. Candès, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207–1223, August 2006.
- [11] E. W. Tramel and J. E. Fowler, "Video compressed sensing with multi-hypothesis," in *Proceedings of the IEEE Data Compression Conference*, J. A. Storer and M. W. Marcellin, Eds., Snowbird, UT, March 2011, pp. 193–202.
- [12] H. Jung and J. C. Ye, "Motion estimated and compensated compressed sensing dynamic magnetic resonance imaging: What we can learn from video compression techniques," *Imaging Systems and Technology*, vol. 20, no. 2, pp. 81–98, June 2010.
- [13] W. B. Johnson and J. Lindenstrauss, "Extensions of Lipschitz mappings into a Hilbert space," *Contemporary Mathematics*, vol. 26, pp. 189–206, 1984.
- [14] N. G. Kingsbury, "Complex wavelets for shift invariant analysis and filtering of signals," *Journal of Applied Computational Harmonic Analysis*, vol. 10, pp. 234–253, May 2001.
- [15] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, April 2004.
- [16] R. Ward, "Compressed sensing with cross validation," *IEEE Transactions on Information Theory*, vol. 55, no. 11, pp. 5773–5782, December 2009.
- [17] J. E. Fowler, "The redundant discrete wavelet transform and additive noise," *IEEE Signal Processing Letters*, vol. 12, no. 9, pp. 629–632, September 2005.
- [18] M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright, "Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems," *IEEE Journal on Selected Areas in Communications*, vol. 1, no. 4, pp. 586–597, December 2007.
- [19] P. Schniter, L. C. Potter, and J. Ziniel, "Fast Bayesian matching pursuit: Model uncertainty and parameter estimation for sparse linear models," *IEEE Transactions on Signal Processing*, 2008, submitted.

TABLE I
PSNR IN dB OF RESIDUAL RECONSTRUCTION (RR) OF VIDEO

	Subrate				
	0.1	0.2	0.3	0.4	0.5
Foreman					
RR w/ MH-TIK	31.5	33.8	35.0	35.9	36.7
RR w/ MH- ℓ_1	28.0	31.0	32.9	34.5	36.3
BCS-SPL	25.1	27.5	29.9	31.8	33.7
Susie					
RR w/ MH-TIK	34.2	36.3	37.3	38.0	38.6
RR w/ MH- ℓ_1	30.7	33.4	35.2	36.7	38.1
BCS-SPL	29.4	31.9	33.6	35.1	36.4
Football					
RR w/ MH-TIK	25.5	27.6	29.1	30.3	31.4
RR w/ MH- ℓ_1	24.6	26.7	28.3	29.7	31.1
BCS-SPL	24.0	25.9	27.2	28.5	29.9
News					
RR w/ MH-TIK	30.9	31.6	32.0	32.1	32.2
RR w/ MH- ℓ_1	23.7	27.0	29.5	31.6	34.1
BCS-SPL	20.2	23.4	26.3	28.9	31.8



Fig. 2. The 512×512 grayscale still images used in the experiments. Top row (left to right): Lenna, Barbara, Barbara2, Goldhill; Bottom row (left to right): Mandrill, Peppers, Boat, Cameraman.




Fig. 3. Barbara (detail) for subrate = 0.1. Top-row (left to right): BCS-SPL, MH-BCS-SPL, MS-BCS-SPL; bottom-row (left to right): MS-GPSR, MH-MS-BCS-SPL, TV.

TABLE II
RECONSTRUCTION TIME FOR LENNA AT SUBRATE 0.3

Algorithm	Time (sec.)
BCS-SPL	14.38
MH-BCS-SPL	146.77
MS-BCS-SPL	12.93
MH-MS-BCS-SPL	45.98
MS-GPSR	138.40
TV	1211.96

TABLE III
IMAGE RECONSTRUCTION PSNR (dB)

		Subrate				
		0.1	0.2	0.3	0.4	0.5
Lenna						
BCS-SPL	Original	27.93	31.27	33.41	35.15	36.71
	MH	29.85	32.85	34.73	36.34	37.82
	MS	31.48	34.59	36.61	37.82	38.94
	MH-MS	31.61	34.88	36.79	38.32	39.74
MS-GPSR		30.30	33.60	35.29	36.29	37.73
TV		29.84	32.93	35.05	36.82	38.43
Barbara						
BCS-SPL	Original	22.28	23.78	25.32	26.88	28.51
	MH	27.89	31.46	33.63	35.68	37.29
	MS	23.84	25.12	26.05	27.27	28.83
	MH-MS	24.28	26.42	27.98	32.95	36.21
MS-GPSR		24.03	25.24	26.06	27.50	29.64
TV		22.95	24.48	26.30	28.42	30.78
Barbara2						
BCS-SPL	Original	23.64	25.56	27.19	28.75	30.38
	MH	26.83	30.19	32.03	33.71	35.27
	MS	25.14	27.31	29.12	30.34	31.66
	MH-MS	25.61	28.87	31.06	33.74	36.30
MS-GPSR		25.28	27.34	28.82	30.19	32.21
TV		23.90	26.24	28.51	30.74	32.97
Goldhill						
BCS-SPL	Original	26.57	28.92	30.42	31.72	33.03
	MH	27.67	30.28	31.82	33.26	34.62
	MS	28.97	31.06	32.75	33.67	34.64
	MH-MS	29.07	31.35	33.06	34.55	36.10
MS-GPSR		28.39	30.55	32.10	32.89	34.20
TV		27.52	29.87	31.63	33.23	34.81
Mandrill						
BCS-SPL	Original	20.45	21.80	22.90	23.96	25.10
	MH	20.47	22.36	24.03	25.36	26.67
	MS	21.45	23.07	24.69	25.54	26.46
	MH-MS	21.66	23.20	24.82	25.81	27.10
MS-GPSR		21.52	22.92	24.27	25.07	26.21
TV		20.53	22.02	23.44	24.94	26.52
Peppers						
BCS-SPL	Original	28.83	31.98	33.72	35.12	36.36
	MH	30.28	32.82	34.32	35.63	36.87
	MS	31.75	34.59	35.76	36.77	37.67
	MH-MS	32.08	34.73	35.96	37.15	38.37
MS-GPSR		29.25	31.90	33.03	34.22	35.75
TV		30.36	33.14	34.70	35.91	37.05
Boat						
BCS-SPL	Original	25.14	27.78	29.57	31.12	32.60
	MH	26.17	29.30	31.18	32.89	34.45
	MS	27.35	30.08	31.98	33.12	34.22
	MH-MS	27.46	30.38	32.23	33.89	35.46
MS-GPSR		26.52	29.21	30.84	32.06	33.64
TV		26.42	29.24	31.28	32.97	34.53
Cameraman						
BCS-SPL	Original	26.02	30.36	33.65	36.42	38.90
	MH	29.86	33.97	36.57	39.28	41.48
	MS	31.05	36.54	39.98	42.91	44.91
	MH-MS	31.71	38.08	42.85	45.63	48.15
MS-GPSR		29.53	34.76	39.59	41.82	45.08
TV		30.87	35.08	38.32	41.09	43.73