

Classification-Based Spatial Error Concealment for Visual Communications

Meng Chen, Yefeng Zheng, and Min Wu

Department of Electrical and Computer Engineering, University of Maryland, College Park, MD 20742, USA

Received 1 March 2005; Revised 11 August 2005; Accepted 22 August 2005

In an error-prone transmission environment, error concealment is an effective technique to reconstruct the damaged visual content. Due to large variations of image characteristics, different concealment approaches are necessary to accommodate the different nature of the lost image content. In this paper, we address this issue and propose using classification to integrate the state-of-the-art error concealment techniques. The proposed approach takes advantage of multiple concealment algorithms and adaptively selects the suitable algorithm for each damaged image area. With growing awareness that the design of sender and receiver systems should be jointly considered for efficient and reliable multimedia communications, we proposed a set of classification-based block concealment schemes, including receiver-side classification, sender-side attachment, and sender-side embedding. Our experimental results provide extensive performance comparisons and demonstrate that the proposed classification-based error concealment approaches outperform the conventional approaches.

Copyright © 2006 Hindawi Publishing Corporation. All rights reserved.

1. INTRODUCTION

Due to the various kinds of distortion and failures, part of a compressed image or video can be damaged or lost during transmission or storage. The widely used block-based visual coding systems have prompted a need of block-based error concealment on the decoder side. A number of concealment approaches have been proposed in recent years [1–8]. The smoothness and continuity properties in spatial or frequency domain, the repeating patterns, and other properties of visual data have been exploited to recover corrupted blocks from the survived surroundings. Through a benchmarking effort on the existing error concealment approaches, we have observed that different approaches are suitable for different image characteristics of a corrupted block and its surroundings, and none of the existing approaches is an all-time champion. This motivates us to explore a classification-based concealment approach that can combine the better performance of two state-of-the-art approaches in the literature. The classification-based approach also helps us achieve a better tradeoff between the concealment quality and the computation complexity on the receiver side. This is because some state-of-the-art approaches have rather high computation demand, and classification allows the computation power to be spent more strategically by performing expensive computations only when they are likely to offer a substantial gain in the concealment quality.

The classification in the proposed new framework of error concealment can be done either on the receiver side or on the sender side. The receiver-side classification uses the survived surrounding pixels to determine which candidate concealment approach would give better concealment quality for each corrupted block. As will be seen in this paper, the proposed receiver-side classification approach does not require side information and the overall concealment quality can outperform each candidate alone. To provide more proactive protection and further exploit the knowledge from the original, uncorrupted image, a few recent works in the literature [9–11] have jointly considered the design of sender and receiver systems to facilitate error concealment. We explore this sender-driven perspective for our classification-based concealment framework by obtaining a small amount of classification data on the sender side. As the classification results need to be delivered as side information from the sender to the receiver, we examine and compare two approaches for delivering the side information, namely, by attaching as part of the file header and by embedding in the image signal.

The paper is organized as follows. Section 2 provides a brief description of the evaluated algorithms and presents benchmarking results on a collection of natural and artificial images. Since the performance on various images shows the advantages and disadvantages of different error concealment techniques, a classification scheme on the receiver side is proposed in Section 3 to take advantages of the sweet spots

of existing techniques. The sender-side classification-based error concealment is proposed in Section 4 to further improve the concealment quality by supplying the ground truth of concealment technique selection to a receiver. We compare the concealment performance, computation complexity, and bandwidth usage of the three proposed schemes as well as their suitable application scenarios in Section 5, and conclude the paper in Section 6.

2. MOTIVATION

2.1. Prior work

Early explorations on spatial domain image concealment were reviewed in [1]. Among them, the multidirectional interpolation (MDI) approach performs pixel-domain interpolation along eight possible edge directions and considers the cases of both single edge and multiple edges [2]; the projection-onto-convex-sets (POCS) approach constrains the feasible solution set based on such prior information as smoothness and neighborhood consistency [3]; and the maximally smooth recovery (MSR) method makes use of the smoothness property of visual signals and formulates the concealment as a constrained energy minimization problem [4].

Three recent works in [5–7] have demonstrated the performance improvement on classic images such as “Lena” or “Barbara” over the earlier approaches. The geometric-structure-based (GSB) error concealment by Zeng and Liu [5] is a directional interpolation scheme, which uses the local geometric information extracted from the surroundings. Two layers of pixels surrounding a corrupted block are converted to a binary pattern to reveal the local geometric structure and to classify the block as flat or nonflat. For flat blocks, the projective interpolation technique of [12] is applied. For nonflat blocks, the edges inside the lost block are estimated by pairing significant transition points from the aforementioned binary pattern, and the lost pixels are recovered by bilinear interpolation along the edge directions.

The orientation adaptive sequential interpolation (OASI) scheme by Li and Orchard [6] employs a linear regression model. It first estimates the local characteristics from a neighborhood of about four layers of uncorrupted pixels, and then uses the model parameters obtained to estimate each missing pixel from its surrounding pixels. More specifically, the interpolation can be characterized by $S = \sum_{k=1}^N \alpha_k s_k$, where S is an estimate of the missing pixel and s_k 's are N neighboring pixels. The interpolation coefficients α_k form a vector α , which can be determined using the classical least-square method from an M -pixel neighborhood M_n with $M > N$, that is, $\alpha = (C^T C)^{-1} C y$. Here, y is an $M \times 1$ vector representing M pixels in the training area M_n ; C is an $M \times N$ matrix, and each of its M rows consists of N neighbors around the corresponding pixel in y . When $C^T C$ is singular, α_k is set to $1/N$.

The long range correlation (LRC) scheme by Zhang and Wang [7] exploits the repeating patterns in an image. It extracts a ring window surrounding the corrupted area,

TABLE 1: The names and the references for the benchmarked approaches.

Acronym	Name	Reference
MDI	Multidirectional interpolation	[2]
POCS	Projection-onto-convex-sets	[3]
MSR	Maximally smooth recovery	[4]
GSB	Geometric-structure-based	[5]
OASI	Orientation adaptive sequential interpolation	[6]
LRC	Long range correlation	[7]

searches for an area in the image that best matches the pattern of this ring in a mean-squared error sense, and replaces the corrupted area with the pattern inside the best matching ring. LRC is also exploited in the recent image inpainting work by Bertalmio et al. [8], where the basic texture synthesis procedure for concealing the lost area is similar to the LRC concealment algorithm. By simultaneously filling in the structure and texture information of missing areas, the inpainting technique demonstrates excellent subjective quality when the missing area is relatively small compared with the size of the whole image. It is worth noticing that the image inpainting technique focuses more on the overall subjective quality and is not designed to optimize an objective error measure of the concealment quality (such as MSE or PSNR) on many small blocks.

2.2. Performance benchmarking

If an image is compressed by a block-based codec and transmitted over an error-prone channel, the error impairments are likely to be in the block domain. We focus on isolated block concealment in this work because block-based codecs are dominant for image or video transmission and the interleaving techniques can be employed in packetization to significantly reduce consecutive block loss [10]. Since various error concealment techniques employ quite different “philosophies,” it was not conclusive from the literature which one is the best. We attempt to address this issue through a benchmarking effort, which also sheds light on the design direction of a new concealment framework that can outperform the existing approaches.

We use a collection of fifteen 8-bit gray-scaled images with different characteristics to evaluate the performance of the six approaches reviewed above, namely, MDI, POCS, MSR, GSB, OASI, and LRC. The names and the corresponding references for these approaches are listed in Table 1. The collection of the 15 images is shown in the upper part of Figure 11. They can be divided into roughly four categories according to the visual content, namely, portraits, artificial images, natural scenery images, and rich texture images. We test the concealment on a typical loss pattern as shown in Figure 1, where a total of 25% blocks are lost in a checkerboard fashion and the block size is 8×8 . This damage pattern is used in all following experiments if not specified

TABLE 2: Comparison of algorithms in concealment quality PSNR (dB). For each image, the scheme achieving the best performance is highlighted in bold font. The Better-2 column lists the concealment quality of the recovered images in which each concealed block is the better one selected between GSB and OASI.

Type	Name	Size	MDI	POCS	MSR	LRC	GSB	OASI	Better-2
Natural	Bassharbor	512×512	29.47	28.12	28.83	27.84	30.69	30.37	31.46
	Blueflower	512×512	27.88	27.55	27.09	26.77	29.68	29.85	31.04
	House	512×512	28.78	26.08	27.00	26.86	29.47	30.00	30.98
	NewYork	512×512	24.25	21.00	23.66	22.80	24.13	24.52	25.29
	Operahouse	512×512	30.91	28.88	28.53	29.08	30.88	31.30	32.38
	Papermachine	512×512	29.77	28.46	25.80	31.78	33.85	33.75	36.12
	Watch	512×512	31.40	29.59	29.41	31.35	33.77	33.99	35.52
Portrait	Lena	512×512	32.28	29.49	29.20	30.64	34.43	35.12	36.08
	Barbara	512×512	27.41	23.35	27.14	29.78	29.26	30.79	31.80
	Kid	480×480	31.86	29.62	29.57	30.21	33.47	33.45	34.98
	Man	512×512	27.59	25.41	26.07	25.60	28.77	29.13	30.12
Artificial	Circletrain	512×512	41.62	34.16	32.11	46.51	48.33	34.90	48.33
	Tulip	512×512	29.74	28.05	26.71	27.61	33.22	33.47	35.13
	Waterfall	512×512	27.92	26.36	26.52	26.18	28.79	29.12	30.20
Texture	Bear	384×384	30.05	29.55	27.99	27.82	32.33	33.30	34.38




FIGURE 1: A checkerboard pattern with 25% block loss used in the concealment experiments.

otherwise. We examine the quality of recovered images in terms of PSNR and the computation complexity in terms of the concealment speed, and summarize the results in Tables 2 and 3, respectively. All algorithms have been implemented in C/C++ with a moderate amount of optimization and the same speed-up settings, and tested on a 1.20 GHz Pentium-4 PC with 256 MB RAM.

We can see from Table 2 that among the three recent techniques reviewed earlier, the LRC approach does not outperform the GSB and OASI approaches on most images. One reason is that the checkerboard error pattern leaves a very limited number of the candidate matching windows that do not suffer from the loss. The LRC approach does not perform well on most natural scenery images either, since there are few repeating patterns. On the other hand, the GSB and

TABLE 3: Comparison of algorithms in speed (seconds) for concealing the “Lena” image using a 1.20 GHz Pentium-4 PC with 256 MB RAM.

	MDI	POCS	MSR	LRC	GSB	OASI
Lena	3.03	219.58	0.59	98.45	0.56	7.12

OASI approaches significantly outperform other approaches on these benchmark images, although neither of the two gives the best performance for all images. The lack of all-time champion suggests that the image characteristics vary significantly from one to another, so a single algorithm based on an assumption about one aspect of the characteristics is not suitable for all images. This motivates us to go one step further and assemble a recovered image in which each concealed block is the better one selected between the GSB and OASI concealment results. As shown in the last column (“Better-2”) of Table 2, this assembled image gives a much higher overall concealment quality than using GSB or OASI alone.

In terms of computation complexity measured in concealment speed, Table 3 shows that MSR and GSB are the fastest. MDI and OASI are about an order of magnitude slower, and LRC and POCS are by far the slowest algorithms. Jointly considering the concealment quality and speed, we see that although GSB and OASI both have high performance on concealment quality, OASI has relatively high computation complexity. If we could choose the OASI method to conceal corrupted blocks only when it provides significant performance gain, we would achieve both higher concealment quality and relatively lower computation complexity. This motivates us to research on an adaptive scheme for selecting error concealment methods to combine the advantages of these two top performing schemes.



FIGURE 2: Illustration of better performing concealment scheme between GSB and OASI on the “Lena” image: (white blocks) OASI performs better; (black blocks) GSB performs better; (gray blocks) GSB and OASI do not have significant performance difference.

2.3. Classification-based concealment

For a receiver to pick the better one between the two state-of-the-art techniques correctly is a nontrivial task. This is because a receiver does not have the original undamaged image to compare with and determine which scheme gives better performance. Available to a concealment system are only the survived pixels that surround each corrupted block. If we could establish the connection between the image characteristics of the survived surrounding pixels and the correct selection between GSB and OASI using a training set, we could make a smart decision on which scheme to choose for a new damaged image.

To help exploring a rule in classifying the survived surrounding pixels, we take a close look at the “Better-2” test from Table 2. For each block, we quantify the error concealment performance of GSB and OASI by

$$\begin{aligned} P1 &= \sum_{i=1}^K |C1_i - O_i|, \\ P2 &= \sum_{i=1}^K |C2_i - O_i|, \end{aligned} \quad (1)$$

where K is the number of pixels in the block and is 64 in our case; O_i is the original value of the i th pixel in the block; $C1_i$ and $C2_i$ are the corresponding recovered pixel values by GSB and OASI, respectively. We visualize in Figure 2 the scheme selection for each lost block of the “Lena” image. The gray blocks indicate that GSB and OASI do not have significant performance difference (i.e., $|P1 - P2| < 96$); the white blocks indicate that $P2$ is much smaller for the corresponding blocks; and the black blocks indicate that $P1$ is much smaller. From Figure 2, we do not observe any obvious trend in determining where GSB and OASI would perform better: the black blocks appear in both edges and some texture areas and so do the white blocks.

We further explore if one could deduce some simple rules from the spatial characteristics of survived pixels surrounding the lost blocks. We define a smoothness feature from

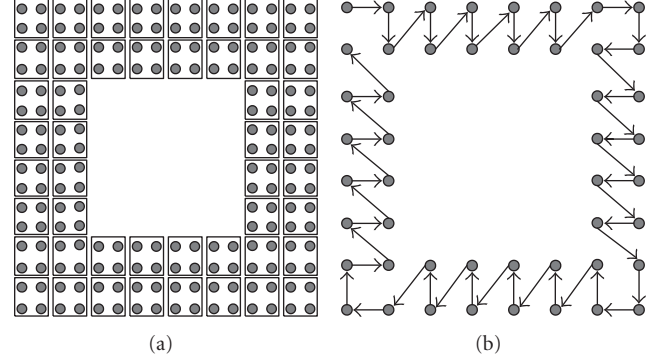


FIGURE 3: Feature extraction from survived surrounding pixels: (a) grouping of survived pixels into small 2×2 segments, (b) scanning order for constructing a feature vector.

four layers of survived surrounding pixels as follows. First, we group the pixels into a total of 48 segments, and each segment has 2×2 pixels, as shown in Figure 3(a). For each segment, we generate a binary value characterizing smoothness: if the range of the pixel intensity in the segment exceeds a predetermined threshold of 15, we use “1” to indicate it as a nonflat segment; otherwise, we use “0.” Next, the binary values from different segments are scanned according to the order in Figure 3(b) to form a feature vector, which is a binary sequence. We count the total number of 1s in the feature vector (i.e., the number of nonflat segments) for each of the 15 images used in our benchmark test. For each possible count of nonflat segments, we also compute the ratio of the number of blocks where OASI performs better versus those where GSB performs better. The relation is visualized in Figure 4, where we can see a general trend that GSB is likely to perform better on smooth blocks, and OASI tends to be better for texture blocks. But the curve is not monotonic and the ratios do not deviate much from one, suggesting that we cannot reliably determine the better performing concealment scheme just based on the nonflat segment count of the surviving surroundings.

The difficulty for a receiver in arriving at a simple rule to determine the better performing scheme can be tackled in two ways. If a decision is to be made solely on the receiver side, there is a need of employing advanced classification tools to group all possible surrounding pixel patterns into two classes, one class favoring the use of OASI for concealment and the other class favoring GSB. Alternatively, we can avoid the difficult task of receiver-side classification by determining the classification information on the sender side where the uncorrupted image is available for providing ground truth, and by sending such extra information to the receiver through attachment or data embedding techniques. In the next two sections, we will present the details of the proposed receiver-side and sender-side schemes, respectively. While we use OASI and GSB as building blocks to investigate our proposed framework of classification-based concealment, the new framework is general so that it can

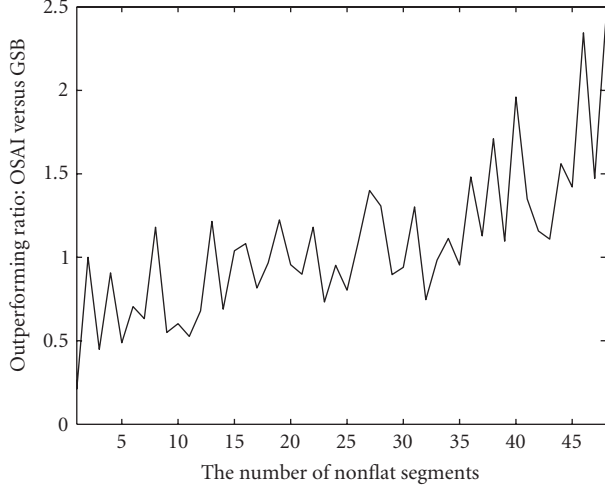


FIGURE 4: Examining the feasibility of a simple smoothness measure for distinguishing the better performing scheme: the x -axis represents the number of nonflat segments in survived surroundings and the y -axis represents the ratio of the block counts where OSAI performs better to those where GSB is better.

be easily extended to incorporate other appropriate concealment schemes and perceptual criteria.

3. RECEIVER-SIDE ADAPTIVE BLOCK CONCEALMENT USING SVM CLASSIFICATION

3.1. Classification based on support vector machine

We formulate a receiver's choice of concealment scheme for each block as a supervised classification problem. Each error concealment method is considered as a class, and a feature vector is extracted from the pixels that surround an image block. In the training stage, we collect a number of feature vectors from training images, and label every feature vector \mathbf{x}_i with a ground truth class corresponding to the best concealment method for the associated block. We train the classifier using these feature-class pairs.

We adopt support vector machine (SVM) classifiers, as they often exhibit good generalization performance [13, 14] with theoretical insights of structural risk minimization [15, 16]. The design of an SVM classifier can be boiled down to a convex quadratic programming problem with global optimal solutions in training. For our two-class pattern classification problem that decides between the GSB and OSAI concealment approaches, two kernel functions have been used to search for the optimal classification solution, namely, a linear kernel function and a radial kernel function.

3.1.1. Linear SVM

The linear SVM determines a linear discriminant function (a hyperplane) that gives the maximum separation margin between the two classes of training data [15]. The optimization

problem can be formulated as

$$\begin{aligned} & \text{minimize } f(\mathbf{w}, b) = \|\mathbf{w}\|^2, \\ & \text{subject to } y_i(\mathbf{x}_i^T \mathbf{w} + b) - 1 \geq 0, \end{aligned} \quad (2)$$

where \mathbf{x}_i is the i th training feature vector and $y_i \in \{-1, 1\}$ represents the corresponding class label. The separating hyperplane is parameterized by a vector \mathbf{w} and a scalar b , where \mathbf{w} is the norm of the separating hyperplane. The Lagrangian multiplier formulation for this constrained optimization problem is

$$L_p = \frac{1}{2} \|\mathbf{w}\|^2 - \sum_{i=1}^l \alpha_i y_i (\mathbf{x}_i^T \mathbf{w} + b) + \sum_{i=1}^l \alpha_i, \quad (3)$$

where $\{\alpha_i\}$ is a set of Lagrangian multipliers. Now, the problem is reduced to minimizing L_p with respect to \mathbf{w} and b under the following restrictions: (i) the derivatives of L_p with respect to all α_i 's vanish and (ii) $\alpha_i \geq 0$. For this convex quadratic programming problem, it is well established that the solution can be obtained through the Karush-Kuhn-Tucker (KKT) conditions or through an easier *dual* problem [15].

When the training data of the two classes is linearly separable, the linear kernel SVM approach gives a classifier in the form of a hyperplane separating the two classes of training data with the largest margin. If the training data is not linearly separable, a positive slack variable ξ_i ($\xi_i \geq 0$) can be introduced to alleviate the sensitivity of noisy training patterns [17]:

$$y_i(\mathbf{x}_i^T \mathbf{w} + b) - 1 + \xi_i \geq 0, \quad (4)$$

$$L_p = \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{i=1}^l \xi_i - \sum_{i=1}^l \alpha_i [y_i(\mathbf{x}_i^T \mathbf{w} + b) - 1 + \xi_i] - \sum_{i=1}^l u_i \xi_i, \quad (5)$$

where C is a parameter adjusting the relative penalty given to the classification errors on the training data.

To use a trained classifier to classify a new test sample \mathbf{z} , we evaluate the sign of the following function:

$$f(\mathbf{z}) = \mathbf{w}^T \mathbf{z} + b = \sum_{i=1}^{N_s} \alpha_i y_i \mathbf{x}_i^T \mathbf{z} + b. \quad (6)$$

Here, \mathbf{w} is explicitly determined by a set of N_s *support vectors*, which are such training vectors that lie closest to the hyperplane separating the two classes [15]. The sign reflects on which side of the decision boundary that \mathbf{z} lies and thus determines the classification result.

3.1.2. Handling nonlinearity

The feature vector as an input to a classifier for the concealment problem can be the pixel pattern surrounding a lost block, or some statistics generated from the pattern (such as the binary feature vector defined in Section 2). The training features for each class may have complicated distributions,




FIGURE 5: Handling the nonlinearity by a divide-and-conquer technique that trains a set of classifiers, one for each subset of the feature space.

and in general are far from separable by a linear discrimination function in the original vector space. The nonseparability by a linear discrimination function can be handled in two ways. One is to extend the linear SVM with the kernel technique and the other is to divide the vector space into groups and find one classifier for each group.

Nonlinear classification functions [15] can be built by replacing the dot product term $\langle \mathbf{x}_i, \mathbf{x}_j \rangle = \mathbf{x}_i^T \mathbf{x}_j$ in the linear kernel SVM by an appropriate kernel function $K(\mathbf{x}_i, \mathbf{x}_j)$. This is equivalent to transforming feature vectors to a higher-dimensional space H through a mapping $\Phi: R^d \rightarrow H$, and then finding a linear SVM classifier in this new space with $K(\mathbf{x}_i, \mathbf{x}_j) = \langle \Phi(\mathbf{x}_i), \Phi(\mathbf{x}_j) \rangle$. The radial basis kernel function in the form of

$$K(\mathbf{x}_i, \mathbf{x}_j) = e^{-\|\mathbf{x}_i - \mathbf{x}_j\|^2 / 2\sigma^2} \quad (7)$$

is commonly used for its good generalization capabilities, especially when very limited information is available about the data distribution and separability for all classes. Here, σ is the width of the radial basis. It affects the classification performance substantially and will be addressed later in this section.

An alternative way of dealing with the nonlinearity is to use a divide-and-conquer technique. The idea is illustrated by the two-dimensional example shown in Figure 5, where the two classes of data represented in Figure 5(a) are not linearly separable. However, if we divide the space into four stripes as shown by the dashed lines in Figure 5(b), the data within each stripe becomes more separable by a linear function. The subdivision of the feature space naturally accommodates the nonlinearity in the class boundary, yet the training process is comprised of training a set of relatively simple linear SVMs. Subdividing the feature space into nonoverlapped subsets can be done through dividing

the dynamic range of some feature elements or according to the norm of the feature vector. The latter reflects the overall smoothness of the surrounding pattern for the feature vector defined in Section 2, as the L_1 norm of the vector gives the total number of nonflat 2×2 segments over the 48 pixel segments surrounding a lost block. Recalling the trend seen in Figure 4 on the classes as a function of the overall smoothness, the subdivision allows us to naturally adapt to the changing characteristics.

The nonlinearity in the classification can also be handled using a combination of the above two approaches. This hybrid approach divides the feature space into subsets and provides a nonlinear SVM (such as the radial kernel function) for each subset. It offers a great amount of flexibility, allowing the subsets to use different kernel parameters (such as σ in the radial basis function) or even different kernels. The nonlinear SVM obtained for each subset of feature space can have a much smaller number of support vectors; hence can be considerably simpler than a nonlinear SVM trained for the entire space. As such, the hybrid approach has a low computational complexity in both the training and test phases.

3.1.3. Determining kernel parameters

In practice, the relation between the classification accuracy on the training set and on test set relies highly on the generalization capability of the classifier. In SVMs, there are several important parameters affecting the generalization capability, such as C in (5) and σ in (7). Choosing SVM kernel parameters can be viewed as a validation process, and evaluating the performance of the trained model on a validation set is a general approach to select kernel parameters. Based on this approach, we propose the following preprocessing procedure for choosing the kernel parameters.

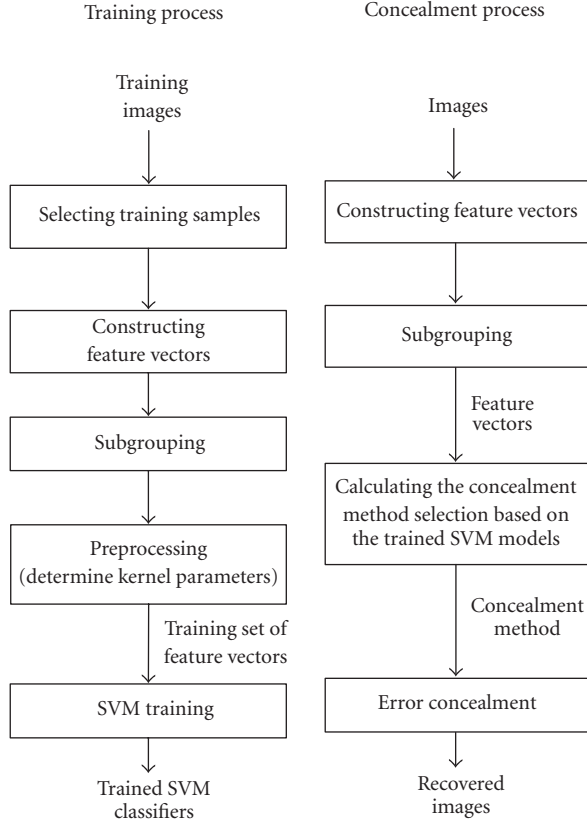


FIGURE 6: Block diagram of the proposed receiver-side classification-based concealment approach.

Step 1. Dividing the training samples into two subsets, \mathcal{A} and \mathcal{B} : in each iteration below, we use set \mathcal{A} for training and set \mathcal{B} for validation.

Step 2. Choosing kernel parameters and constructing a new training set \mathcal{R} : we adjust kernel parameters $\sigma^{(1)}$ and $C^{(1)}$ so that the sum of training errors on \mathcal{A} and validation errors on \mathcal{B} is minimized. More generally, we may employ an objective function using a weighted sum of the two types of errors, and low error rate on the validation set is often desirable to ensure a good generalization capability of the classifier. Since SVM is known to generalize well and does not usually suffer from overfitting problem as much as the conventional classifiers do, we choose to minimize the sum of errors (i.e., with equal weights) for simplicity. A new training set \mathcal{R} is then generated consisting of the support vectors from set \mathcal{A} and the successfully classified samples from set \mathcal{B} .

Step 3. Switching subsets: we switch set \mathcal{A} with set \mathcal{B} and repeat Step 2. We record the kernel parameters as $\sigma^{(2)}$ and $C^{(2)}$ and denote the new training set as \mathcal{S} . The union of set \mathcal{R} and set \mathcal{S} becomes the final training set \mathcal{T} .

Step 4. Determining kernel parameters: the kernel parameters obtained from the two iterations above provide a search

range for determining the final parameters. For example, $\sigma^{(1)}$ and $\sigma^{(2)}$ specify a range over which we will search for the final value of σ that can minimize the training error on set \mathcal{T} . Other kernel parameters can be jointly determined through the search.

In addition to determining kernel parameters, we also filter out the samples that have very similar values but different class labels. These samples are usually located in such region of the feature space that is difficult to classify and they can make the classification boundary very complex. Removing them from the training set helps improve the generalization capability of the classifier.

3.2. Overall algorithm

The overall algorithm of our proposed receiver-side classification-based block concealment is summarized in Figure 6. Below we explain a few additional details of the training and concealment processes.

3.2.1. Selection of training data

We choose a set of training images that represent a variety of characteristics. Because of the spatial correlation in most natural images, we use about one fourth of blocks in the checkerboard pattern from each training image as candidates to form a training set. As discussed earlier, we further filter out the blocks where different concealment schemes do not give substantially different performance.

3.2.2. Construction of feature vectors

Since different spatial block concealment techniques may use different sets of surrounding pixels, the feature vectors derived for classification should come from the union of the sets of pixels used by these techniques. For example, GSB often uses two surrounding layers to extract the geometric structure information, while OASI uses four surrounding layers to compute the interpolation coefficients. The classification region should therefore include four surrounding layers of pixels. For block size of 8×8 , 192 pixels are involved in classification.

While pixels can be used directly as features, they often require a sophisticated kernel function to ensure separability and thus incur high computation complexity. We generate a more compact feature vector from pixel values using a similar approach as described in Section 2.3 and summarized as follows. We first partition the four surrounding layers of pixels into segments, as illustrated in Figure 3(a). For the i th segment of four pixels, the feature value v_i characterizes the smoothness of the segment and is computed as

$$v_i = \text{floor} \left[\left(\max \{p_k\} - \min \{p_k\} - s \right) / Q_v \right] + 1, \quad (8)$$

where $\{p_k\}$ are the pixels in the i th segment, the floor function returns the largest integer less than or equal to the input. The two parameters s and Q_v control the sensitivity of the feature. We choose $s = 15$ and $Q_v = 50$ based on our experimental results. We then put these feature values into

TABLE 4: Overall classification accuracy on the 13 test images.

	1 group	16 subgrouping	48 subgrouping	48 subgrouping with preprocessing
Linear SVM	50.55%	65.96%	66.26%	67.11%
Radial SVM	65.54%	66.75%	67.17%	70.16%

a vector. The ordering of features in the feature vector does not affect the performance of a trained SVM classifier since the kernel functions widely used in SVM classification are invariant with respect to the ordering of features.

3.2.3. Subgrouping

As discussed earlier, to handle the nonlinearity of the class boundary, we divide the feature space into n subsets and train an SVM classifier for each subset. We use a simple empirical partitioning rule based on the number of nonzero values in a feature vector.

3.2.4. Preprocessing of training samples

The feature vectors we used for training are divided into sets \mathcal{A} and \mathcal{B} . Each set includes images from all four representative categories mentioned before, namely, portraits, artificial images, natural scenery images, and rich texture images. We determine in this step the kernel parameters and training set using the approaches described in Section 3.1.3.

3.2.5. Concealment process

After the training process is performed off-line, the parameters of trained SVM classifiers are stored in the receiver system. To conceal a corrupted image block, the receiver system use the same approach as in the training process to construct feature vector and identify to which subgroup the feature vector belongs. The classification result will then determine which concealment scheme to use.

3.3. Experimental results and performance analysis

In this section, we present the experimental results on the proposed block concealment method using receiver-side classification. We use the SVM^{light} toolkit [18] to accomplish this classification task. SVM^{light} is an implementation of SVM based on the optimization algorithm in [19].

A total of 15 images are used for training and 13 for testing, which are shown in Figure 11. There are a total of 5 562 blocks in the training images and 3 804 blocks in the test images having substantially different concealment performance by GSB and OASI. These blocks are used to evaluate the classification accuracy.

We first train a linear SVM using the 48-dimension feature vectors of all training blocks. The classification accuracy of this trained linear SVM on the test blocks is only 50.55%. The failure of this classification experiment indicates the high nonlinearity in the boundary of the two classes.

We then examine the effects of various approaches in handling the nonlinearity. The simulation results of this exploration are shown in the first row of Table 4. We compare the cases of no subgrouping, 16-group subgrouping, and 48-group subgrouping. For these three cases, the kernel parameters are chosen that can provide the highest classification accuracy on three of the training images, “Lena,” “Barbara,” and “Bassharbor.” We also consider the case of applying preprocessing with 48-group subgrouping for thorough selection of kernel parameters and filter out noisy samples, using the approaches described in Section 3.1.3. As shown in the table, subgrouping significantly improves the classification accuracy by more than 15%; and preprocessing and finer subgrouping can further improve the classification accuracy.

Based on results from the above exploration, we adopt 48 subgroups with preprocessing procedure for our training process and examine the concealment performance of the proposed receiver-side classification-based scheme on the thirteen 8-bit gray-scaled test images. The classification accuracy for each subgroup ranges from 58.82% to 83.09%, and the overall classification accuracy is 67.11%. From the comparison of concealment results with that of GSB [5] and OASI [6] in Table 5, we can see that the classification-based method with a linear kernel has up to 0.84 dB gain when compared to the GSB method and up to 1.06 dB gain when compared to the OASI method.

We then train a radial basis kernel SVM to evaluate how well it handles the nonlinearity of training data. The preprocessing and subgrouping are also evaluated for this nonlinear kernel. As with the linear kernel, the radial basis kernel can also benefit from the preprocessing and finer subgrouping for improving the classification accuracy, although the improvement due to grouping is less significant on the radial basis kernel than on the linear kernel. This latter aspect is expected as the radial basis kernel has a good capability of handling the nonlinear classification boundary even without subgrouping. The classification accuracy for each group ranges from 60.00% to 80.53%, and the overall classification accuracy is 70.16%. As shown in Table 5, the classification-based method using the radial basis kernel SVM has up to 0.94 dB gain compared to the GSB method and up to 1.26 dB gain when compared to the OASI method. The proposed scheme consistently outperforms the two prior algorithms on all test images. As an example, we show a portion of the “Nickel” image in Figure 7, and we can see that the proposed concealment scheme provides better visual quality and leaves fewer artifacts.

It is worth noting that a radial basis kernel gives about 3% higher classification accuracy than a linear kernel, under the same 48-group subgrouping and preprocessing procedure.

TABLE 5: Comparison of concealment quality in PSNR (dB) of existing concealment schemes and the proposed receiver-side classification-based approaches.

Type	Name	Size	GSB	OASI	Better-2	Linear kernel	Radial kernel
Natural	Fishingboat	512×512	30.93	31.10	32.28	31.36	31.64
	Goldhill	512×512	32.35	32.41	33.52	32.63	32.84
	Peppers	512×512	35.18	35.55	36.72	36.02	35.79
	Skylinearch	400×400	32.01	31.34	33.22	32.40	32.60
	Lochness	512×512	32.74	32.33	33.40	32.78	32.78
	Bellflower	512×512	33.27	33.70	35.57	34.12	34.21
	Brandyrose	512×512	39.47	39.27	40.42	39.86	39.80
	Lake	512×512	28.54	28.73	30.14	29.10	29.04
Portrait	F14	496×496	38.64	38.86	39.88	38.75	39.05
	Elaine	512×512	35.17	35.93	36.35	35.85	35.96
	Couple	512×512	30.74	31.06	32.22	31.49	31.43
Artificial	Nickel	256×256	29.05	28.55	30.53	29.33	29.58
Texture	Baboon	512×512	26.11	26.48	27.12	26.62	26.62

The small improvement in classification accuracy, however, does not always translate into the improvement of concealment quality. For example, we can see from Table 5 that radial basis kernel provides slightly better concealment for some test images, while linear kernel is better for others. This is because the set of accurately classified blocks may be different by the two kernel techniques, and the quality gain on the slightly bigger set of accurately classified blocks may not always offset the quality loss on the falsely classified ones. On the other hand, we see that the classification-based schemes give consistently higher concealment quality than the two current state-of-the-art algorithms. With more accurate classification, the concealment quality can be further improved. Along the line of seeking more accurate classification information, we are inspired by the growing importance of involving both sender and receiver in efficient and reliable multimedia communications. In the next section, we investigate what role the sender system can play in facilitating classification-based concealment.

4. BLOCK CONCEALMENT WITH SENDER-SUPPLIED CLASSIFICATION INFORMATION

The receiver-side classification algorithm proposed in Section 3 outperforms the conventional error concealment approaches. Coming with such benefit is the increase in computation complexity at receiver-side for performing classification. The increased complexity may pose a challenge for systems that have very limited computation resources and/or stringent real-time rendering constraints. If some parts of the concealment task could be moved to the sender side, it would help reduce the computation burden on the receiver side, as demonstrated in several recent works [9, 10].

An important benefit of moving the classification task from a receiver to a sender is that it allows for an easy access of the perfect classification information. This is because the sender has full reference to the original, uncorrupted image,

and can compare the concealment quality by various techniques to obtain the ground truth about which technique works better. The higher accuracy of the classification information can further improve the overall concealment quality upon what we have achieved in Section 3, which is an even more attractive advantage than the reduced receiver-side computation complexity.

In this section, we extend the classification-based concealment framework from a sender-driven perspective to design and evaluate error concealment schemes with sender-supplied classification information. We will examine two main approaches to conveying the classification information from a sender to a receiver: one is to attach the side information in the header and the other is to embed the side information in the image signal using data hiding technique.

4.1. Conveying classification information by attachment

A quite straightforward way to convey the classification information from the sender to the receiver is to transmit the information along with the image, for example, in the image header. The side information requires extra bandwidth, and therefore, the appropriateness of the attachment approach depends on the application and the image/video size. An alternative approach to avoid the increase in bandwidth is to encode the image at a lower rate to spare room for side information. This would reduce the image quality, leading to a similar tradeoff as in the data embedding approach to be discussed in the next subsection.

We present the system block diagram of the sender-side attachment scheme in Figure 8. On the sender side, in addition to encoding an image as usual, the system would perform the following tasks:

- (1) perform error concealment on each block or on selected blocks using multiple error concealment methods;




FIGURE 7: Visual quality comparison of three concealment schemes: (a) original image; (b) corrupted image; (c) recovered image using GSB; (d) recovered image using OASI; and (e) recovered image using the proposed classification-based method.

- (2) compare the quality of the images obtained by these concealment methods and classify each block according to the winning technique;
- (3) encode the classification information for each block, possibly using lossless compression techniques;
- (4) attach the classification information to the compressed image bit stream.

On the receiver side, upon detecting the corrupted blocks, the receiver will extract the classification information from the received stream and use this side information to select the appropriate method for concealing each corrupted block. We can further apply forward error correction coding with appropriate strengths to protect the image stream and the side information.

Regarding the detailed encoding method for side information, we denote the side information for the GSB concealment method as “0” and that for OASI as “1.” The side information for all blocks can be put together as a binary sequence. Recall that GSB concealment has lower computation complexity than OASI. So as before, we choose the error concealment technique with lower computation complexity for the blocks where the performance of the two concealment methods are not significantly different. This also helps give a long run of “0” in the side-information encoding. We then apply run-length coding and arithmetic coding to compress the binary sequence of classification information.

It can be seen that the attachment scheme trades additional bandwidth for improved concealment quality. The tradeoff can be adjusted as follows. For each block, the performance of each algorithm ($P1$ and $P2$) is calculated according to (1). The binary-valued side information L for the block is determined by

$$L = \begin{cases} 1, & \text{if } P1 - P2 > \Delta_{th}, \\ 0, & \text{otherwise,} \end{cases} \quad (9)$$

where Δ_{th} is a threshold. An experiment with different settings of Δ_{th} is performed on the JPEG-compressed “Lena” image with quality factor $Q = 80\%$, where the image size is 512×512 and the JPEG file size is 303 072 bits. As shown in Figure 9, the larger Δ_{th} we choose, the lower PSNR we get. On the other hand, since more blocks are labeled as “0” with a larger Δ_{th} , compressing the classification information using run-length coding and arithmetic coding will achieve a higher compression ratio. The results in Figure 9 shows that when Δ_{th} is around 96, the gain in error concealment quality is significant, yet the additional bandwidth for classification side information is quite moderate and only about one percent of the image file size. Thus we use this value to evaluate the overall concealment quality.

The simulation results of the attachment scheme are listed in Table 6. The results suggest that our proposed concealment scheme by attaching classification information outperforms each individual receiver-side concealment approach. The error concealment quality can be improved by about 1 ~ 2 dB when compared to the better one between the two individual methods. Readers may notice that the attachment scheme has 0 dB gain on the “Circletrain” image when compared to GSB. As shown in Figure 11, this artificial image has uniform background and smooth edges. GSB gives better concealment quality in terms of PSNR for every recovered blocks, so we cannot get any improvement compared to GSB.

4.2. Conveying classification information by embedding

Although the attachment scheme has excellent performance, the additional bandwidth for side information may not be available or too pricey in some systems. Recoding the image part to a slightly lower rate requires a nontrivial amount of computation complexity to ensure that the total bandwidth

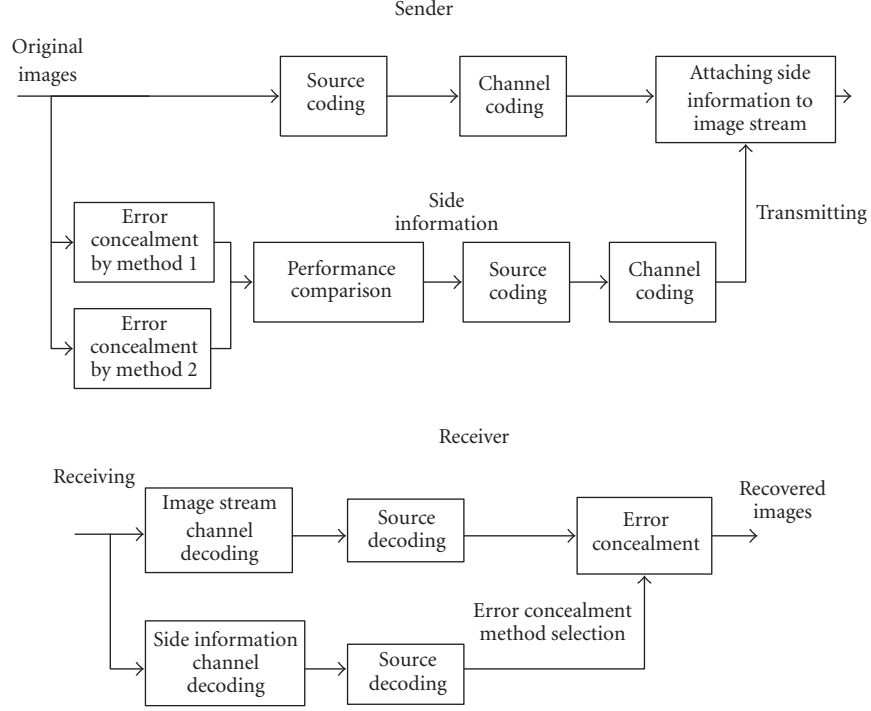




FIGURE 8: Block diagram of the sender-side attachment approach.



(a)



(b)

FIGURE 9: Relation of the threshold Δ_{th} versus the concealment quality and the bandwidth required for side information, respectively, when applying the sender-side attachment approach on the “Lena” image.

of the image plus the side information is unchanged. A viable alternative to convey side information with little additional bandwidth is embedding it in the image. More specifically, we embed 1-bit classification information of a block into its neighboring block. The embedding will be incorporated

in the visual communication system along with interleaved packetization mentioned at the beginning of the paper, so that the neighboring blocks are packed into different packets. In such a way, it is unlikely for a block and its neighbor holding its classification information to be corrupted

TABLE 6: Performance evaluation of the sender-side attachment approach.

Image type	Image name	JPEG file size (bytes)	Side-information size (bytes)	Quality gain over GSB (dB)	Quality gain over OASI (dB)
Natural	Bassharbor	50 867	368	0.52	1.14
	Blueflower	53 528	495	0.87	0.90
	House	46 975	361	1.28	1.26
	Newyork	73 830	436	0.89	0.67
	Operahouse	48 666	365	1.09	0.99
	Papermachine	41 773	285	1.95	2.07
	Watch	41 773	293	1.23	1.09
Portrait	Lena	37 884	287	0.99	0.93
	Barbara	50 867	424	2.21	1.12
	Kid	30 791	257	1.12	1.08
	Man	61 810	431	0.80	0.79
Artificial	Circletrain	15 709	124	0	11.19
	Tulip	48 641	437	1.45	1.68
	Waterfall	44 734	292	0.93	0.75
Texture	Bear	26 089	280	1.32	1.12

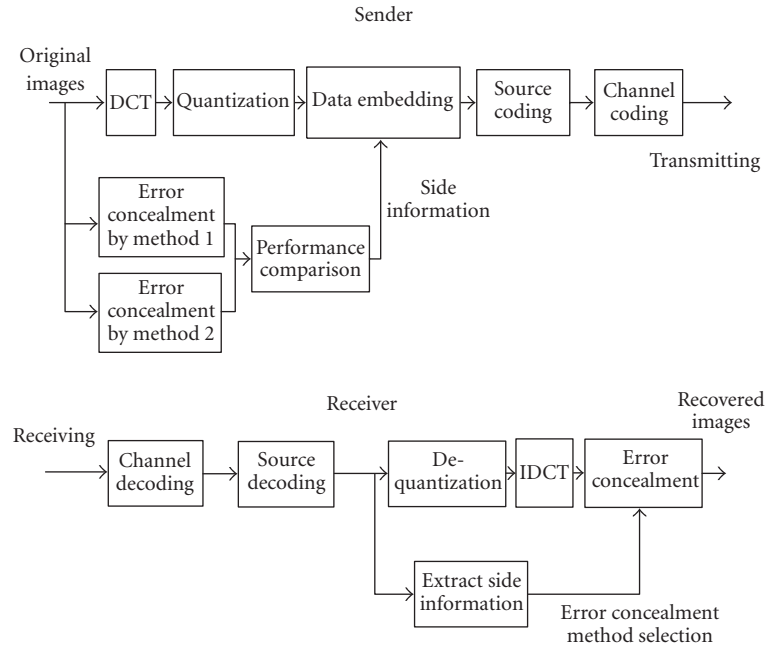


FIGURE 10: Block diagram of the sender-side embedding approach.

simultaneously. As we will see later in this subsection, the embedding in the neighboring block has additional advantage when dealing with smooth blocks. We summarize the system block diagram in Figure 10 and explain a few details of embedding below.

As can be seen from the previous subsection, the amount of classification information is on the order of a couple of thousand bits, which calls upon an embedding technique with quite high embedding rate. Unlike many copyright

protection applications, there is no major adversary to circumvent the embedded data in error concealment application, where the side information helps improve the performance of image communications [20]. The quantization-based data embedding is a viable choice to meet these requirements [21].

We use a simple version of quantization embedding, known as the *odd-even embedding* technique [22], to embed the classification information in an image. To avoid a



Bassharbor (512 × 512)



Blueflower (512 × 512)



House (512 × 512)



Newyork (512 × 512)




Operahouse (512 × 512)



Papermachine (512 × 512)



Watch (512 × 512)



Lena (512 × 512)




Barbara (512 × 512)




Kid (480 × 480)



Man (512 × 512)



Circletrain (512 × 512)



Tulip (512 × 512)



Waterfall (512 × 512)



Bear (384 × 384)



Fishingboat (512 × 512)



Goldhill (512 × 512)



Peppers (512 × 512)



Skylinearch (400 × 400)



Lochness (512 × 512)



Bellflower (512 × 512)



Brandyrose (512 × 512)



Lake (512 × 512)



F 14 (496 × 496)




Elaine (512 × 512)



Couple (512 × 512)



Nickel (256 × 256)



Baboon (512 × 512)

FIGURE 11: The first fifteen 8-bit gray-scaled images are used for training and the next thirteen 8-bit gray-scaled images are used for testing in the classification-based concealment. The image sizes are listed in parentheses after the image names.

TABLE 7: Performance evaluation of the sender-side embedding approach. Images are in JPEG format with quality factor $Q = 80\%$.

Image type	Image name	PSNR of image after embedding (dB)	Concealment gain over GSB (dB)	Concealment gain over OASI (dB)
Natural	Bassharbor	41.89	0.14	0.76
	Blueflower	41.73	0.77	0.80
	House	42.01	1.00	0.98
	Newyork	38.25	0.74	0.52
	Operahouse	40.57	0.63	0.53
	Papermachine	42.42	1.02	1.14
	Watch	42.82	0.74	0.60
Portrait	Lena	43.21	0.30	0.24
	Barbara	42.25	1.94	0.85
	Kid	43.16	0.60	0.56
	Man	39.48	0.49	0.48
Artificial	Circletrain	47.36	-2.80	8.39
	Tulip	42.31	0.78	1.01
	Waterfall	40.47	0.62	0.44
Texture	Bear	43.91	0.63	0.43

substantial impact on compression size and visual quality of the image, the classification information for each block is embedded into the last quantized nonzero DCT coefficient in a zigzag scan order. The coefficient is forced to be an even value if we want to embed “0,” or an odd value if to embed “1,” and the embedding tries to make minimum necessary changes to enforce such a relation. If all the quantized AC coefficients in a block are zero, which we would encounter for smooth blocks, we will not make any changes on the coefficients. In this case, the receiver would consider a “0” to be embedded in the block based on the above-mentioned rules, and apply the concealment technique of lower computation complexity (i.e., GSB) for the corrupted block. Such an arrangement works well in practice. This is because GSB usually performs better for blocks with relatively “flat” surrounding; in the mean time, the characteristics of nearby blocks are likely to be similar and can be fully exploited by neighborhood embedding presented earlier, where classification information is embedded into neighboring block.

The experimental results of the embedding scheme are shown in Table 7. The improvement of concealment quality on most images is significant: we have a 0.14 ~ 1.94 dB gain compared to GSB and 0.24 ~ 1.14 dB gain compared to OASI. For most images, GSB performs better on some blocks and OASI performs better on some other blocks. As such, the quality degradation introduced by the embedding procedure is overcome by the substantial concealment gain compared to either GSB or OASI alone. An interesting exception appears on the “Circletrain” image. Different from other images, GSB is the better selection for all blocks in the “Circletrain” image and the concealed quality is very high (with the PSNR dB value being in high forty). The sender-supplied classification information thus provides no gain when compared to using GSB alone. On the other hand, the embedding technique inevitably introduces a moderate amount of

quality degradation. As a result, for the “Circletrain” image, the embedding scheme achieves a net loss of 2.8 dB in PSNR compared to GSB, although little visual difference could be visible at such high PSNR levels. In comparison with OASI, the gain over OASI is over 8 dB and is much more noticeable.

5. COMPARISONS AND DISCUSSIONS

In the previous two sections, we have proposed three classification-based error concealment schemes to improve the concealment quality. Among the three schemes, one performs classification on the receiver side using an SVM classifier and features derived from the survived pixels surrounding a corrupted block, and the other two schemes convey the sender-supplied classification information to receiver by attachment and embedding, respectively. As we can see from Tables 5, 6, and 7, they all improve the concealment quality quite substantially. In this section, we compare the three schemes, discuss their advantages and shortcomings, and identify the application scenarios that each scheme is suitable for. We also discuss a few directions for further extension and generalization.

We first compare the quality of concealed images by these three schemes and show the results in Table 8. For each image, we use the uncorrupted JPEG compressed version with a quality factor of 80% as reference. Since the attachment scheme provides the ground truth of concealment technique selection to the receiver, it gives the highest concealment quality among the three schemes. The improvement over the individual concealment schemes is in the range of 0.5 ~ 1.5 dB. While the embedding scheme also provides the ground truth of most blocks to receiver (except for some very smooth blocks), its performance is lower than the attachment scheme by about 0.3 ~ 0.5 dB. This small loss is due to the distortion introduced by embedding, a price paid

TABLE 8: Comparison of concealment quality in PSNR (dB) by the receiver-side and sender-side approaches. Images are in JPEG format with quality factor $Q = 80\%$.

Image type	Image name	GSB	OASI	Receiver-side classification	Sender-side embedding	Sender-side attachment
Natural	Fishingboat	30.81	30.87	31.03	31.55	32.02
Portrait	Elaine	35.47	35.18	35.43	35.84	36.22
Artificial	Nickel	28.48	28.41	28.71	29.40	29.93
Texture	Baboon	26.02	26.19	26.25	26.45	26.74

for sending side information without additional bandwidth. The receiver-side classification scheme has the smallest improvement over individual scheme because the classification result at the receiver is not always accurate.

In addition to the visual quality of concealed image, other important issues include computational complexity, bandwidth usage, and complexity associated with overall system deployment. The receiver-side classification-based error concealment requires neither side information to be sent nor any special involvement of a sender. It can be therefore integrated in a standard-compliant coding system. The training involves a large amount of computation but can be performed off-line. A moderate amount of run-time computation power is required from the receiver to extract features and feed them into a trained SVM classifier to determine which concealment scheme to use, and this is done only for corrupted blocks. As the classification results are not always perfect and depend heavily on the generalization capability of the classifier, the concealment performance may vary substantially from one image to another. This scheme is suitable for applications where there is limited design flexibility on the sender side.

The schemes with sender-supplied classification information provide more *proactive* protection. They require a significant amount of computation power and cooperation on the sender side to perform concealment, provide ground truth on the concealment scheme to use for *every* block, and encode or embed the classification information with the image. The attachment scheme requires additional bandwidth to deliver the ground truth of classification. After such an attachment, the resulting media stream may not be standard-compliant. In contrast, the embedding scheme can maintain standard compliance of the resulting media stream. This is at an expense of minor reduction of the perceptual quality in the transmitted image, even when the transmission is free from error. On the other hand, the more accurate sender-supplied classification information provides substantial improvement in concealment quality and also eliminates the computation needed on the receiver side for classification. These schemes are suitable for applications with powerful sender and simple receiver and for scenarios where the visual data is encoded once but delivered and consumed by many users.

The spatial concealment schemes investigated in this paper can be used for both image and video transmissions. They can be applied to each corrupted video frame and can be used in conjunction with other temporal concealment

methods [23, 24]. The schemes that maintain standard-compliance of the transmitted video, such as the receiver-end classification and the embedding schemes, allow image/video to be handled by a number of existing visual communication systems that support the standard, with few additional changes to the system.

In addition to conveying side information to facilitate concealment, data embedding can also be used for detecting corrupted blocks [25]. For this error detection purpose on each block, the parity information or some known patterns should be embedded *inside* the corresponding block. The receiver will check the correctness of the parity or the integrity of the patterns to determine whether the block is corrupted. On the other hand, the side information of a block for facilitating its concealment must be stored *outside* that block, as seen in the algorithm presented in Section 4.

We have so far assumed that the block damage is isolated (i.e., all neighboring blocks of the damaged one are correctly received). Since consecutive block damage is a challenge to most error concealment techniques, interleaving techniques have been suggested in packetization to avoid packing neighboring blocks together [1, 10]. As such, consecutive block losses rarely happen at a moderate loss rate. In case when there remain some consecutive block losses, both GSB and OASI techniques have been demonstrated to handle a small number of consecutive blocks [5, 6]. The classification can also be extended to cope with this case, for example, to incorporate the loss of two horizontal or vertical neighboring blocks by training additional classifiers. And since what we have proposed is a general framework, it can be further extended to incorporate other concealment techniques and accommodate more than two candidate techniques.

6. CONCLUSIONS

In this paper, we present a new, classification-based spatial error concealment framework for visual communications. Our proposed framework takes advantages of state-of-the-art concealment techniques and adaptively selects the best suitable one for concealing each corrupted block. Under this new framework, we have proposed a receiver-side classification scheme to combine the sweet spots of several current state-of-the-art techniques, while maintaining standard compliance and requiring no special involvement on transmitter. We have also examined a sender-driven perspective to provide perfect classification information to a receiver through attachment or embedding, and thus further enhance

the error concealment performance. The advantages of each of the three proposed schemes have been analyzed and the suitable application scenarios suggested. Our experiments on a diverse set of images have shown that the proposed classification-based concealment framework provides up to 2 dB higher concealment quality over the current state-of-the-art algorithms.

ACKNOWLEDGMENTS

This research was supported in part by research grants from US National Science Foundation CCR-0133704 (CAREER). Preliminary results of this paper were presented in the IEEE International Conference on Image Processing (ICIP'03) [26]. The authors would like to thank Professor Wenjun Zeng of the University of Missouri, Columbia, for providing the source code of the GSB error concealment algorithm.

REFERENCES

- [1] Y. Wang and Q.-F. Zhu, "Error control and concealment for video communication: a review," *Proceedings of the IEEE*, vol. 86, no. 5, pp. 974–997, 1998.
- [2] W. Kwok and H. Sun, "Multi-directional interpolation for spatial error concealment," *IEEE Transactions on Consumer Electronics*, vol. 39, no. 3, pp. 455–460, 1993.
- [3] H. Sun and W. Kwok, "Concealment of damaged block transform coded images using projections onto convex sets," *IEEE Transactions on Image Processing*, vol. 4, no. 4, pp. 470–477, 1995.
- [4] Y. Wang, Q.-F. Zhu, and L. Shaw, "Maximally smooth image recovery in transform coding," *IEEE Transactions on Communications*, vol. 41, no. 10, pp. 1544–1551, 1993.
- [5] W. Zeng and B. Liu, "Geometric-structure-based error concealment with novel applications in block-based low-bit-rate coding," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 9, no. 4, pp. 648–665, 1999.
- [6] X. Li and M. T. Orchard, "Novel sequential error-concealment techniques using orientation adaptive interpolation," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 12, no. 10, pp. 857–864, 2002.
- [7] D. Zhang and Z. Wang, "Image information restoration based on long-range correlation," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 12, no. 5, pp. 331–341, 2002.
- [8] M. Bertalmio, L. Vese, G. Sapiro, and S. Osher, "Simultaneous structure and texture image inpainting," *IEEE Transactions on Image Processing*, vol. 12, no. 8, pp. 882–889, 2003.
- [9] P. Yin, B. Liu, and H. H. Yu, "Error concealment using data hiding," in *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP '01)*, vol. 3, pp. 1453–1456, Salt Lake City, Utah, USA, May 2001.
- [10] P. Yin, M. Wu, and B. Liu, "Robust error-resilient approach for MPEG video transmission over Internet," in *Visual Communications and Image Processing (VCIP '02)*, vol. 4671 of *Proceedings of SPIE*, pp. 103–111, San Jose, Calif, USA, January 2002.
- [11] Y. Liu and Y. Li, "Error concealment of digital images using data hiding," in *Proceeding of 9th IEEE DSP Workshop*, Hunt, Tex, USA, October 2000.
- [12] K.-H. Jung, J.-H. Chang, and C. W. Lee, "Error concealment technique using projection data for block-based image coding," in *Visual Communications and Image Processing (VCIP '94)*, vol. 2308, part 3 of *Proceedings of SPIE*, pp. 1466–1476, Chicago, Ill, USA, September 1994.
- [13] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, John Wiley & Sons, New York, NY, USA, 2nd edition, 2001.
- [14] T. Joachims, "Text categorization with support vector machines: learning with many relevant features," in *Proceedings of 10th European Conference on Machine Learning (ECML '98)*, pp. 137–142, Chemnitz, Germany, April 1998.
- [15] C. J. C. Burges, "A tutorial on support vector machines for pattern recognition," *Data Mining and Knowledge Discovery*, vol. 2, no. 2, pp. 121–167, 1998.
- [16] V. N. Vapnik, *The Nature of Statistical Learning Theory*, Springer, Berlin, Germany, 1995.
- [17] A. J. Smola, P. Bartlett, B. Schölkopf, and D. Schuurmans, *Advances in Large Margin Classifiers*, MIT Press, Cambridge, Mass, USA, 1999.
- [18] T. Joachims, "SVM^{light} Support Vector Machine V5.00," 2002, <http://svmlight.joachims.org>.
- [19] T. Joachims, "Making large-scale SVM learning practical," in *Advances in Kernel Methods—Support Vector Learning*, B. Schölkopf, C. J. C. Burges, and A. J. Smola, Eds., pp. 169–184, MIT Press, Cambridge, Mass, USA, 1999.
- [20] M. Wu and B. Liu, *Multimedia Data Hiding*, Springer, New York, NY, USA, 2002.
- [21] M. Wu and B. Liu, "Data hiding in image and video: Part-I Fundamental issues and solutions," *IEEE Transactions on Image Processing*, vol. 12, no. 6, pp. 685–695, 2003.
- [22] M. Wu, H. H. Yu, and A. Gelman, "Multi-level data hiding for digital image and video," in *Multimedia Systems and Applications II*, vol. 3845 of *Proceedings of SPIE*, pp. 10–21, Boston, Mass, USA, September 1999.
- [23] C. Dvornik, D. Tull, and P. Ramanathan, "Hybrid spatial/temporal loss concealment for packet video," in *Proceedings of the 9th International Packet Video Workshop (PVW '99)*, New York, NY, USA, April 1999.
- [24] Y.-C. Lee, Y. Altunbasak, and R. M. Mersereau, "A temporal error concealment method for MPEG coded video using a multi-frame boundary matching algorithm," in *Proceedings of IEEE International Conference on Image Processing (ICIP '01)*, vol. 1, pp. 990–993, Thessaloniki, Greece, October 2001.
- [25] M. Chen, Y. He, and R. L. Lagendijk, "A fragile watermark error detection scheme for wireless video communications," *IEEE Transactions on Multimedia*, vol. 7, no. 2, pp. 201–211, 2005.
- [26] M. Chen, M. Wu, and Y. Zheng, "Classification-based spatial error concealment for images," in *Proceedings of IEEE International Conference on Image Processing (ICIP '03)*, vol. 2, pp. 675–678, Barcelona, Spain, September 2003.

Meng Chen received the B.E. and M.E. degrees from Xi'an Jiaotong University, China, in 1994 and 1997, respectively, and the M.S. degree from University of Maryland, College Park in 1999, all in electrical engineering. Currently she is a Senior Design Engineer at Spirent Communications, and pursuing her Ph.D. degree in signal processing and communications at University of Maryland, College Park. Previously she worked at Hughes Network Systems as an Embedded Software Engineer from 1999 to 2001. Her research interests include multimedia signal processing and network performance evaluation. She received the Distinguished Teaching Assistant Award from University of Maryland in 1999 and Scholarships for Academic Excellence from Xi'an Jiaotong University from 1989 to 1995.



Yefeng Zheng received the B.E. and M.E. degrees from the Department of Electronic Engineering, Tsinghua University, Beijing, China, in 1998 and 2001, respectively. Currently, he is pursuing the Ph.D. degree in the Department of Electrical and Computer Engineering at the University of Maryland, College Park. His research interests include document image analysis, pattern recognition, and computer vision. As a codeveloper of an Asian OCR system during his Master's program, he won the National Scientific and Technological Progress Award (2nd class) of China in 2003.



Min Wu received the B.E. degree in electrical engineering and the B.A. degree in economics (both with the highest honors) from Tsinghua University, Beijing, China, in 1996, and the Ph.D. degree in electrical engineering from Princeton University in 2001. Since 2001, she has been an Assistant Professor of the Department of Electrical and Computer Engineering and the Institute of Advanced Computer Studies at the University of Maryland, College Park. Previously she was with NEC Research Institute and Panasonic Laboratories. She coauthored a book *Multimedia Data Hiding* (Springer, 2003) and holds five US patents. Her research interests include information security and forensics, multimedia signal processing, and multimedia communications. She received an NSF CAREER Award in 2002, a George Corcoran Education Award from University of Maryland in 2003, an MIT Technology Review's TR100 Young Innovator Award in 2004, and an ONR Young Investigator Award in 2005. She is a corecipient of the 2004 Best Paper Award from the EURASIP Journal on Applied Signal Processing. She is an Associate Editor of the IEEE Signal Processing Letters and served as a Guest Editor of a special issue in EURASIP JASP and Publicity Chair of 2003 IEEE International Conference on Multimedia and Expo.

