

Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/authorsrights>



Contents lists available at ScienceDirect

Signal Processing: *Image Communication*journal homepage: www.elsevier.com/locate/image

Perceptually optimized blind repair of natural images



Anush Krishna Moorthy*, Anish Mittal, Alan Conrad Bovik

Laboratory for Image and Video Engineering, Department of Electrical and Computer Engineering, The University of Texas at Austin, USA

ARTICLE INFO

Article history:

Received 10 September 2012

Received in revised form

8 July 2013

Accepted 21 August 2013

Available online 6 September 2013

Keywords:

Image quality

Image repair

Distortion identification

Perceptual image repair

ABSTRACT

We define the new idea of blind image repair as a process of correcting one or more different and unknown types of distortions afflicting an image. These distortions could introduce linear or non-linear degradations, compression artifacts, noise, etc., or combinations of these. Thus the concept encompasses denoising, deblurring, deblocking, deringing, and other post-acquisition image improvement processes that address distortions. The problem is distortion-blind when the natures of the distortion processes are unknown prior to analyzing the image. Towards solving this problem, we describe a new framework for repairing an image that has undergone an unknown set of distortions, based on identifying the distortion(s) present in the image (if any) and applying possibly multiple distortion-specific image repair algorithms. Our philosophy is based on the principle that the task of general purpose image repair is one of agglomeration, i.e., the algorithm should embody multiple high-performing distortion-specific repair modules such that seamless general purpose image repair is achieved. Our proposed framework – the GEneral-purpose No-reference Image Improver (GENII) – enables the design of algorithms that are blind to distortion type as well as to distortion parameters, and only requires as input the distorted image to be repaired. The GENII framework is modular and easily extensible to image repair problems beyond those considered here. GENII operates by using natural scene statistic models to identify distortion, to perceptually optimize the distortion parameter(s), to assess the quality of the intermediate repaired images, and to perceptually optimize the repair processes. We explain the general purpose image repair framework and one specific realization, dubbed GENII-1, which assumes that the image has been affected by one or more of four possible distortion types. The performance of GENII-1 is evaluated on 4000 distorted images, and shown to deliver substantial improvements in both quantitative and qualitative visual quality.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Image repair is the process of correcting one or more possibly different types of distortions afflicting an image. The general purpose image repair problem is formulated as

$$\mathbf{y} = \mathbf{H} \cdot f(\mathbf{x}) + \mathbf{n} \quad (1)$$

where \mathbf{y} is the observed distorted image, \mathbf{x} is the original pristine image that we seek to recover, \mathbf{n} is the additive noise, $f(\cdot)$ is a local non-linearity and \mathbf{H} is a matrix that models multiplicative distortion (e.g., a low-pass filter) [1,2]. The model in (1) is not restricted to the spatial domain (where the vectors \mathbf{x} and \mathbf{y} would be columnized versions of the 2D image) and we do not assume that the models for image repair are limited to the spatial domain.

The general image repair problem is ill-posed, and in order to solve the problem, certain assumptions are usually made about the structure of f , \mathbf{H} and \mathbf{n} . For example, for image denoising, fix f to be an identity

* Corresponding author. Tel.: +1 512 415 0213.

E-mail address: anushmoorthy@gmail.com (A.K. Moorthy).

transform, and $\mathbf{H} = \mathcal{I}$, where \mathcal{I} is the identity matrix, and assume a distribution and correlation structure on the noise \mathbf{n} [3]. For deblurring (deconvolution), assume a zero-mean noise model with known variance, then estimate \mathbf{x} from the observed \mathbf{y} [4] and so on. Solving the problem (1) has recently become quite compelling, given the proliferation of digital consumer grade cameras and the vast numbers of digital pictures being taken (about 80 billion/year [5]) and uploaded onto personal computers and social networking sites. These pictures are commonly distorted, often severely, and often by multiple distortions such as low-light noise and blur, over and under-exposure, motion effects, etc. Being able to automatically identify and fix these distortions completely blindly is highly desirable.

Image repair algorithms have been broadly partitioned into blind and non-blind classes. Blind algorithms do not assume prior knowledge of the distortion parameters, while non-blind models assume that the parameters are known. Given the ill-posed nature of the problem, there has been more activity and success on various non-blind image repair problems than on blind image repair problems [6,7]. While the general field of image repair has seen quite a bit of research, especially on single distortion problems such as denoising [6,7], deconvolution [8,4] and deblocking [9,10], the general purpose blind image repair problem, where the specific distortion(s) afflicting the image are unknown, has been little studied. There has been some work dealing with two image distortions simultaneously (and blindly, if the blur/noise parameters are unknown), the classic example being the image restoration problem [8] of simultaneously deblurring and denoising an image. The common theme of these approaches is that it is known *a priori* what the distortions are that afflict the image. A blind algorithm then seeks to discover the parameters of the distortion (noise, blur) and then ameliorate them.

We take a different approach to the problem of general purpose image repair. We begin by assuming that the distortion(s) (if any) afflicting a given image are unknown and possibly multiple, although they are assumed to come from a finite population of possible image distortions. We refer to such a problem framework as *distortion blind*. We also recognize that for specific image distortions of general interest, there exist algorithms that ostensibly correct that particular distortion reasonably well. This is not always the case, of course; for example, image restoration (deblur and denoise simultaneously) is a particularly difficult inverse problem that requires precise modeling to achieve worthwhile results. This is often impossible given that blur nearly always arises from a non-linear process. In any case, in this work we do not attempt to improve upon the state of the art of any type of image repair problem. Rather, we propose the new idea of preprocessing the image to determine the distortion(s) afflicting it and any unknown parameters of the distortion(s), then once done, ameliorate these using the best algorithm available. Of course, if the distortion identification stage is particularly effective (e.g., by better parameter estimation), then performance may be notably improved.

Thus we define the task of a general purpose image repair algorithm to be agglomeration, i.e., automatically

deploy any of multiple high-performing image repair algorithms towards achieving seamless general purpose image repair across a wide variety of distortion types. Such a general purpose image repair algorithm should perform as well as the best algorithms on each included subclass of distortion (since it embodies these algorithms in its architecture). As mentioned, given the ill-posed nature of many inverse image repair problems, it is natural that repair algorithms (even the best of them) would fail in certain situations. In these cases we further posit that the general-purpose image repair algorithm should be able to detect the failure and act on it so that the repaired image is given the best *perceptual* quality at the output, thereby rejecting failures by the internal repair algorithms. Further, given that each of the subclass repair algorithms may also introduce new artifacts (e.g., deblocking can introduce blur in the image), the general-purpose algorithm should enable iterative distortion correction within the set of subclasses that it encompasses. Finally, if some of the best repair algorithms are non-blind, the general-purpose image repair algorithm could include blind parameter estimation modules so that these non-blind algorithms operate using these estimated parameters, towards solving a class-specific blind image repair.

Here we propose both a general design framework – the General-purpose No-reference Image Improver (GENII) – for general purpose distortion-blind image repair as well as an example working model and algorithm dubbed GENII-1. Our framework and exemplar models are based on using natural scene statistics (NSS) [11,12,3] to identify distortions by type and severity. The specific exemplar model, GENII-1, is capable of restoring images distorted by additive noise, Gaussian blur, JPEG compression or JPEG2000 compression, without knowing in advance which (if any) of the distortions impairs the image, or the parameters of the distortion. Given a distorted image, the algorithm uses natural scene statistic (NSS) features to first identify whether the image has been distorted, and if so, identify (a) the distortion that afflicts the image, (b) the associated distortion parameter (e.g., noise variance) and (c) the perceptual quality of the image. The algorithm then proceeds to apply an appropriate off-the-shelf image repair algorithm based on the identified distortion category. The perceptual quality of the repaired image so obtained is repeatedly evaluated and the general-purpose image repair loop continues until a maximum level of objectively determined perceptual quality is obtained. Thus our model seeks to guarantee that the final repaired image will not only be distortion-reduced, but will also present the best possible perceptual quality. The entire process is completely blind both to the distortion types and the distortion parameters. The only information available to the algorithm is the fact that the image it is trying to repair belongs to the category of natural images¹ and that the distortions (if any) belong to one of the multiple diverse distortions.

¹ Natural images are those images captured by a camera, and do not include computer-generated renders of the visual world.

In order to achieve these goals, we use realizations of a previously proposed two-stage framework for image quality assessment [13] that first identifies the distortion that afflicts the image [14] and then proceeds to assess quality. Two realizations of this two-step framework have been developed – one in the wavelet domain [15] and the other in the spatial domain [16] – both capable of accurate distortion-identification as well as blind parameter estimation and blind image quality assessment. Either can be combined with off-the-shelf distortion-specific image repair algorithms to perform distortion-blind general-purpose, image repair. We are unaware of any approach that takes this approach to the general image repair problem and to the best of our knowledge, the proposed model is the first of its kind.

Before we proceed, we define the terminology used in the rest of this paper. A distortion *class* refers to a particular kind of distortion afflicting the image, e.g., blur or noise. A distortion *type* refers to a particular type of a distortion class, e.g., Gaussian noise, or spatially invariant Gaussian blur. Image repair refers to eliminating distortions arising from any one of the multiple classes of distortions. An exemplar implementation of the GENII framework (GENII-1) ameliorates distortions arising from a finite set of multiple distortion types.

The diverse subfields of image repair that we address have been well explored and broadly surveyed and hence we refrain from a thorough review of these techniques. Instead, we summarize some key algorithms and where appropriate point the reader to other literature in the field. We (very briefly) summarize relevant research on the following subclasses of image repair: (1) deringing, (2) deblocking, (3) denoising and (4) deblurring/deconvolution. While we are unaware of algorithms that tackle all the four distortions considered here, some algorithms tackle more than one of these distortions and this is noted in the summary below.

Deringing: Existing approaches include iterative projection onto convex sets (POCS) [17], total variation [18–20], anisotropy [21], bilateral filtering and its variants [22–25] and quadtree decompositions [26–28].

Deblocking: Prominent deblocking algorithms include those that use a field of experts model for natural images [9,29], those operating in the DCT domain [30], those that use local smoothing filters [10], and block processing [31,32] and those based on POCS [33]. A shape adaptive DCT algorithm for denoising has also been shown to perform well at deblocking images [34].

Denoising: Algorithms for denoising include subband methods [3,35–38], those that use sparse coding [39] and those based on collaborative filtering and local shape adaptation in the DCT domain [34,40]. Reviews and analysis may be found in [6,7].

Deblurring: Deconvolution algorithms include approaches based on collaborative Wiener filtering [41], statistics of natural images in the gradient domain [4,42], color statistics of natural images [8], space-variant Gaussian scale mixture (GSM) statistical modeling of wavelet coefficients [43,44] and those that tackle spatially varying blur [45]. The approaches in [8,43,44] (among others) are capable of performing both denoising and deblurring.

The main contributions of this paper are: (1) a method to use perceptual image quality assessment predictions to perform distortion-blind and distortion-severity-blind image repair; (2) a method to perceptually optimize the parameters of distortion repair algorithms so that the repaired image attains maximum perceptual quality and (3) a general purpose framework that allows for iterative distortion repair to maximize output perceptual quality. The GENII approach to image repair is markedly different from conventional image processing thinking, yielding fertile ground for addressing many questions in image optimization and perceptual image processing.

2. Distortion blind image repair

2.1. Distortion-identification and quality assessment

Recently, a two-stage model was proposed for designing distortion-blind no-reference image quality assessment (IQA) algorithms that first seeks to identify the distortion(s) that are present in the image, then proceeds to perform distortion-specific quality assessment [13]. A combination of the two stages leads to a general-purpose blind IQA model. Two realization of this framework both based on natural scene statistics (NSS) but operating in different domains (wavelet and space) were later proposed in [15,16]. The first of these, the Distortion Identification-based Image Integrity and Verity Evaluation (DIIVINE) index operates in the wavelet domain [15], while the Blind/Referenceless Image Spatial Quality Evaluator (BRISQUE) is defined in the spatial domain [16].

Both algorithms generally proceed as follows. Given a distorted image, the algorithms extract a set of features from the image and ‘learn’ the mapping from feature space to distortion-category and for each distortion class, learn the map to quality. The same set of features is used for distortion-identification and for quality assessment. The extracted statistical features quantitatively capture ‘naturalness’ in the image, and are hence useful for detecting and assessing unnatural distortions that cause perceptual annoyance. This philosophy is supported by the widely understood notion that visual cortical neurons have adapted to these ‘natural’ statistics [46]. From a performance point of view, both DIIVINE and BRISQUE correlate well with human judgements of quality as measured against large public databases of distorted images and subjective scores [47]. Fig. 1 illustrates this two-stage concept.

The algorithms are detailed in [15,16]; both are capable of identifying distortions in images that they have been

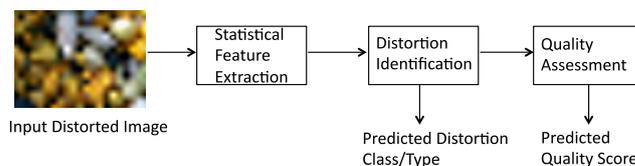


Fig. 1. An illustration of the DIIVINE/BRISQUE approach to image quality assessment [15,16] which consists of two stages: probabilistic distortion identification followed by distortion-specific quality assessment.

trained for. With respect to computation, BRISQUE is a far more efficient implementation than DIIVINE.

2.2. General purpose image repair

Our approach to general purpose image repair is summarized as follows. Given an input (possibly) distorted image, first extract statistical DIIVINE/BRISQUE features from the image, then use these features to attempt to identify the distortion afflicting the image. Once a distortion class has been posited, the same features are used to predict the perceptual quality of the image using the second stage of quality assessment. If the predicted quality of the image lies above a certain threshold, then the quality of the image is deemed to be high enough that repair does not need to be performed; in which case, the algorithm halts, yielding as output the input image. If the predicted quality falls below this threshold, the algorithm continues.

Given identified distortion(s), the same DIIVINE/BRISQUE features are used to perform blind parameter estimation for the corresponding image repair problem, for example, these features might predict the noise variance in the image if the distortion is predicted to be Gaussian noise. The algorithm then proceeds to invoke the appropriate image repair algorithm, providing as input to this algorithm the distorted image to be repaired and the associated parameter(s) that the off-the-shelf (possibly non-blind) algorithm may require. The repaired intermediate image so obtained is then passed back into the loop in order to evaluate quality and identify distortion. This loop continues until the obtained intermediate image has the highest possible quality or if a specified finite number of

repair iterations have been performed. While this procedure does not guarantee convergence, we have not found any example among the 4000 distorted images that we tested on (see below) that produced convergence issues.

Generally the question of convergence of a general repair algorithm of this sort is complex. Indeed, whether a general repair algorithm depends on the component repair modules would require a careful analysis of each algorithm and its convergence under a perceptual criterion. More pragmatically, it is likely sufficient to simply place a limit on the maximum number of iterations, unless a maxima is attained, or a desired level of quality is reached, which need not be DMOS=0 (highest achievable quality).

Finally, the algorithm determines whether the repaired image has a higher predicted visual quality than the input distorted image, and only outputs the repaired image if it has a higher visual quality, thereby avoiding unpalatable distortions that the repair algorithm may introduce as well as compensating for the failure of the repair algorithm. An illustration of this general purpose image repair scheme is diagrammed in Fig. 2.

Note that the approach that we have proposed is highly modular in the sense that any repair scheme can be replaced by another repair scheme deemed to be more effective than the one on the system. Further, the model that we have described may deploy either blind or non-blind repair algorithms. When non-blind, our features can be used to predict repair algorithm parameters.

In this paper, we demonstrate the efficacy of our general-purpose image repair framework using four distortion repair subtypes. This four-distortion blind repair system is both unique and unprecedented in its ability to

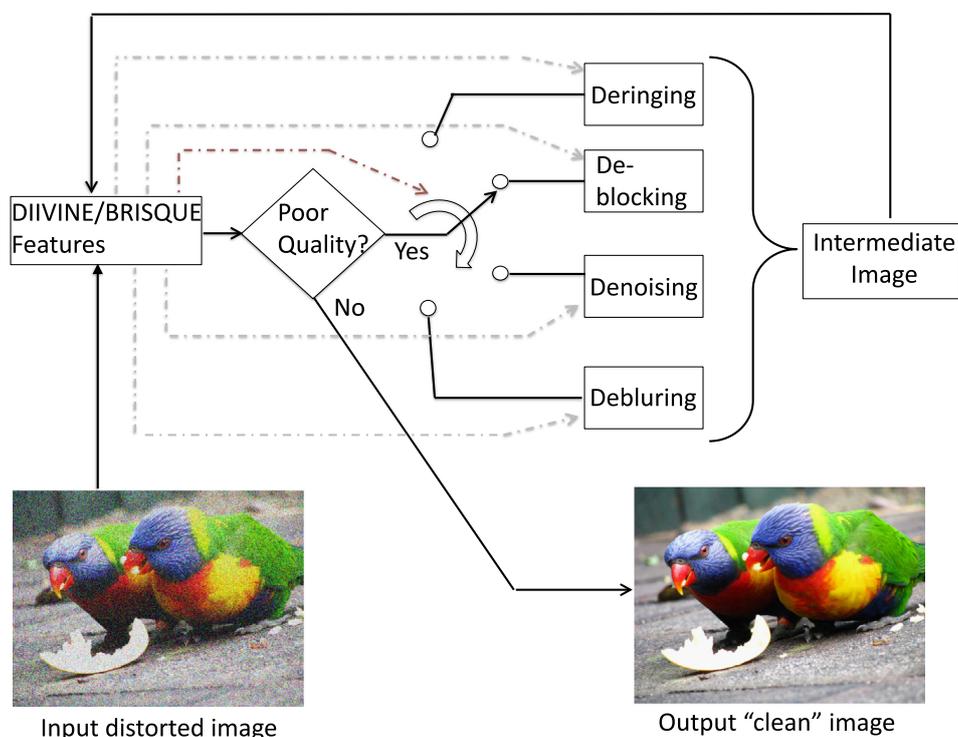


Fig. 2. An illustration of the GENII framework. DIIVINE/BRISQUE features are used to predict the distortion class, the visual quality, and the distortion parameters that may serve as inputs to a possibly non-blind repair algorithm. The intermediate repaired image is fed back to the system until the best possible quality is achieved at the output.

handle (deduce and repair) multiple unknown distortions. However, the GENII framework encompasses the possibility of more extensive implementations capable of handling more than the four distortion classes considered here, diverse subtypes of the distortions (different types of blur, noise, missing data, or other artifacts), and appropriate combinations of distortion. The resulting algorithm, GENII-1, is responsive to these distortions although blind to which (if any) of these distortions occur. We envision that future research on the GENII idea will result in repair models that handle even larger number of distortion types. However, this will first require the development of blind IQA algorithms that can handle more distortion types, which in turn will require that large human studies be conducted on these distortion on which blind IQA algorithm can be trained.

2.3. Image repair algorithms

As described in the introduction, GENII uses a two-stage framework, where once distortions are identified, off-the-shelf image repair algorithms are deployed to conduct distortion-specific image repair. Now we shall summarize the exemplar repair algorithms that we have decided to use to demonstrate the principle of general purpose image repair, in our prototype implementation GENII-1. The algorithms were chosen since they were either readily available online or were easy to implement, have a previously demonstrated high-level of performance, and adequate computational efficiency without sacrificing performance. The modularity of our model implies that any one of these algorithms could be substituted for by a suitable alternative.

The algorithms below are used to ameliorate distortions stemming from JPEG compression, JP2K compression, additive white Gaussian noise and spatially invariant Gaussian blur. The corresponding repair schemes used are deblocking, deringing, denoising and deblurring. Note that this scheme assumes that the compression schemes of JPEG and JP2K introduce mainly blocking and ringing artifacts.

2.3.1. Deblocking

We use the simple algorithm proposed in [48], which iteratively applies JPEG compression at the quality level at which the distorted image was compressed, to shifted versions of the distorted image. The resulting collection of images is averaged to produce a final deblocked image. The premise behind this approach is explained in [48], and we have found that the algorithm efficiently reduces blocking artifacts in a perceptually satisfying manner. The input parameter required, since the algorithm is not blind, is the quality factor at which the image was compressed. This can be read from the JPEG header, or it can be estimated, as we demonstrate below.

2.3.2. Deringing

We use the trilateral filter described in [49] to remove ringing artifacts from the image. The trilateral filter is an extension of the bilateral filter [23], which first computes a texture map from the gradient information and then filters

the image using a locally adaptive filtering procedure, where the filter kernels are functions of the image intensity and the textural information at each location. We tried other deringing approaches (for example, the one in [50]). However, while these approaches reduced ringing artifacts, the images produced had poorer quality than the distorted image, both by visual inspection and by quantitative QA analysis [51], whereas the trilateral filter produced higher quality images. This algorithm does not need any input parameter, i.e., the algorithm is blind.

2.3.3. Denoising

We use the Block matching 3D (BM3D) algorithm for denoising [40]. The BM3D algorithm operates as follows. Given a distorted image, with known noise variance, a set of groups of 2D image patches is created via block-matching to produce a 3D group of image patches, each of which are then denoised in a sparse transform domain using a popular wavelet shrinkage based approach [52]. These denoised patches yield a basic estimate of the denoised image, which is then used to perform re-grouping, followed by collaborative Wiener filtering, where the ‘collaboration’ is between the image patches in the group. The algorithm, which was designed for Gaussian noise, is not blind and the input parameter is the noise variance.

2.3.4. Deblurring

The approach proposed in [4] is used for deconvolution. Local image gradients are modeled using a heavy-tailed distribution, which forms a natural image prior. A *maximum a posteriori* (MAP) problem is solved using the iterative re-weighted least squares (IRLS) approach [53]. The algorithm requires the blur kernel as prior information to be able to perform deblurring. Since we consider spatially invariant Gaussian blur, the parameter to be estimated is simply the variance of the blur kernel.

3. Implementation and performance evaluation

Since there are a variety of stages involved in the exemplar image repair algorithm, GENII-1, we evaluate each stage individually, then the overall performance. Further, in order to demonstrate the robustness the GENII concept, instead of evaluating it only on a standard set of images, we perform a more complete analysis of the repair performance of GENII-1 on a much larger database. To this end, we created a large database of distorted images spanning a wide range of distortion levels encompassing the four types that GENII-1 has been designed to repair: JPEG2000 compression (JP2K), JPEG compression (JPEG), additive white noise (WN) and Gaussian blur (Blur).

A total of 300 reference images from the Berkeley image segmentation database [54] were distorted at 10 different degrees of severity for each distortion type to produce a total of 12,000 distorted images (3000 per class). JPEG compression was implemented using MATLAB's `imwrite` command; JPEG2K was implemented using the Kakadu encoder [55]; zero-mean WN was added to the image using MATLAB's `imnoise` command; and Blur was simulated using a Gaussian kernel to filter the image. The various control parameters for these

Table 1

Distortion parameters and their minimum and maximum values used for inducing distortions.

Distortion type and parameter	Min. value	Max. value
JP2K (bit-rate)	0.05	0.25
JPEG (quality parameter)	7.5	20
WN (σ^2 of filter)	0.001	0.05
Gblur (σ of filter)	1	10

distortions and the ranges of these parameters for the simulated distortion levels are listed in Table 1. The distortion levels were uniformly sampled on a log-scale between the minimum and maximum parameter values. Fig. 3 plots some sample reference and associated distorted images to give a sense of the distortion levels created.

In the discussion above, we described how GENII-1 performs distortion identification and blind parameter estimation. To achieve this, GENII-1 requires a training phase in which the extracted features are mapped onto the associated distortion type as well as to the distortion parameters. In order to train our algorithm, we split the above database randomly, based on image content, such that 200 reference images (and the associated 8000 distorted images) are used for training and the remaining 100 reference images (and the associated 4000 distorted images) are used for testing. All results to follow are reported on this testing set. This train-test split procedure ensures that there is no content overlap between the training and test sets. Next, we summarize the training procedure.

3.1. Training the model

3.1.1. Classification

A multi-class support vector machine (SVM) [56,57] was trained to classify the distorted images into one of four distortion types using DIIVINE or BRISQUE features as inputs and the labels associated with the distortion types as the outputs. The parameters of the SVM are set via cross-validation. Once trained, when fed with DIIVINE/BRISQUE features, the SVM returns a distortion type and a probability distribution over all distortion types which corresponds to the predicted type and the confidence associated with the classification respectively. This predicted distortion type is used to select the right image repair algorithm.

3.1.2. Quality assessment

The confidence associated with the prediction (probability estimates) is used in conjunction with regression modules to accomplish quality assessment. This procedure is described in great detail in [15,13]. Supposing n distortion types (in GENII-1, $n=4$), n regression modules (support vector regression (SVRs) [56]) are trained, taking as input DIIVINE or BRISQUE features, and then regressed onto (known) quality scores for each of the distortion types independently. Since human opinion scores are not available for the database that we created, we instead use the multi-scale structural similarity (MS-SSIM) index [51]

to supply quality scores. MS-SSIM produces quality predictions that correlate quite well with human judgments of quality of images impaired by these and many other types of distortions [47]. As such it is a useful proxy for human opinion scores. However, MS-SSIM correlates non-linearly with human judgments of quality. Therefore, instead of using the MS-SSIM scores directly, the MS-SSIM scores are remapped to human opinion scores obtained from the LIVE IQA database [58]. These remapped scores have a range of [0, 100], where '0' is the best possible subjective quality. This procedure is detailed in the Appendix. In order that the distinction between the MS-SSIM scores and the remapped MS-SSIM scores is clear, the remapped scores are labelled MS-SSIM_D.

During the test phase, given an input distorted image, the algorithm uses the classifier module to produce probability estimates for the types p_i , and for each distortion type produces a quality scores q_i . The final DIIVINE/BRISQUE predicted quality score is then $\sum_i p_i q_i$.

3.1.3. Parameter estimation

The repair algorithms used by GENII-1 are mostly non-blind and require as input certain parameters described in Section 2. To predict these repair parameters, for each distortion type and for each image in the training set, we use DIIVINE/BRISQUE features to train a regression module (SVR [56]) to perceptually optimize the parameter of interest (e.g., variance of the blur kernel). Given the input test image, the image is first classified by the distortion type and the appropriate (trained) regression module is queried to output the estimated parameter of that distortion. This parameter and the appropriate repair algorithm are used to produce a repaired image.

Denoising using BM3D [40] is handled in a slightly different manner to promote an improved perceptual result. Although BM3D requires noise variance as the input, when fed with the actual noise variance in the distorted image, the algorithm tends to over-smooth the image resulting in lower quality than when the algorithm is fed with a different (albeit incorrect) input noise variance [59]. For example, Fig. 4 shows a noisy image, its repaired version using the correct noise variance input to BM3D, and the perceptually optimized approach detailed below, which uses a different variance parameter as input to BM3D. BM3D tends to oversmooth images when provided with the actual noise variance, and is capable of producing better quality images when supplied with a different input parameter. We modify BM3D to improve visual quality using an important new aspect of GENII: perceptually optimized training of distortion repair parameters. The training procedure for the denoising module is modified in the following way.

To maximize the visual quality of the denoised image, we use the training procedure outlined in [59]. Specifically, during the training phase, the distorted image is denoised multiple times using BM3D with different input noise variances. The resulting repaired image is then quality-assessed using the perceptually relevant MS-SSIM [51] index. For each distorted image, the value of the input noise variance that maximizes the visual quality (as gauged by MS-SSIM) is the parameter value to train the

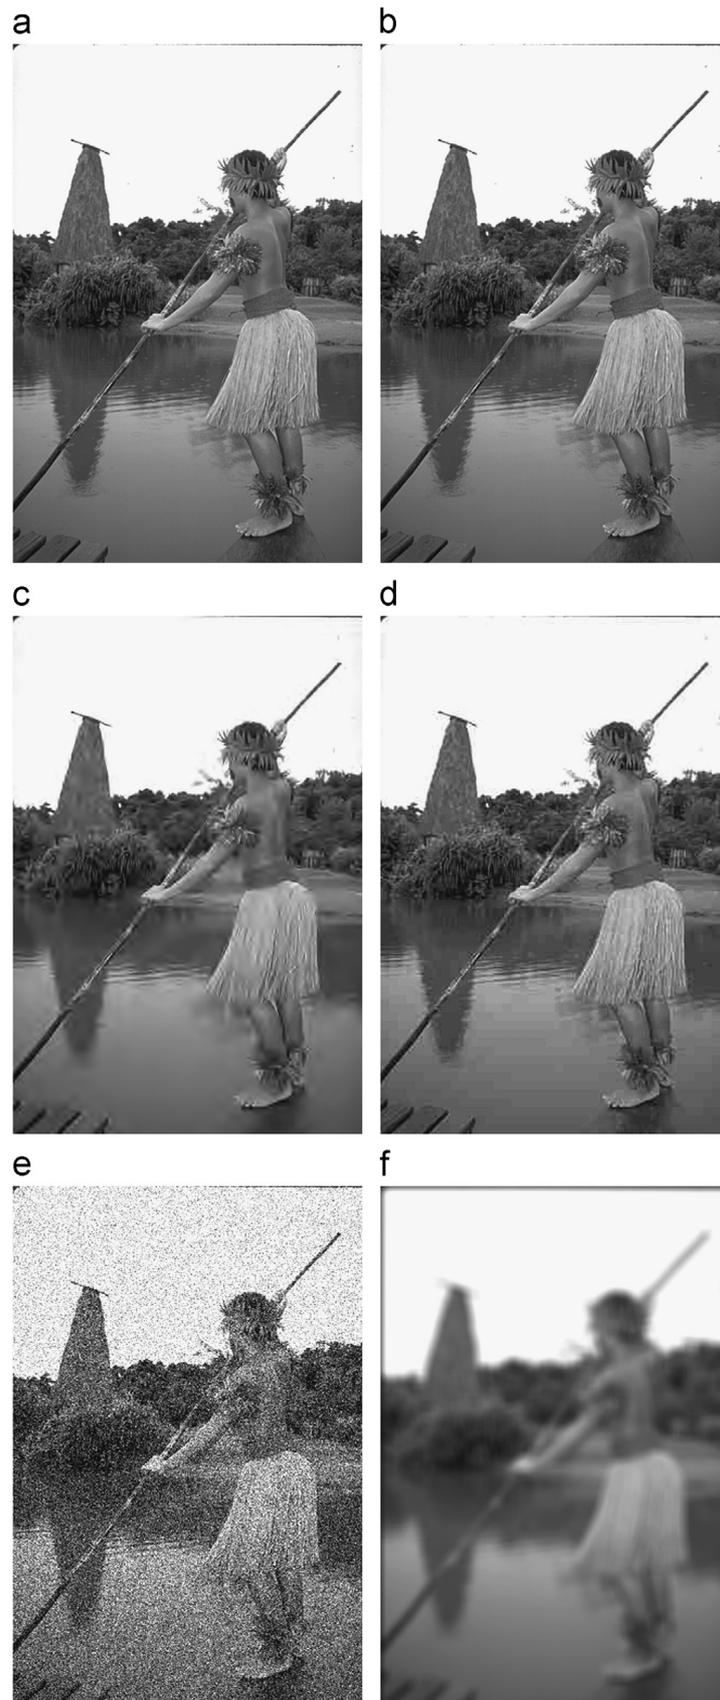


Fig. 3. Sample simulated distorted images (crops) from the Berkeley image segmentation database [54]. (a) Reference image, (b) reference crop, (c) distorted crop: JP2K, (d) distorted crop: JPEG, (e) distorted crop: WN and (f) distorted crop: Blur.



Fig. 4. Accurate noise variance as input to the algorithm in [40] produces poorer quality denoised images: (a) noisy image ($\sigma = 0.0158$, $MS\text{-}SSIM_D = 107.26$), (b) denoised with $\sigma = 0.0158$ ($MS\text{-}SSIM_D = 64.00$) and (c) denoised with $\sigma = 0.0040$ ($MS\text{-}SSIM_D = 53.82$).

noise parameter regression module. In summary, instead of using the actual noise variance as input to the regression module, a different perceptually optimized value is chosen.

Note that this training procedure is specific to the BM3D denoising algorithm and may not be needed for other algorithms that might be used to replace BM3D in an improved GENII implementation. On the other hand, perceptual optimization of an image repair parameters is a powerful option.

In summary, we train one classification module of GENII-1 that outputs a probability distribution over distortion types, four regression modules that output quality scores independently for each of the classes, and four parameter estimation regression modules which output appropriate distortion parameters, all of which utilize as input only natural scene statistic DIIVINE/BRISQUE features. During the test phase, the DIIVINE/BRISQUE features are extracted just once and from them all of the necessary outputs are produced using the trained modules to (1) classify the image, (2) assess quality, and (3) estimate distortion parameters. The predicted dominant (most possible) distortion type and the distortion parameters are used to perform image repair. As we noted before, this entire process can be repeated until a maximum quality is achieved.

We first consider the case of a single pass through the system i.e., the output repaired image is not fed back into the loop for further correction. Later we also demonstrate iterative image repair. Fig. 5 shows an example of the operation of the GENII-1, where deconvolution was performed.

3.2. Performance evaluation

Next, we evaluate the specific instantiation of our general purpose image repair framework, GENII-1, on the test image database. Specifically, we evaluate the accuracy achieved when predicting distortion type, intermediate image quality, distortion parameters and the final improvement in visual quality obtained after repair.

3.2.1. Classification and quality assessment

Table 2 reports classification accuracy for each distortion type as well as the overall classification accuracy over the 4000 distorted images using both DIIVINE and BRISQUE features. Excellent classification accuracy was achieved; there was very little confusion between distortion types ($< 1\%$) and hence, for brevity, we do not report these numbers here.

We computed the Spearman's rank ordered correlation coefficient (SROCC) between the predicted quality scores from DIIVINE and BRISQUE and the objective predicted by $MS\text{-}SSIM_D$ on the test images across all distortions. We also computed the commonly used *full-reference* peak signal-to-noise ratio (PSNR) as an additional comparison. The SROCC values relative to $MS\text{-}SSIM_D$ that were observed were – PSNR=0.8066, DIIVINE=0.9308, BRISQUE=0.9490. True $MS\text{-}SSIM$ would be at an SROCC of 1.0. As expected, DIIVINE and BRISQUE both correlate much better with the perceptually relevant $MS\text{-}SSIM_D$ than does PSNR.

To verify that the predicted SSIM scores relate to quality well, we also computed SROCC between the DIIVINE predicted $MS\text{-}SSIM$ scores and the quality scores from another high-performing full-reference algorithm – the visual information fidelity (VIF) index, and obtained an SROCC of 0.7757. True $MS\text{-}SSIM$ values achieve an SROCC of 0.8138, while PSNR achieves an SROCC of 0.6302. We also trained the features on VIF scores and tested the correlation with respect to VIF and $MS\text{-}SSIM_D$; the SROCC so obtained for DIIVINE with respect to VIF is 0.9316 and 0.7471 with respect to $MS\text{-}SSIM_D$ while PSNR achieves a correlation of 0.6302 with respect to VIF.

Thus, the NR measures predict quality as well as $MS\text{-}SSIM$, and exhibit good correlation with human perception.

3.2.2. Parameter estimation

Fig. 6 plots the mean estimated parameters for each distortion type as a function of the actual input parameters. The associated standard error bars across the 100 different contents in the test set are given for both DIIVINE and BRISQUE. The figure also lists the root mean-squared-error (RMSE) between the actual value and the predicted value. Both DIIVINE and BRISQUE do a good job of predicting the distortion parameter, and hence their predictions can be used as inputs to non-blind repair algorithms.

Note that in GENII-1, for the denoising task, we do not actually use the predicted noise variance as the parameter for BM3D, for reasons explained earlier. Also, the deringing algorithm used does not require any input parameter. The plot simply demonstrates that the DIIVINE/BRISQUE framework is capable of predicting these distortion parameters with a high degree of accuracy.

3.2.3. Image repair

Having demonstrated that DIIVINE/BRISQUE features are capable of classifying images according to distortion

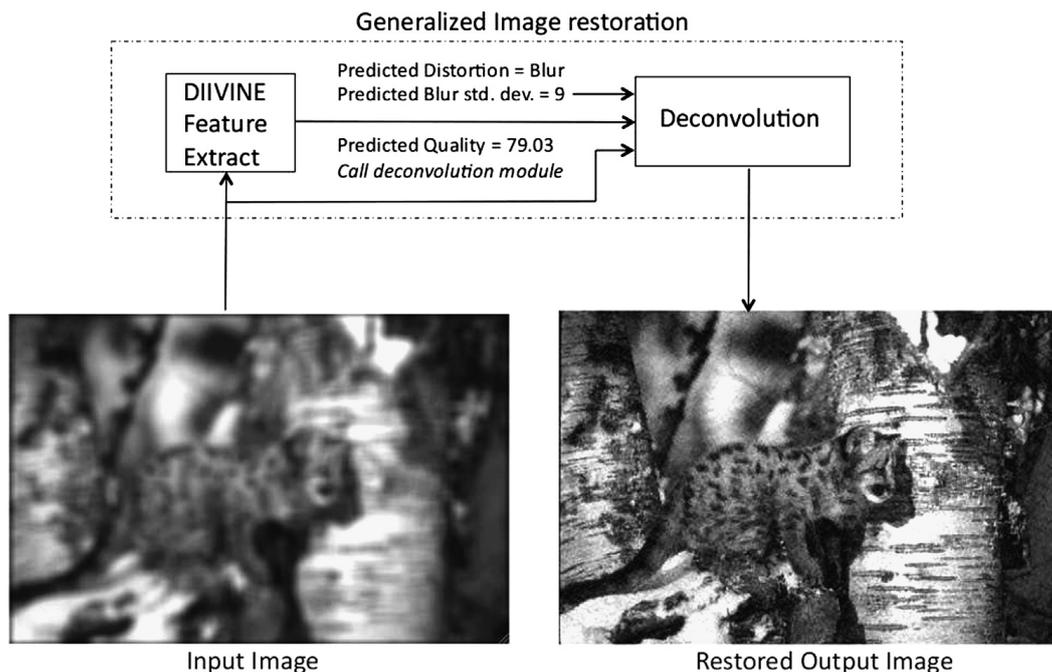


Fig. 5. Illustration of the operation of GENII-1 using DIIVINE features extracted from the input image. These features are used to identify the distortion, predict the quality and estimate the blur kernel standard deviation. The distorted image and the blur kernel are then fed to the appropriate repair scheme – deconvolution – to produce the output repaired image.

Table 2
Classification accuracies of DIIVINE and BRISQUE.

	JP2K (%)	JPEG (%)	WN (%)	Blur (%)	All (%)
DIIVINE	98.10	98.20	100	99.30	98.90
BRISQUE	99.30	100	99.90	99.30	99.63

types, quality assessment, and blind parameter estimation, we now demonstrate how these stages can be combined to fully realize GENII-1 using off-the-shelf algorithms for image repair. Since we are unaware of any other general purpose image repair technique similar to GENII-1, comparisons are impossible. Hence, we report the mean increment in MS-SSIM_D quality (and the standard error bars) after the repair process over the baseline MS-SSIM_D quality of the distorted image in Fig 7 for each distortion type, and across distortion classes. Fig. 7 shows results for a single-pass of GENII-1, which only repairs the image once based on the distortion category and then calls the appropriate repair algorithm. This single-pass implementation of GENII-1 also checks the image quality at the output and returns the image (repaired or input distorted) having the higher quality, as predicted by the image quality index being used (DIIVINE or BRISQUE). This quality check guarantees the best quality at the output, and accounts for distortions possibly introduced in the repair process which may have reduced the perceptual quality (although the original distortion may have been repaired, e.g., deblocking leading to blur).

For comparison purposes, Fig. 7 also plots the quantitative quality improvement obtained when using each image repair algorithm with perfect knowledge of the distortion (i.e., without the classifier stage), and perfect knowledge of the input distortion parameters (i.e., using non-blind algorithms). In order to demonstrate that the use of MS-SSIM_D

does not bias the results, we also use another high-performance quality measure: the visual information fidelity index (VIF) [60] as a validating measure. Note that VIF scores are not on the DMOS scale, and hence the differences are non-linearly related to quality, thereby magnifying them. The differences in gains across the two quality measures are a function of how each of them correlate with human perception for that particular distortion and how well the algorithms correlate with each other (SROCC=0.8138).

The ‘perfect’ baseline could be viewed as being unfair to the GENII since it has perfect information. Alternatively, one could view GENII as having an unfair advantage owing to its ‘quality-aware self correction’. Thus, while not the ideal algorithm to compare against the ‘perfect’ repair, it provides insights into the performance of the proposed approach.

Fig. 7 indicates that GENII-1 performs quite well predicting and ameliorating the distortions present in the image. Note that the improvement in quality that GENII-1 delivers is limited by the performance of the repair modules that it uses. A better repair algorithm will lead to greater increases in visual quality, as evidenced by the large gains in MS-SSIM_D obtained for deconvolution and denoising.

In order to study the effect of the amount of distortion on the performance of GENII-1, Fig. 8, plots the gain in quality (as measured by MS-SSIM_D) as a function of distortion severity for each distortion type. Again, the gains that would be obtained by the (unfair) baseline approach equipped with perfect knowledge of the distortion type and the input distortion parameters are also plotted for comparison purposes. For all distortions, an increase in severity reduces gains in objective quality (MS-SSIM_D). The results indicate that the perceptually optimized GENII-1 performs as well as (if not better than) the ‘perfect’ repair scheme, even though the ‘perfect’ scheme has full knowledge of the distortion parameters.

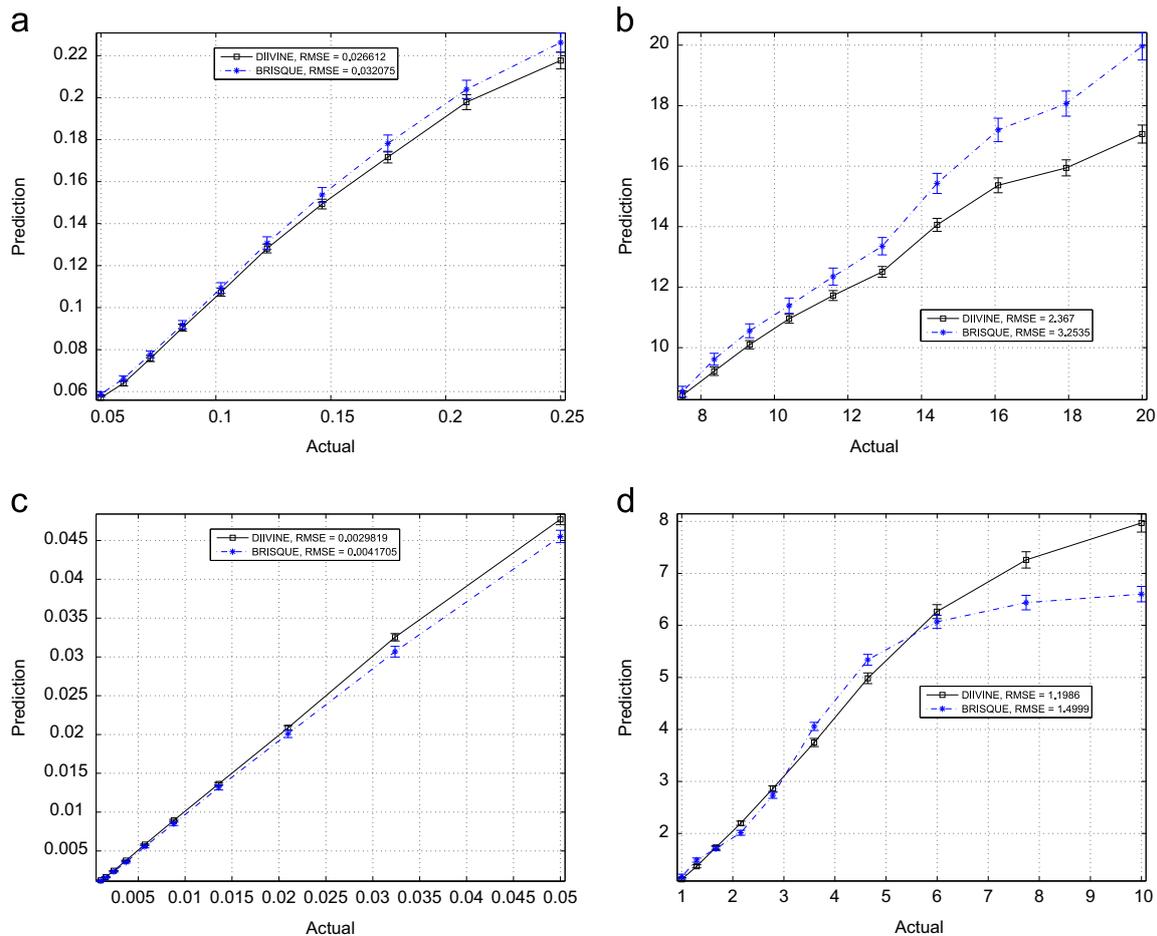


Fig. 6. Parameter estimation using DIIVINE and BRISQUE: plots of (mean) predicted vs. actual parameters and the standard error bars of distortions considered here. Each subfigure indicates the distortion type and the root mean-squared-error (RMSE) between the actual and predicted values. (a) JP2K: bit rate, (b) JPEG: quality parameter, (c) WN: noise standard deviation and (d) Blur: standard deviation of blur kernel.

The case of JP2K requires further explanation. Although the perfect repair algorithm is a somewhat unfair baseline, it does not have the added advantage of quality-driven self-correction and so, in many cases the output images obtained are of inferior quality relative to the input distorted image.

The training procedure for WN was also modified to provide perceptually optimized denoised images, and the perfect reconstruction baseline does not have this training-based advantage. Although the gains obtained are not reflected in the mean quality-gain plots, on individual images, such a training procedure does indeed produce better quality, as exemplified by Fig. 4.

Finally, to provide a visual illustration of the results, Fig. 9 plots samples of distorted images from the test set and their repaired versions using GENII-1 with DIIVINE features along with the quantitative changes in objective quality that were obtained. Additional result is plotted in the accompanying Supplementary Material.

3.2.4. Iterative image repair

Iterative repair using GENII-1 can proceed as illustrated in Fig. 2, where the repair is performed in a loop until a stopping condition is reached. This condition could be a pre-fixed threshold on quality (which may not always be achieved) or one that assesses the amount of improvement

in quality, and stops when the improvement becomes small, ceases to be positive or some combination.

The improvement in quality obtained could be computed relative to the original distorted image at each iteration, or as a difference between the quality at the current iteration and the previous one. Since the repair chain is not guaranteed to produce a steady improvement in quality at each iteration, the best solution would be the former, where one computes the difference between the current quality and the distorted image quality and continues the loop until the improvement is negative. The algorithm would then pick the intermediate image which yields the highest predicted perceptual quality. While this solution is optimal in the current setup, it is time consuming, and stopping if the quality change is negative as compared to the previous iteration may be an attractive alternative in a practical implementation. We now demonstrate two examples of iterative repair using these stopping criteria in Figs. 10 and 11.

Fig. 10 illustrates a case where the quality increases steadily with iteration count reaching a well-defined minimum at the fifth repair iteration. Beyond this point, deconvolution failure occurs at iteration 13, and the algorithm stops. Due to our design, the best quality image (iteration 5) is produced as output. In this case, the stopping criterion which checks only the previous iteration would have produced the same result.

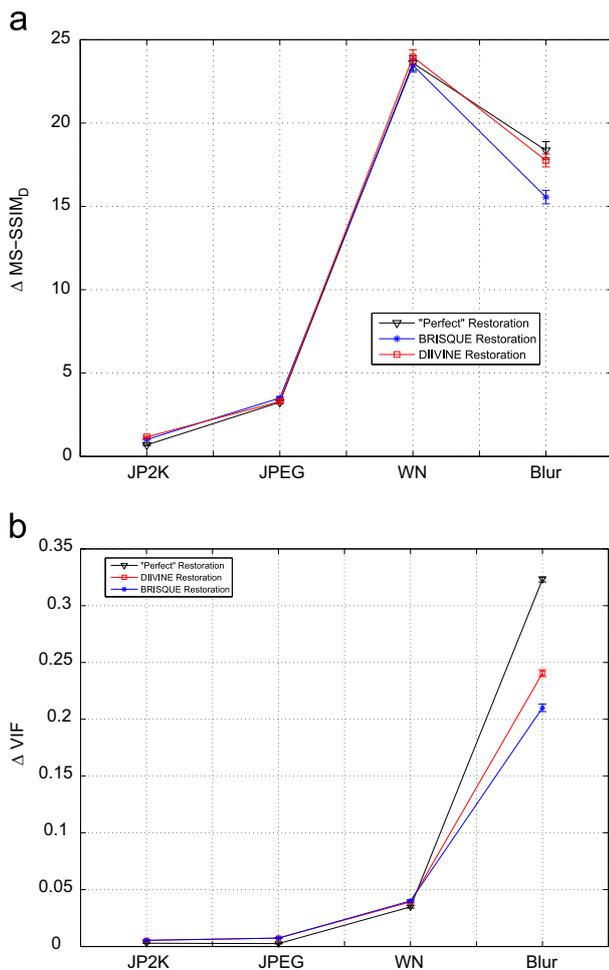


Fig. 7. Mean increments in quality and the standard error bars for perfect repair, DIIVINE-based and BRISQUE-based GENII-1 algorithms using (a) MS-SSIM_D and (b) VIF as quality measures.

Fig. 11 plots a case where the objective quality degrades in iteration 3 as compared to iteration 2. The simpler stopping criterion would have produced as output the image at iteration 2 (Fig. 11(b)). However, the quality score at iteration 4 is far lower (better) than at iterations 2 and 3. If the more exhaustive stopping criterion was used, the image at iteration 4 would be produced as the output. Deconvolution failure occurs at iteration 6 (image not shown), where the algorithm stops.

3.3. Images with multiple distortions

Images with multiple distortions have generally not been explored much in the context of image repair except for the image restoration problem (blur+noise) [8]. There has also been some work on quality assessment of images distorted with multiple distortions [61]. Understanding how multiple distortions in an image interact and mask each other and their effect on perceptual quality and image repair schemes is a large field that deserves considerable study (certainly beyond the scope of this article). However, we conducted a preliminary analysis of general purpose image repair for the case of additive white Gaussian noise+JPEG compression. This scenario is common; for

example, a low-light image compressed and uploaded on one of the photo-sharing websites (noise+JPEG).

We randomly chose a reference image from the test set and simulated the multi-distorted image, where the individual distortions were applied in sequence. The noise variance was set at $\sigma^2 = 0.01$ and JPEG quality (MATLAB) was set at $Q=5$. We then applied our iterative image repair process using GENII-1. Since there are multiple distortions in the image, the distortion type predicted to have the maximum probability of being present (as indicated by the classifier) is picked. Ties are broken randomly. The iterative process of repair continues as outlined before, and is stopped after two repair stages. Fig. 12 plots the multi-distorted image, the output image at the first iteration and the final repaired image.

The first distortion identified is blocking, which is ameliorated by the deblocking algorithm. The deblocked image is less noisy due to smoothing of the deblocking algorithm, which in turn introduces blur, so the deblocked image is identified as having blur. The deconvolution step fails, and GENII-1 produces the deblocked image.

As Fig. 12 indicates, GENII-1 performs well on identifying the distortions present in the image, even when the image is distorted with multiple distortions. Naturally more robust deconvolution algorithm would possibly fail less often, producing better quality output images. Note that we did not re-train the classifier to accommodate multi-distorted images. Generally, modeling multiply distorted images and their repair processes remains an appealing area for future work.

4. Discussion and conclusion

We have introduced a distortion-blind perceptually optimizable general-purpose image repair paradigm called GENII that repairs images distorted by any of multiple distortions by using natural scene statistics to (1) identify the likely distortion(s) impairing the image, (2) to estimate the quality of the distorted image, (3) and to estimate the parameters (i.e., distortion severity) of the distortion, and (4) based on these estimated data, selects an appropriate (possibly non-blind) repair module. Steps (1)–(3) are performed using NSS features extracted from the image either in the wavelet domain or in the spatial domain [15,16]. We demonstrated a working prototype, dubbed GENII-1, capable of repairing images impaired by any of JPEG200, JPEG, additive Gaussian white noise or linear Gaussian blur distortions.

GENII is the first system of its kind that first identifies any of multiple unknown distortions coming from a trained set and then attempts to fix the image. It is modular and easily extensible to distortions beyond those considered here. The modularity of the approach implies that one could replace any of the repair modules used (e.g., by GENII-1) with better or more appropriate repair algorithms leading to even better performance. New modules could be added for additional distortion classes and/or types, including multi-distortions. This will require design and creation of suitable distorted image databases and associated human studies. Further, the GENII framework was designed such that the output image quality will

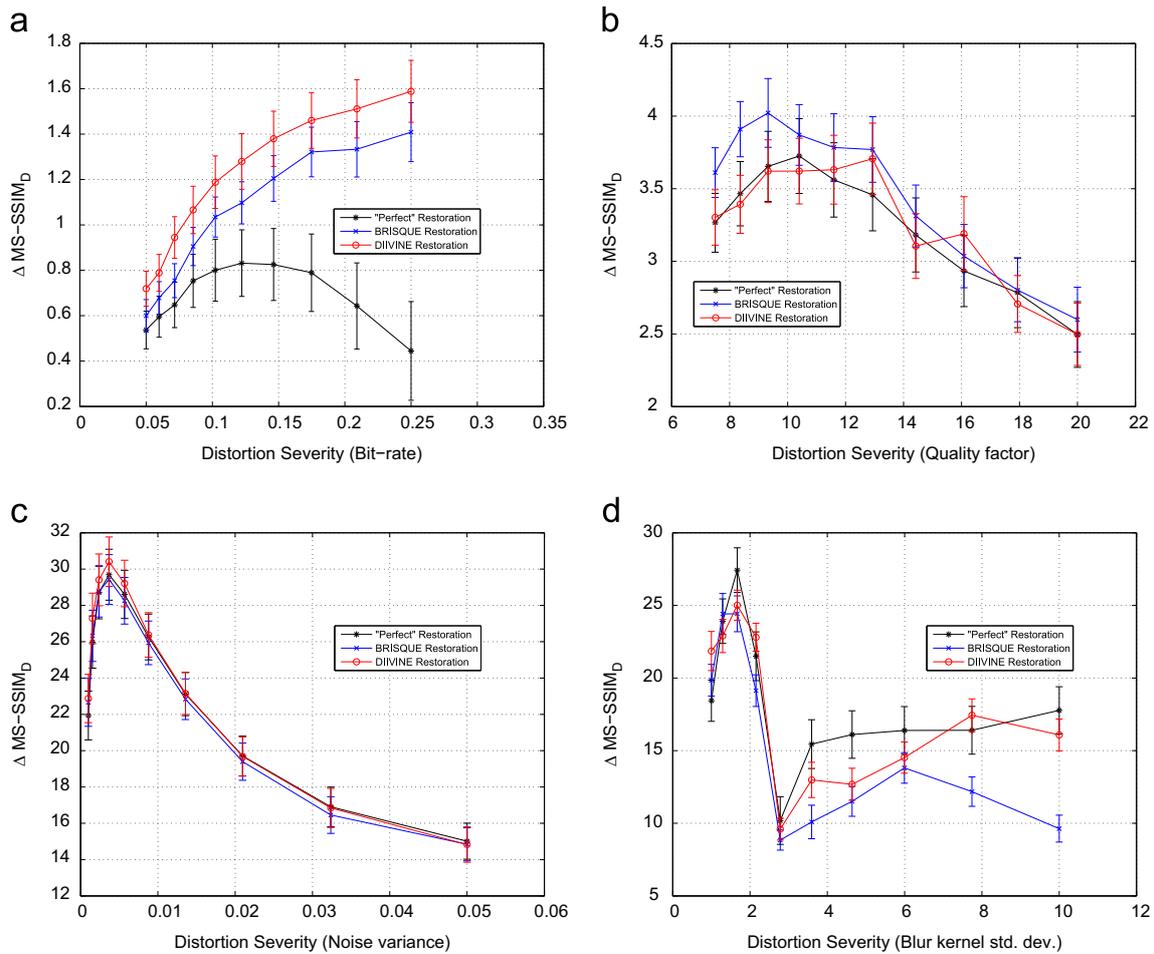


Fig. 8. Mean changes in objective quality ($MS-SSIM_D$) and the standard error bars for perfect repair, DIIVINE-based and BRISQUE-based generalized repair as a function of distortion severity for: (a) JP2K (deringing), (b) JPEG (deblocking), (c) WN (denoising) and (d) Blur (deblurring).

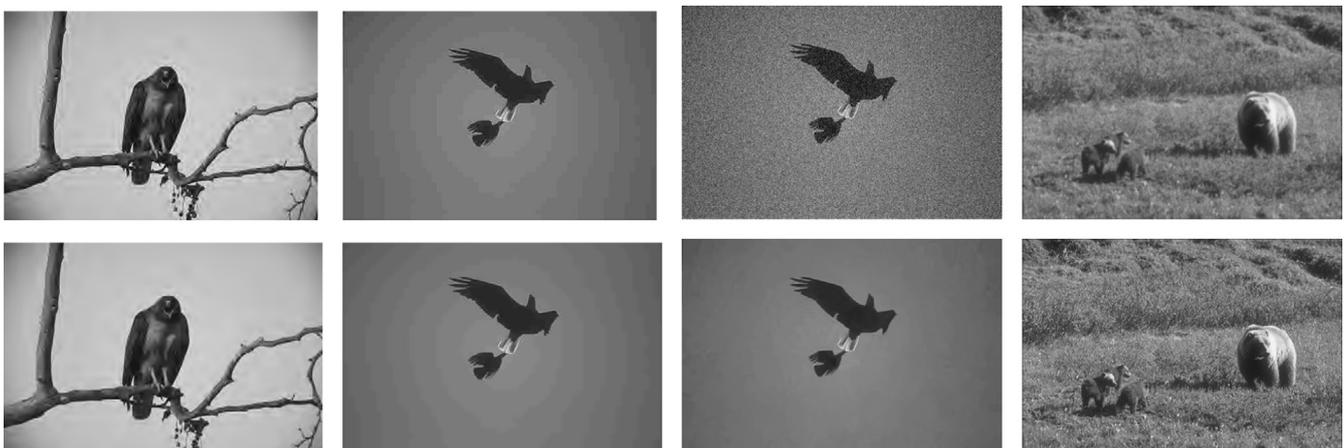


Fig. 9. Sample distorted images (first row) and their repaired versions (second row) obtained using the proposed blind general purpose image repair framework. Distortions (quality gains), left–right: JP2K (7.90), JPEG (15.27), WN (70.60), Blur (51.18).

always be at least as good as the input image quality, as measured by a high-quality, objective but perceptually relevant image quality assessment algorithm, thereby accounting for failures of the repair modules. The iterative nature of GENII implies that distortions introduced by the repair algorithms may also be eliminated, thereby increasing the quality of the output image.

The GENII framework is a radically different approach to image repair, that seeks to maximize the visual quality of the images, as measured by a no-reference image quality assessment algorithm, instead of simply targeting the distortion(s) present in the image. We showed that this new approach to image repair leads to significant improvements in output quality both visually and

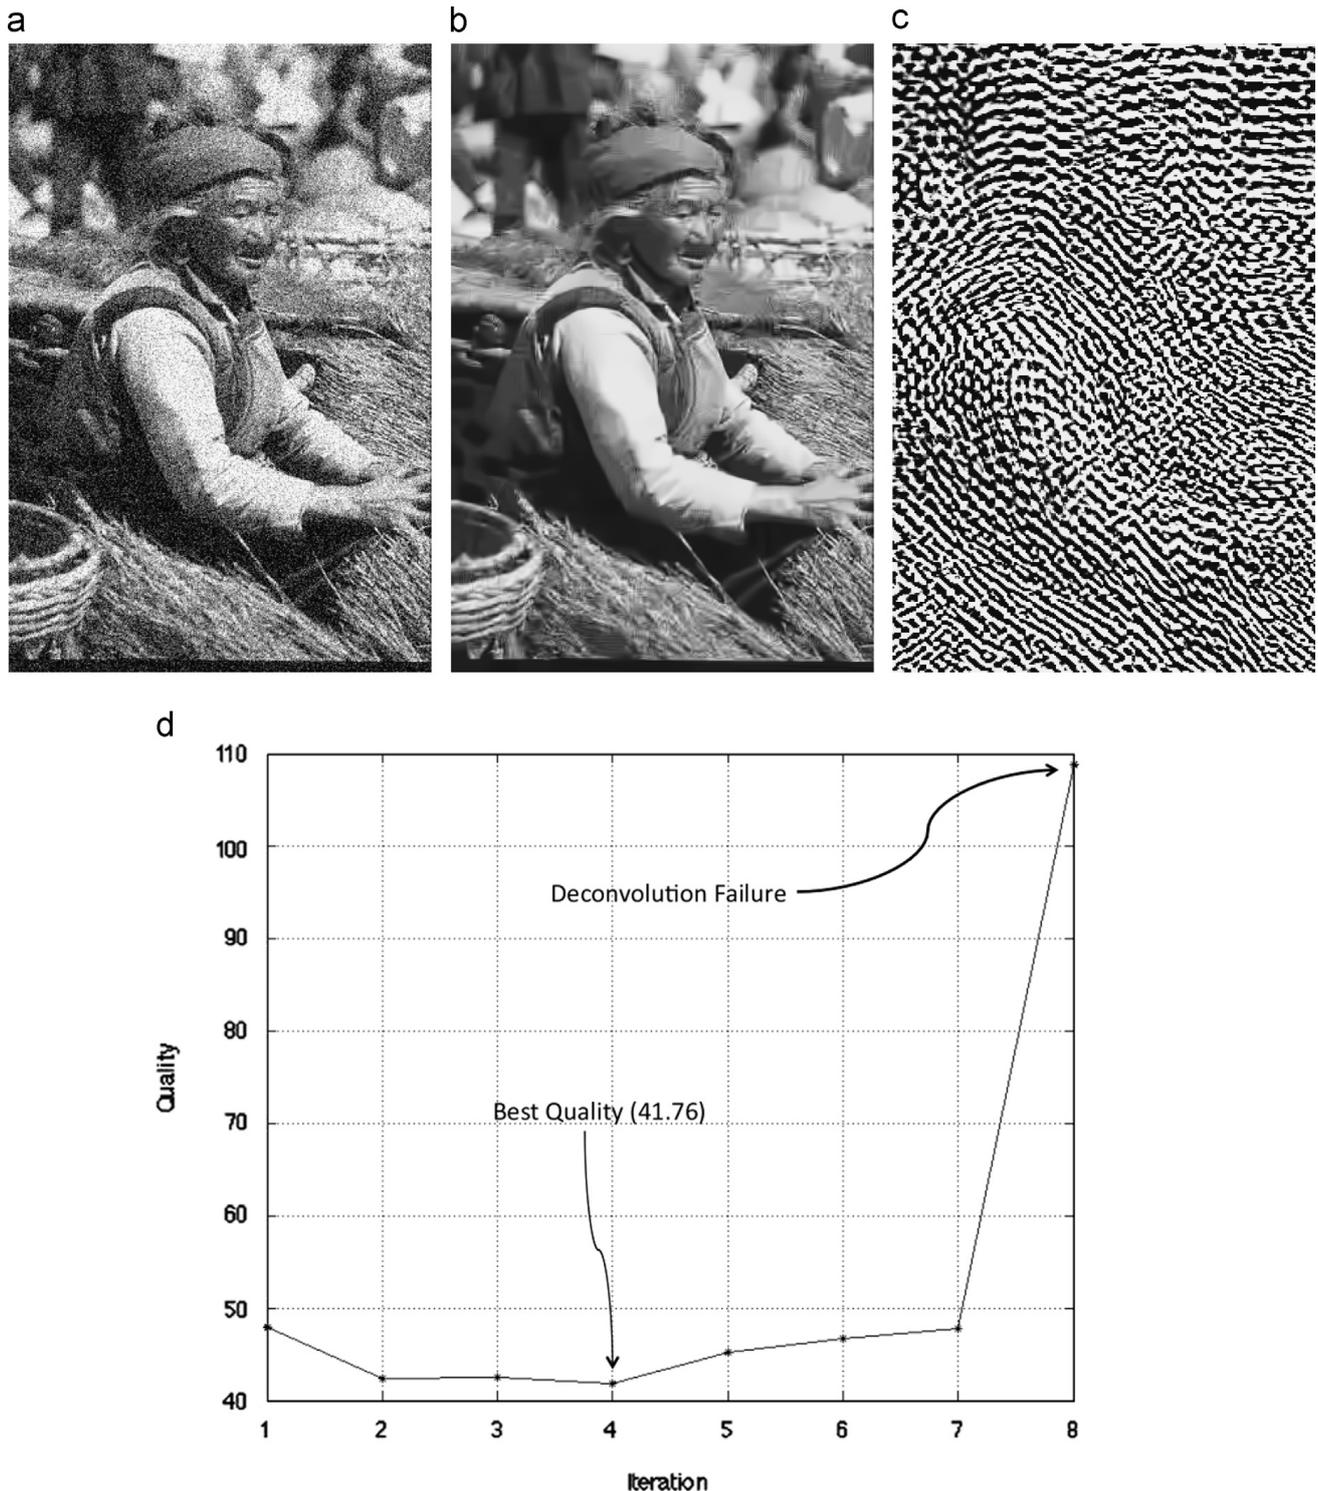


Fig. 10. Example iterative image repair using GENII-1 driven by DIIVINE features for image distorted with white noise, see text for explanation. (a) Distorted image, (b) best quality repaired image, (c) deconvolution failure at iteration 13, (d) quality as a function of repair iterations with predicted distortion type labels. GENII-1 outputs (b).

quantitatively as measured by a high-performance full-reference objective image quality assessment algorithm. To the best of our knowledge, the proposed framework is the first of its kind to approach the general image repair problem from a perceptual optimization point of view although related problems such as objective quality-driven models for image restoration [62], denoising [63], compression [64,65], and deblocking [66] have been studied.

It is also the first model to combine a no-reference image quality index with distortion identification to perform general purpose image repair.

GENII uses off-the-shelf algorithms for image repair, and an important point to address is the choice of these algorithms. In the exemplar model GENII-1, we chose algorithms that are computationally efficient and that result in repaired images exhibiting higher levels of

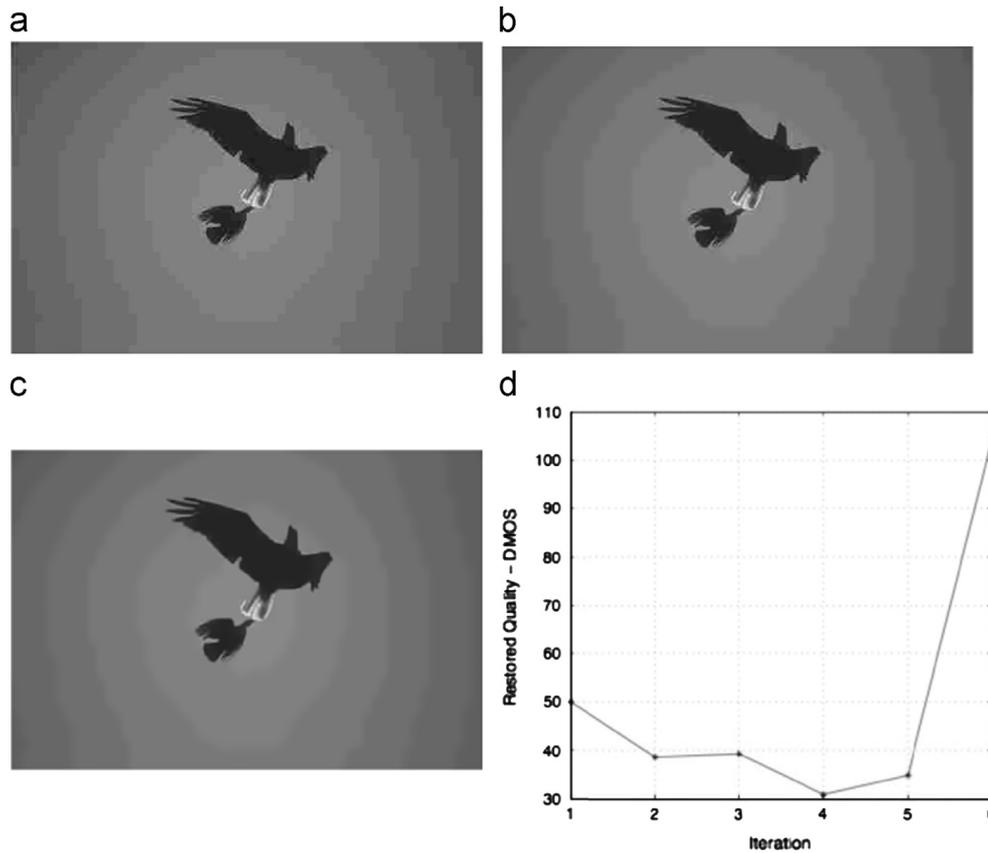


Fig. 11. Example iterative image repair using GENII-1 driven by DIIVINE features, see text for explanation. (a) Distorted image ($MS-SSIM_D=50.1$), (b) repaired image at iteration 2 ($MS-SSIM_D=38.51$), (c) repaired image at iteration 4, highest quality ($MS-SSIM_D=30.81$), (d) quality as a function of repair iterations. GENII-1 outputs (c).



Fig. 12. Multiply distorted images and their repaired versions. GENII-1 outputs (b). (a) Noise+JPEG, (b) deblocked and (c) deblurred.

(predicted) visual quality. Since the repair performance of GENII is as good as the individual algorithms used, improvements in the performance of GENII can be obtained as image repair algorithms continue to evolve. The choice of the repair algorithms may also be governed by the complexity and the requirements of the application where GENII is deployed. The use of state-of-the-art algorithms would lead to a GENII performance at least as good as these algorithms.

GENII is extendible to distortions beyond those considered here. While we have not explicitly demonstrated this here, we have previously shown that NSS-based classification is capable of handling other distortions such as packet loss [15,16]. We are currently studying NSS-based classification for other a wider variety of distortions. As we state this, it may be prudent to temper expectations.

In its current form, GENII-1 is capable of repairing images distorted with only four classes of distortions. Even amongst these classes, the distortion types are limited to what may be viewed as simplistic distortions: additive white Gaussian noise and spatially invariant Gaussian blur. While we remain cautiously optimistic about the general GENII philosophy, we nevertheless realize that GENII is not a panacea to all distortions classes and types.

The results we have obtained indicate that misclassification very rarely occurs amongst the distortions that GENII-1 handles, and when it does, it generally occurs on high quality images where distortion(s) may be only weakly perceptible [14]. If one were to set the threshold of quality at the outset before conducting repair (as described in GENII, but not implemented in GENII-1), then the frequency of misclassification would be reduced.

Further, since a quality check is performed on the output, incorrectly repaired images are discarded yielding output image quality that is at least as good as the input image.

Our current and future work involves extending the GENII concept to distortions even more diverse and extensive than those considered here and modeling images distorted by multiple coincident distortion and suitable repair processes for such images. The development of GENII models capable of handling multiple distortions will require a number of new developments. First, since the GENII framework is NSS-based, studies of NSS of multiply distorted images will need to be undertaken. This will necessarily include studies of human subjective judgements of these multi-distortions, since the complex way in which the distortion may interact, both in terms of the way in which they modify image structure and the way they are perceived, is likely to be nonlinear and complicated. In essence, they must be viewed as new distortions. Such a deeper NSS and perception-based analysis will enable modifications of current no-reference IQA indices, such as DIIVINE and BRISQUE, to handle such multi-distortions.

Acknowledgment

This research was supported by the National Science Foundation under grants CCF-0728748, IIS-1116656 and by Intel and Cisco corporation under the VAWN program.

Appendix A. Mapping MS-SSIM to DMOS

Instead of directly using the MS-SSIM scores to quantify quality, we re-map the MS-SSIM scores to the more easily interpreted perceptual scale of differential mean opinion scores (DMOS), obtained from human subjective studies such as that in [47]. We use the human DMOS obtained from [47], and map MS-SSIM scores via a logistic function fit (2), where the parameters β_i $\{i = 1, 2, \dots, 5\}$ are estimated via a nonlinear optimization procedure (MATLAB function `nlinfit`) between the DMOS and the MS-SSIM scores. This was done for each of the four distortions targeted by GENII-1, to produce MS-SSIM_D.

$$f(x) = \beta_1 \left[\frac{1}{2} \frac{1}{1 + \exp(-\beta_2(x - \beta_3))} \right] + \beta_4 x + \beta_5 \quad (2)$$

The non-linear fitting procedure detailed here is identical to that used in [47] prior to computing linear correlation and root-mean squared error between algorithm scores and DMOS.

While such a remap using a database is limited by the database and is specific to it, the LIVE IQA database of [47] incorporates a wide variety of distortions levels and spans a good range of visual quality and hence, the re-mapped scores obtained are reasonable representations of visual quality on the linear DMOS scale, where 0 indicates the best quality and 100 indicates the worst quality.

Appendix B. Supplementary material

Supplementary data associated with this article can be found in the online version at <http://dx.doi.org/10.1016/j.image.2013.08.007>.

References

- [1] H.C. Andrews, B.R. Hunt, Digital Image Restoration, Prentice-Hall Signal Processing Series, vol. 1, Prentice-Hall, Englewood Cliffs, 1977.
- [2] M. Banham, A. Katsaggelos, Digital image restoration, IEEE Signal Processing Magazine 14 (2) (1997) 24–41.
- [3] J. Portilla, V. Strela, M.J. Wainwright, E.P. Simoncelli, Image denoising using scale mixtures of Gaussians in the wavelet domain, IEEE Transactions on Image Processing 12 (11) (2003) 1338–1351.
- [4] A. Levin, R. Fergus, F. Durand, W.T. Freeman, Deconvolution using natural image priors, in: ACM SIGGRAPH, 2007.
- [5] A. Fiegl, Image obsessed, National Geographic 221 (2012) 35.
- [6] P. Chatterjee, P. Milanfar, Is denoising dead? IEEE Transactions on Image Processing 19 (4) (2010) 895–911.
- [7] A. Buades, B. Coll, J.M. Morel, A review of image denoising algorithms with a new one, Multiscale Modeling and Simulation 4 (2) (2006) 490–530.
- [8] N. Joshi, C.L. Zitnick, R. Szeliski, D.J. Kriegman, Image deblurring and denoising using color priors, in: IEEE Conference on Computer Vision and Pattern Recognition, 2009, pp. 1550–1557.
- [9] D. Sun, W.K. Cham, Postprocessing of low bit-rate block DCT coded images based on a fields of experts prior, IEEE Transactions on Image Processing 16 (11) (2007) 2743–2751.
- [10] O. Bryt, M. Elad, Improving the k-SVD facial image compression using a linear deblocking method, in: IEEE Convention of Electrical and Electronics Engineers in Israel, 2008, pp. 533–537.
- [11] A. Srivastava, A.B. Lee, E.P. Simoncelli, S.C. Zhu, On advances in statistical modeling of natural images, Journal of Mathematical Imaging and Vision 18 (1) (2003) 17–33.
- [12] D.L. Ruderman, The statistics of natural images, Network Computation in Neural Systems 5 (4) (1994) 517–548.
- [13] A.K. Moorthy, A.C. Bovik, A two-step framework for constructing blind image quality indices, IEEE Signal Processing Letters 17 (2) (2010) 587–599.
- [14] A.K. Moorthy, A.C. Bovik, Statistics of natural image distortions, in: IEEE International Conference on Acoustics Speech and Signal Processing (ICASSP), 2010, pp. 962–965.
- [15] A. Moorthy, A. Bovik, Blind image quality assessment: from natural scene statistics to perceptual quality, IEEE Transactions on Image Processing 20 (12) (2011) 3350–3364.
- [16] A. Mittal, A.K. Moorthy, A.C. Bovik, Blind/referenceless image spatial quality evaluator, in: Asilomar Conference on Signals, Systems and Computers, 2011.
- [17] X. Li, Improved wavelet decoding via set theoretic estimation, IEEE Transactions on Circuits and Systems for Video Technology 15 (1) (2005) 108–112.
- [18] A. Krylov, A. Nasonov, Adaptive total variation deringing method for image interpolation, in: IEEE International Conference on Image Processing, 2008, pp. 2608–2611.
- [19] A.V. Nasonov, A.S. Krylov, Adaptive image deringing, in: Proceeding of Graphicon, 2009.
- [20] A.V. Nasonov, A.S. Krylov, Scale-space method of image ringing estimation, in: IEEE International Conference on Image Processing, 2009, pp. 2793–2796.
- [21] P. Bourdon, B. Augereau, C. Olivier, C. Chatellier, A PDE-based method for ringing artifact removal on grayscale and color JPEG2000 images, in: IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003.
- [22] C. Tomasi, R. Manduchi, Bilateral filtering for gray and color images, in: International Conference on Computer Vision, IEEE, 1998, pp. 839–846.
- [23] M. Elad, On the origin of the bilateral filter and ways to improve it, IEEE Transactions on Image Processing 11 (10) (2002) 1141–1151.
- [24] V. Khryashchev, I. Apalkov, L. Shmaglit, A novel smart bilateral filter for ringing artifacts removal in JPEG2000 images, in: International Conference on Image Processing, Computer Vision and Pattern Recognition, 2010.
- [25] A. Nikitin, V. Solovyev, V. Khryashchev, A. Priorov, Adaptive bilateral filter for JPEG2000 deringing, in: IEEE International Conference on Image Processing, Computer Vision and Pattern Recognition, 2011.

- [26] G. Zhai, J. Cai, W. Lin, X. Yang, W. Zhang, Image deringing using quadtree based block-shift filtering, in: International Symposium on Circuits and Systems, IEEE, 2008, pp. 708–711.
- [27] Y.Y. Chen, Y.W. Chang, W.C. Yen, Design of a de-ringing filter for wavelet-based compressed image, in: International Technical Conference on Circuits/Systems, Computers and Communications, 2008, pp. 1265–1268.
- [28] Y. Chen, S. Tai, C. Wang, K. Lin, Design of a filter against artifacts for JPEG2000, *Journal of Electronic Imaging* 14 (2005) 043002.
- [29] D. Sun, W.K. Cham, An effective postprocessing method for low bit rate block DCT coded images, in: IEEE International Conference on Acoustics, Speech, and Signal Processing, 2007, pp. 705–708.
- [30] R. Zhang, Y.L. Fong, W.K. Cham, Image deblocking by the dual adaptive FIR wiener filter and overcomplete representation, in: IEEE International Conference on Information, Communications and Signal Processing, 2009, pp. 1–4.
- [31] G. Zhai, W. Zhang, X. Yang, W. Lin, Y. Xu, Efficient image deblocking based on postfiltering in shifted windows, *IEEE Transactions on Circuits and Systems for Video Technology* 18 (1) (2008) 122–126.
- [32] Z. Li, E.J. Delp, Block artifact reduction using a transform-domain Markov random field model, *IEEE Transactions on Circuits and Systems for Video Technology* 15 (12) (2005) 1583–1593.
- [33] A.W.C. Liew, H. Yan, N.F. Law, POCS-based blocking artifacts suppression using a smoothness constraint set with explicit region modeling, *IEEE Transactions on Circuits and Systems for Video Technology* 15 (6) (2005) 795–800.
- [34] A. Foi, V. Katkovnik, K. Egiazarian, Pointwise shape-adaptive DCT for high-quality denoising and deblocking of grayscale and color images, *IEEE Transactions on Image Processing* 16 (5) (2007) 1395–1411.
- [35] H. Rabbani, M. Vafadoost, I. Selesnick, S. Gazor, Image denoising based on a mixture of bivariate gaussian models in complex wavelet domain, in: IEEE International Summer School on Medical Devices and Biosensors, 2006, pp. 149–153.
- [36] H. Rabbani, M. Vafadoost, I. Selesnick, S. Gazor, Image denoising employing a mixture of circular symmetric Laplacian models with local parameters in complex wavelet domain, in: IEEE International Conference on Acoustics, Speech and Signal Processing, vol. 1, 2007, pp. 805–808.
- [37] E. Le Pennec, S. Mallat, Sparse geometric image representations with bandelets, *IEEE Transactions on Image Processing* 14 (4) (2005) 423–438.
- [38] J.L. Starck, E.J. Candes, D.L. Donoho, The curvelet transform for image denoising, *IEEE Transactions on Image Processing* 11 (6) (2002) 670–684.
- [39] M. Elad, M. Aharon, Image denoising via sparse and redundant representations over learned dictionaries, *IEEE Transactions on Image Processing* 15 (12) (2006) 3736–3745.
- [40] K. Dabov, A. Foi, V. Katkovnik, K. Egiazarian, Image denoising by sparse 3-D transform-domain collaborative filtering, *IEEE Transactions on Image Processing* 16 (8) (2007) 2080–2095.
- [41] K. Dabov, A. Foi, V. Katkovnik, K. Egiazarian, Image restoration by sparse 3D transform-domain collaborative filtering, in: SPIE Electronic Imaging, Citeseer, 2008.
- [42] X. Li, Fine-granularity and spatially-adaptive regularization for projection-based image deblurring, *IEEE Transactions on Image Processing* 99 (2011). 1–1.
- [43] J. Guerrero-Colón, L. Mancera, J. Portilla, Image restoration using space-variant gaussian scale mixtures in overcomplete pyramids, *IEEE Transactions on Image Processing* 17 (1) (2008) 27–41.
- [44] J.A. Guerrero-Colón, J. Portilla, Deblurring-by-denoising using spatially adaptive gaussian scale mixtures in overcomplete pyramids, in: IEEE International Conference on Image Processing, 2006, pp. 625–628.
- [45] J. Bardsley, S. Jefferies, J. Nagy, R. Plemmons, Blind iterative restoration of images with spatially-varying blur, *Optics Express* 14 (2) (2006) 1767–1782.
- [46] E.P. Simoncelli, B.A. Olshausen, Natural image statistics and neural representation, *Annual Review of Neuroscience* 24 (1) (2001) 1193–1216.
- [47] H.R. Sheikh, M.F. Sabir, A.C. Bovik, A statistical evaluation of recent full reference image quality assessment algorithms, *IEEE Transactions on Image Processing* 15 (11) (2006) 3440–3451.
- [48] A. Nosratinia, Enhancement of JPEG-compressed images by re-application of JPEG, *The Journal of VLSI Signal Processing* 27 (1) (2001) 69–79.
- [49] T. Wang, G. Zhai, JPEG2000 image postprocessing with novel trilateral deringing filter, *Optical Engineering* 47 (2008) 027005.
- [50] A. Nosratinia, Postprocessing of JPEG-2000 images to remove compression artifacts, *IEEE Signal Processing Letters* 10 (10) (2003) 296–299.
- [51] Z. Wang, E.P. Simoncelli, A.C. Bovik, Multi-scale structural similarity for image quality assessment, in: Proceedings of IEEE Asilomar Conf on Signals, Systems, and Computers, Asilomar, 2003.
- [52] D.L. Donoho, De-noising by soft-thresholding, *IEEE Transactions on Information Theory* 41 (3) (1995) 613–627.
- [53] P. Meer, Robust techniques for computer vision, *Emerging Topics in Computer Vision* (2004) 107–190.
- [54] D. Martin, C. Fowlkes, D. Tal, J. Malik, A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics, in: Proceedings of 8th International Conference on Computer Vision, vol. 2, 2001, pp. 416–423.
- [55] D.S. Taubman, M.W. Marcellin, JPEG2000: Image Compression Fundamentals, Standards, and Practice, Kluwer Academic Publishers, 2001.
- [56] V. Vapnik, *The Nature of Statistical Learning Theory*, Springer Verlag, 2000.
- [57] C. Chang, C. Lin, LIBSVM: a Library for Support Vector Machines. URL (<http://www.csie.ntu.edu.tw/~cjlin/libsvm/>), 2001.
- [58] H.R. Sheikh, Z. Wang, L.K. Cormack, A.C. Bovik, LIVE Image Quality Assessment Database. URL (<http://live.ece.utexas.edu/research/quality>), 2007.
- [59] A. Mittal, A.K. Moorthy, A.C. Bovik, Automatic parameter prediction for image denoising algorithms using perceptual quality features, in: SPIE Proceedings Human Vision and Electronic Imaging, 2012.
- [60] H.R. Sheikh, A.C. Bovik, Image information and visual quality, *IEEE Transactions on Image Processing* 15 (2) (2006) 430–444.
- [61] V. Kayargadde, J. Martens, Perceptual characterization of images degraded by blur and noise: model, *Journal of the Optical Society of America A* 13 (1996) 1178–1188.
- [62] S.S. Channappayya, A.C. Bovik, C. Caramanis, R.W. Heath, Design of linear equalizers optimized for the structural similarity index, *IEEE Transactions on Image Processing* 17 (6) (2008) 857–872.
- [63] S.S. Channappayya, A.C. Bovik, R.W. Heath, A linear estimator optimized for the structural similarity index and its application to image denoising, in: IEEE International Conference on Image Processing, IEEE, 2006, pp. 2637–2640.
- [64] S.S. Channappayya, A.C. Bovik, R.W. Heath, Rate bounds on SSIM index of quantized images, *IEEE Transactions on Image Processing* 17 (9) (2008) 1624–1639.
- [65] Y.H. Huang, T.S. Ou, P.Y. Su, H.H. Chen, Perceptual rate-distortion optimization using structural similarity index as quality metric, *IEEE Transactions on Circuits and Systems for Video Technology* 20 (11) (2010) 1614–1624.
- [66] C. Yim, A.C. Bovik, Quality assessment of deblocked images, *IEEE Transactions on Image Processing* 20 (1) (2011) 88–98.