

Impact of deep learning-based image super-resolution on binary signal detection

Xiaohui Zhang,^{a,†} Varun A. Kelkar,^{b,†} Jason Granstedt,^c Hua Li,^{a,d,e}
and Mark A. Anastasio^{© a,b,c,*}

^aUniversity of Illinois at Urbana–Champaign, Department of Bioengineering, Urbana, Illinois, United States

^bUniversity of Illinois at Urbana–Champaign, Department of Electrical and Computer Engineering, Urbana, Illinois, United States

^cUniversity of Illinois at Urbana–Champaign, Department of Computer Science, Urbana, Illinois, United States

^dUniversity of Illinois at Urbana–Champaign, Cancer Center at Illinois, Urbana, Illinois, United States

^eCarle Foundation Hospital, Carle Cancer Center, Urbana, Illinois, United States

Abstract

Purpose: Deep learning-based image super-resolution (DL-SR) has shown great promise in medical imaging applications. To date, most of the proposed methods for DL-SR have only been assessed using traditional measures of image quality (IQ) that are commonly employed in the field of computer vision. However, the impact of these methods on objective measures of IQ that are relevant to medical imaging tasks remains largely unexplored. We investigate the impact of DL-SR methods on binary signal detection performance.

Approach: Two popular DL-SR methods, the super-resolution convolutional neural network and the super-resolution generative adversarial network, were trained using simulated medical image data. Binary signal-known-exactly with background-known-statistically and signal-known-statistically with background-known-statistically detection tasks were formulated. Numerical observers (NOs), which included a neural network-approximated ideal observer and common linear NOs, were employed to assess the impact of DL-SR on task performance. The impact of the complexity of the DL-SR network architectures on task performance was quantified. In addition, the utility of DL-SR for improving the task performance of suboptimal observers was investigated.

Results: Our numerical experiments confirmed that, as expected, DL-SR improved traditional measures of IQ. However, for many of the study designs considered, the DL-SR methods provided little or no improvement in task performance and even degraded it. It was observed that DL-SR improved the task performance of suboptimal observers under certain conditions.

Conclusions: Our study highlights the urgent need for the objective assessment of DL-SR methods and suggests avenues for improving their efficacy in medical imaging applications.

© The Authors. Published by SPIE under a Creative Commons Attribution 4.0 International License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.JMI.8.6.065501](https://doi.org/10.1117/1.JMI.8.6.065501)]

Keywords: deep learning-based image super-resolution; objective image quality assessment; numerical observers; Rayleigh detection task.

Paper 21178R received Jul. 2, 2021; accepted for publication Oct. 27, 2021; published online Nov. 16, 2021.

*Address all correspondence to Mark A. Anastasio, maa@illinois.edu

†These authors contributed equally to this work.

1 Introduction

Single-image super-resolution (SISR) is a classic image restoration operation that seeks to estimate a high-resolution (HR) image from an observed low-resolution (LR) one.¹ A variety of methods have been developed to achieve this goal, such as filtering and interpolation-based approaches² and more formal regularized inverse problem-based formulations,^{3,4} to name a few. Recently, deep learning-based image super-resolution (DL-SR) methods have been widely employed and have shown great promise for SISR in terms of traditional image quality (IQ) metrics such as mean square error (MSE), structural similarity index metric (SSIM), and peak-signal-to-noise ratio (PSNR).^{5–8}

In medical imaging, images are often acquired for specific purposes, and the use of objective measures of IQ is widely advocated for assessing imaging systems and image processing algorithms.^{9–15} Although DL-SR algorithms can improve traditional IQ metrics,^{16–21} it is well-known that such metrics may not always correlate with objective task-based IQ measures.^{22–25} Despite this, relatively few studies have objectively assessed image super-resolution methods.^{19,26–28} Dai et al.²⁷ evaluated six image super-resolution methods on popular vision tasks such as edge detection and semantic image segmentation and found that the standard perceptual metrics correlated well with the usefulness of image super-resolution to these tasks. Jaffe et al.²⁸ conducted a study in which the aesthetic IQ that DL-SR methods sought to improve did not necessarily increase classification accuracy. However, none of these studies were carried out with images, tasks, or observers relevant to medical imaging. Additionally, the data processing inequality indicates that the performance of an ideal observer (IO) on a particular task cannot be improved using image processing transformations.²⁹ The scenarios under which DL-SR may improve the performance of a suboptimal observer on a specified task have not been thoroughly investigated. The purpose of this work is to evaluate DL-SR methods using task-based measures as a preliminary attempt to address the issues raised above. For this study, two canonical DL-SR networks were identified for the analysis. A variety of mathematical and learning-based numerical observers (NOs) were computed on the HR images, the LR images, and the images resolved by the DL-SR methods. Receiver operating characteristics (ROC) analysis was employed to quantify the performance of these NOs. Two stylized binary signal detection tasks were designed to evaluate the DL-SR networks systematically and comprehensively under known statistical conditions. Specifically, a signal-known-exactly and background-known-statistically (SKE/BKS) Rayleigh discrimination task^{30,31} was employed to assess the ability of a DL-SR to resolve two small adjacent objects. The inherent detectability of the signal was varied, and its effect on the utility of DL-SR for improving detection task performance was studied. The impact of the depth of a DL-SR network on NO performance was investigated to see if the deep learning mantra “deeper is better” holds true for signal detection performance.³² Additionally, a signal-known-statistically and background-known-statistically (SKS/BKS) microcalcification (MC) cluster detection task was employed to investigate under what circumstances DL-SR techniques may improve the binary signal detection performance of a suboptimal observer.

The remainder of this paper is organized as follows. Section 2 describes the relevant background on linear imaging systems, the basic theory relating to binary signal detection tasks, NOs, and DL-SR. Section 3 describes the setup for the numerical studies, and Sec. 4 describes the results of the proposed evaluation. Section 5 presents a discussion on the salient findings, and Sec. 6 concludes this paper.

2 Background

Many imaging systems are approximately described by a continuous-to-discrete (C-D) linear imaging model:⁹

$$\mathbf{g} = \mathcal{H}f(\mathbf{r}) + \mathbf{n}, \quad (1)$$

where $f \in \mathbb{L}_2(\mathbb{R}^d)$ is the true object of interest that is a function of the d -dimensional spatio-temporal coordinate \mathbf{r} and $\mathbf{g} \in \mathbb{E}^m$ is a vector that describes the measurement data. The mapping $\mathcal{H}: \mathbb{L}_2(\mathbb{R}^d) \rightarrow \mathbb{E}^m$ denotes the C-D forward operator that represents the data-acquisition process,

and $\mathbf{n} \in \mathbb{E}^m$ denotes the measurement noise. In practice, discrete-to-discrete (D-D) models for the imaging system are often employed, in which case the object $f(\mathbf{r})$ is approximated by a vector $\mathbf{f} \in \mathbb{E}^n$, $n \in \mathbb{N}$ and a D-D approximation $\mathbf{H} \in \mathbb{E}^{m \times n}$ is employed in place of \mathcal{H} .⁹

2.1 Binary Signal Detection Tasks

A binary signal detection task requires an observer to classify the image as satisfying either hypothesis H_0 or hypothesis H_1 :

$$H_0: \mathbf{g} = \mathbf{H}\mathbf{f}_0 + \mathbf{n} = \mathbf{H}(\mathbf{f}_b + \mathbf{f}_{s_0}) + \mathbf{n}, \quad (2)$$

$$H_1: \mathbf{g} = \mathbf{H}\mathbf{f}_1 + \mathbf{n} = \mathbf{H}(\mathbf{f}_b + \mathbf{f}_{s_1}) + \mathbf{n}, \quad (3)$$

where $\mathbf{f}_b \in \mathbb{E}^n$ denotes the background, $\mathbf{f}_{s_0} \in \mathbb{E}^n$ and $\mathbf{f}_{s_1} \in \mathbb{E}^n$ represent the signal under the two hypotheses, $\mathbf{H} \in \mathbb{E}^{m \times n}$ refers to the D-D imaging operator, and $\mathbf{n} \in \mathbb{E}^m$ denotes the measurement noise. The special case of $\mathbf{f}_{s_0} = \mathbf{0}$ corresponds to a task of detecting the presence or absence of the signal \mathbf{f}_{s_1} in an image. When \mathbf{f}_b is a random vector drawn from a certain nondegenerate distribution and \mathbf{f}_{s_0} and \mathbf{f}_{s_1} are fixed known signals, the detection task is known as a SKE/BKS detection task. Alternatively, if \mathbf{f}_{s_0} and \mathbf{f}_{s_1} are also random, then the detection task is known as a signal-known-statistically and background-known-statistically (SKS/BKS) detection task. Both of these tasks are considered in this work.

2.2 Numerical Observers for IQ Assessment

A NO for a signal detection task maps a given set of measurements \mathbf{g} or, alternatively, an image estimate $\hat{\mathbf{f}} \in \mathbb{E}^n$ of the object obtained from \mathbf{g} to a scalar test statistic t that is used to determine whether \mathbf{g} or $\hat{\mathbf{f}}$ satisfies H_0 or H_1 based on comparison with a predetermined threshold τ . The NOs employed in this study are described below.

2.2.1 Ideal observer and ResNet-based observer

The IO is an observer that utilizes all available statistical information about the task at hand to maximize task performance. An IO test statistic $t_{IO}(\hat{\mathbf{f}})$ is any monotonic function of the likelihood ratio:⁹

$$\Lambda(\hat{\mathbf{f}}) = \frac{p(\hat{\mathbf{f}}|H_1)}{p(\hat{\mathbf{f}}|H_0)}, \quad (4)$$

where $p(\hat{\mathbf{f}}|H_0)$ and $p(\hat{\mathbf{f}}|H_1)$ are the conditional probability density functions that describe image estimate $\hat{\mathbf{f}}$ under hypotheses H_0 and H_1 . The exact computation of an IO test statistic based on $\Lambda(\hat{\mathbf{f}})$ is intractable in general, and Markov-chain Monte Carlo techniques have been proposed to approximate it.^{33,34} Recently, it has been empirically shown that the IO can be approximated by a neural network-based observer.¹⁴ In this study, a residual neural network-based (ResNet-based) classifier of sufficient capacity trained on a large labeled training dataset was employed to approximate the IO. This will henceforth be referred to as the ResNet-IO. Note that, if this network does not possess the capacity to accurately approximate $t_{IO}(\hat{\mathbf{f}})$, the resulting NO will be simply referred to as a ResNet-based observer. In this case, the ResNet-based observer is a suboptimal observer.

2.2.2 Hotelling observer and regularized Hotelling observer

The Hotelling observer (HO) is the optimal NO under the condition that the employed test statistic is a linear function of the data.⁹ The test statistic for the HO is defined as

$$t_{\text{HO}}(\hat{\mathbf{f}}) = \mathbf{w}_{\text{HO}}^{\top} \hat{\mathbf{f}}, \quad (5)$$

where

$$\mathbf{w}_{\text{HO}} = \mathbf{K}(\hat{\mathbf{f}})^{-1} \Delta \bar{\mathbf{f}} \quad (6)$$

is known as the Hotelling template and

$$\mathbf{K}(\hat{\mathbf{f}}) = \frac{1}{2} (\mathbf{K}_0(\hat{\mathbf{f}}) + \mathbf{K}_1(\hat{\mathbf{f}})). \quad (7)$$

Here, $\mathbf{K}_0(\hat{\mathbf{f}})$ and $\mathbf{K}_1(\hat{\mathbf{f}})$ denote the covariance matrices of $\hat{\mathbf{f}}$ under the hypotheses H_0 and H_1 , respectively, and $\Delta \bar{\mathbf{f}} = \mathbb{E}(\hat{\mathbf{f}}|H_1) - \mathbb{E}(\hat{\mathbf{f}}|H_0)$ is the difference between the condition mean of $\hat{\mathbf{f}}$ under the two hypotheses.

In some cases, the covariance matrix $\mathbf{K}(\hat{\mathbf{f}})$ can be ill-conditioned, and therefore its inverse cannot be stably computed. To address this, a regularized Hotelling observer (RHO) is employed. The singular value decomposition of \mathbf{K} is written as

$$\mathbf{K} = \sum_{i=1}^R \sigma_i \mathbf{v}_i \mathbf{u}_i^{\dagger}, \quad (8)$$

where R is the rank of \mathbf{K} , $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_R$ are the singular values of \mathbf{K} , \mathbf{u}_i and \mathbf{v}_i are the right and left singular vectors, respectively, and \dagger denotes the complex conjugate transpose operation. The truncated pseudoinverse $\mathbf{K}_{\lambda}^{\dagger}$ of \mathbf{K} is employed as a stable approximation of \mathbf{K}^{-1} :

$$\mathbf{K}_{\lambda}^{\dagger} = \sum_{i=1}^P \frac{1}{\sigma_i} \mathbf{u}_i \mathbf{v}_i^{\dagger}, \quad (9)$$

where λ is a threshold for singular value and P is chosen to satisfy $\sigma_P \geq \lambda \sigma_1 > \sigma_{P+1}$. The truncated pseudoinverse is then used to construct the RHO template, which is then used to obtain the RHO test statistic:

$$t_{\text{RHO}}(\hat{\mathbf{f}}) = \mathbf{w}_{\text{RHO}}(\lambda)^{\top} \hat{\mathbf{f}} = (\mathbf{K}_{\lambda}^{\dagger} \Delta \bar{\mathbf{f}})^{\top} \hat{\mathbf{f}}. \quad (10)$$

2.2.3 Gabor channelized Hotelling observer

To compute a channelized Hotelling observer (CHO) template, the image data $\hat{\mathbf{f}}$ is first transformed into a vector $\mathbf{v} \in \mathbb{E}^q$, $q < n$, known as the channel output, via a transformation $\mathbf{v} = \mathbf{T} \hat{\mathbf{f}}$, where $\mathbf{T} \in \mathbb{E}^{q \times n}$ is known as the channel matrix. The test statistic of CHO is then computed as

$$t_{\text{CHO}}(\hat{\mathbf{f}}) = \mathbf{w}_{\text{CHO}}^{\top} \mathbf{v}, \quad (11)$$

where $\mathbf{w}_{\text{CHO}} = \mathbf{K}_{\mathbf{v}}^{-1} \Delta \bar{\mathbf{v}}$ and $\mathbf{K}_{\mathbf{v}} = \frac{1}{2} (\mathbf{K}_{\mathbf{v},0} + \mathbf{K}_{\mathbf{v},1})$ is the covariance matrix of the channelized image data. Here $\mathbf{K}_{\mathbf{v},0}$ and $\mathbf{K}_{\mathbf{v},1}$ denote the covariance matrices of \mathbf{v} under the two hypotheses H_0 and H_1 . The CHO with Gabor channels (Gabor CHO) can be considered an anthropomorphic observer.^{9,35-37} The channel matrix \mathbf{T} employed in the Gabor CHO is specified as follows. A Gabor function C_i corresponding to the i 'th row of \mathbf{T} is defined in the spatial domain by multiplying a sinusoidal wave with a Gaussian function:

$$C_i(x, y) = \exp\left(-4 \ln 2 \frac{x^2 + y^2}{w_i^2}\right) \cos[2\pi\nu_i(x \cos \theta_i + y \sin \theta_i) + \phi_i], \quad (12)$$

where w_i is the channel width, ν_i is the central frequency, θ_i is the orientation, and ϕ_i is the phase. The element v_i of the channel vector $\mathbf{v} = \mathbf{T} \hat{\mathbf{f}}$ is then given by the scalar product of the discretized version of C_i with the 2D image representation of $\hat{\mathbf{f}}$.

2.3 Deep Learning-Based Image Super-Resolution

In the context of an image super-resolution problem, an LR image $\mathbf{f}_{\text{LR}} \in \mathbb{E}^{n'}$, $n' \in \mathbb{N}$, $n' \leq n$ can be formally thought of as being related to the sought-after HR image $\mathbf{f}_{\text{HR}} \in \mathbb{E}^n$ via the following equation:

$$\mathbf{f}_{\text{LR}} = \mathbf{H}_{\text{blur}} \mathbf{f}_{\text{HR}} + \mathbf{n}, \quad (13)$$

where $\mathbf{H}_{\text{blur}} \in \mathbb{E}^{n' \times n}$ represents a degradation operator that removes the higher spatial frequencies from \mathbf{f}_{HR} and \mathbf{n} denotes the noise. Given a specific LR image, an estimate $\mathbf{f}_{\text{SR}} \in \mathbb{E}^n$ of the original HR image is obtained using image super-resolution methods. However, this is a challenging ill-posed inverse problem. In recent years, deep learning has been widely applied to achieve image super-resolution.⁵⁻⁸ A popular class of deep learning-based approaches calls for establishing a mapping from the space of LR images to the space of HR images:

$$\mathbf{f}_{\text{SR}} = S_{\theta}(\mathbf{f}_{\text{LR}}), \quad (14)$$

where S_{θ} is a deep neural network parametrized by θ . For several supervised learning approaches, a training dataset of size D consisting of paired LR and HR images, $\{(\mathbf{f}_{\text{LR}}^{(i)}, \mathbf{f}_{\text{HR}}^{(i)})\}_{i=1}^D$, is utilized. A loss function is constructed based on a distance metric $\mathcal{L}(S_{\theta}(\mathbf{f}_{\text{LR}}^{(i)}), \mathbf{f}_{\text{HR}}^{(i)})$ between a super-resolved (SR) image and an HR image, and the optimal parameters $\hat{\theta}$ are estimated by approximately minimizing the loss function over the dataset:

$$\hat{\theta} = \arg \min_{\theta} \frac{1}{D} \sum_{i=1}^D \mathcal{L}(S_{\theta}(\mathbf{f}_{\text{LR}}^{(i)}), \mathbf{f}_{\text{HR}}^{(i)}). \quad (15)$$

Various loss functions such as ℓ_1 or ℓ_2 loss, or a perceptual loss,³⁸ can be used to define \mathcal{L} . Additionally, an adversarial loss that attempts to match the distribution of SR images to the distribution of original HR images can also be employed.⁸ The two DL-SR networks considered in this study are the super-resolution convolutional neural network (SRCNN)⁶ and the super-resolution generative adversarial network (SRGAN).⁸

The architectures of these two networks are shown in Fig. 1. The architecture of the SRCNN consists of feed-forward convolutional layers interspersed with pointwise rectified linear unit (ReLU) nonlinearities.^{6,39} The SRGAN architecture consists of a generative network, which is an image-to-image mapping network consisting of convolutional residual blocks interspersed with pointwise ReLU nonlinearities. A discriminator network is jointly trained along with the generative network and provides the adversarial loss for matching the distribution of generated SR images to the distribution of HR images.⁸

3 Numerical Studies

Computer-simulation studies were employed to objectively evaluate the DL-SR methods described above with two binary signal detection tasks: (i) a Rayleigh detection task and (ii) an MC cluster detection task. The NOs described in Sec. 2.2 were computed on the SR images, as well as the LR and true HR images, to objectively assess the impact of DL-SR on the considered tasks.

3.1 Clustered Lumpy Background

The CLB model was developed by Bochud et al.⁴⁰ to generate random backgrounds that resemble mammographic textures. The value of a CLB image at position \mathbf{r} is

$$\mathbf{f}_b(\mathbf{r}) = \sum_{k=1}^K \sum_{n=1}^{N_k} l(\mathbf{r} - \mathbf{r}_k - \mathbf{r}_{kn}, \mathbf{R}_{\theta_{kn}}), \quad \text{where } l(\mathbf{r}, \mathbf{R}_{\theta}) = \exp\left(-\alpha \frac{\|\mathbf{R}_{\theta} \mathbf{r}\|^{\beta}}{L(\mathbf{R}_{\theta} \mathbf{r})}\right). \quad (16)$$

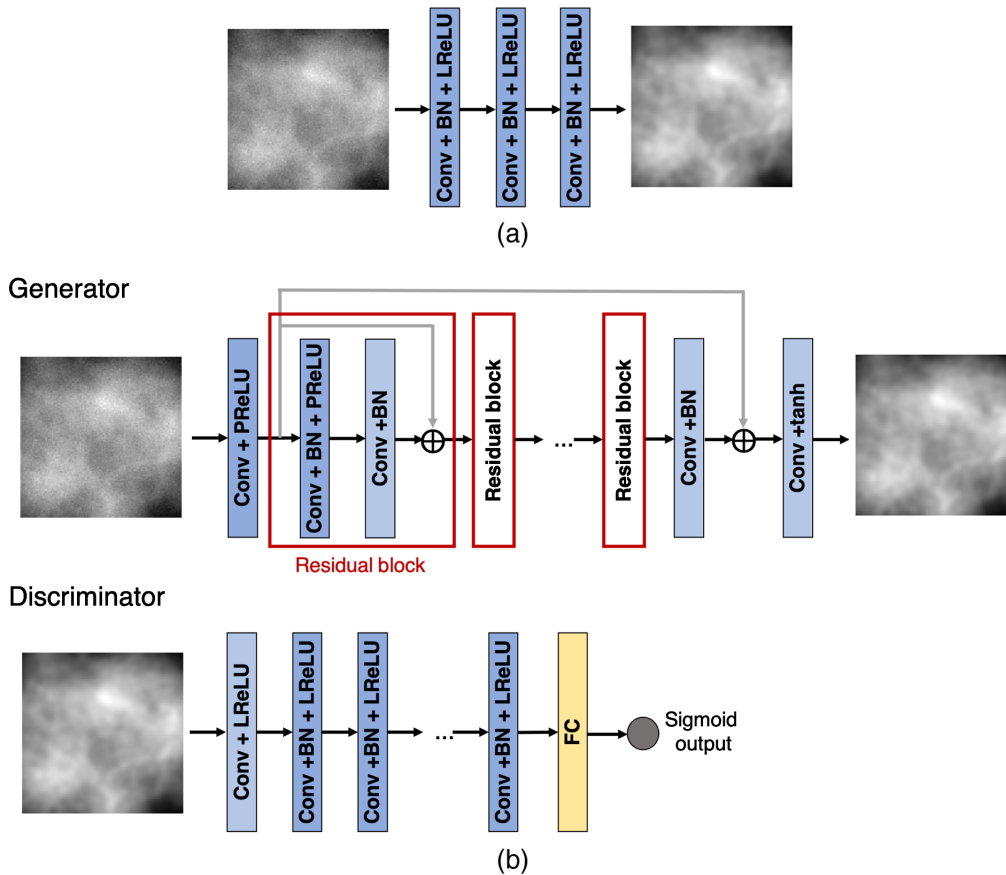


Fig. 1 Architecture of the super-resolution networks employed in our study: (a) SRCNN and (b) SRGAN.

Table 1 Parameters for generating CLB images.

\bar{K}	\bar{N}	L_x	L_y	α	β	σ
150	20	5	2	2.1	0.5	12

Here $l(\mathbf{r}, \mathbf{R}_\theta)$ is known as the blob function. The integer K denotes the number of clusters that was sampled from a Poisson distribution with a mean of $\bar{K}: K \sim \text{Poiss}(\bar{K})$, N_k specifies the number of blobs in the k 'th cluster sampled from a Poisson distribution with the mean of $\bar{N}: N \sim \text{Poiss}(\bar{N})$, \mathbf{r}_k indicates the center location of the k 'th cluster sampled uniformly over the field of view, and \mathbf{r}_{kn} represents the center location of the n 'th blob in the k 'th cluster sampled from a Gaussian distribution with the center of \mathbf{r}_k and standard deviation of σ . The matrix $\mathbf{R}_{\theta_{kn}}$ represents the rotation corresponding to the angle θ_{kn} sampled from a uniform distribution between 0 and 2π , $L(\mathbf{r})$ refers to the radius of the ellipse with half-axes L_x and L_y , and α and β are adjustable coefficients. The parameters of the CLB model employed in both the Rayleigh detection task and MC cluster detection task are shown in Table 1.

3.2 Rayleigh Detection Task with a Clustered Lumpy Background Model

The Rayleigh detection task is a natural task for assessing the resolution properties of imaging systems and has been employed previously for optimizing tomographic imaging systems.^{30,31} This is a binary signal detection task, in which hypothesis H_0 corresponds to a signal \mathbf{f}_{s_0} consisting of two adjacent point objects and hypothesis H_1 corresponds to a signal \mathbf{f}_{s_1} consisting of a single-line object.

3.2.1 Simulated image data for Rayleigh detection task

Given the definition of signals \mathbf{f}_{s_0} and \mathbf{f}_{s_1} provided above, the generation of LR images under H_0 and H_1 is written as

$$H_0: \mathbf{f}_{LR} = \mathbf{H}_{\text{blur}}\mathbf{f}_0 + \mathbf{n} \equiv \mathbf{H}_{\text{blur}}(\mathbf{f}_b + \mathbf{f}_{s_0}) + \mathbf{n}, \quad (17)$$

$$H_1: \mathbf{f}_{LR} = \mathbf{H}_{\text{blur}}\mathbf{f}_1 + \mathbf{n} \equiv \mathbf{H}_{\text{blur}}(\mathbf{f}_b + \mathbf{f}_{s_1}) + \mathbf{n}, \quad (18)$$

where \mathbf{f}_b denotes a CLB image of size 128×128 with parameters defined in Table 1 and \mathbf{n} denotes the measurement noise. Given an adjustable parameter L , termed the signal length, \mathbf{f}_{s_0} is specified by first defining two Kronecker delta functions separated by a distance of $L - 2$, and convolving them with a Gaussian function of standard deviation 1.375 pixels. The signal \mathbf{f}_{s_1} is specified by first defining a horizontal line of length L , which is subsequently convolved with the same Gaussian function. The signals are inserted such that the centers of the signals coincide with the center of the image. The Rayleigh detection task was performed independently on the following datasets, where the HR dataset consists of images of the type:

$$\mathbf{f}_{HR} = \mathbf{f}_i + \mathbf{n}, \quad i = 0, 1, \quad (19)$$

the LR dataset consists of images of the type

$$\mathbf{f}_{LR} = \mathbf{H}_{\text{blur},1}\mathbf{f}_i + \mathbf{n}, \quad i = 0, 1, \quad (20)$$

and the SR dataset consists of images of type

$$\mathbf{f}_{SR} = S(\mathbf{H}_{\text{blur},1}\mathbf{f}_i + \mathbf{n}), \quad i = 0, 1, \quad (21)$$

where $\mathbf{H}_{\text{blur},1}$ represents a Gaussian filter with a standard deviation of 1.5 pixels and $\mathbf{f}_i = \mathbf{f}_b + \mathbf{f}_{s_i}$, as defined in Eq. (17). Here, S denotes the DL-SR operation performed by either the SRCNN or the SRGAN, and \mathbf{n} denotes the sum of pixel-wise independent and identically distributed (IID) Poisson noise with a standard deviation scaled by $\sigma_p = 0.013$ and IID. Gaussian noise with a standard deviation $\sigma_g = 0.35$. The simulation of an example LR image according to the described procedure is shown in Fig. 2.

Two separate studies were formulated based on the Rayleigh detection task.

1. *Signal length variation study.* In this study, the signal length parameter L , which pertains to the distance between the two point objects in \mathbf{f}_{s_0} or the length of the line in \mathbf{f}_{s_1} , was varied to investigate the resolving power of the DL-SR algorithms. The signal lengths of $L = \{5, 6, 7, 8, 9\}$ were employed in this study as shown in Fig. 3.
2. *Network complexity variation study.* To investigate how the DL-SR network complexity correlates with the task performance for a fixed object model and task design, a network complexity variation study in which the number of layers of a DL-SR network was varied




Fig. 2 CLB (\mathbf{f}_b), signals (\mathbf{f}_{s_0} , \mathbf{f}_{s_1}), and combined images of the Rayleigh detection task.




Fig. 3 Example signal image ROIs with respect to different signal lengths.

was conducted. The SRGAN employs an additional tunable parameter controlling the trade-off between the MSE loss and the discriminative loss, the optimal value of which may depend, among other factors, on the number of layers in the network. Hence, only SRCNN was employed in this study.

3.2.2 Training details for the DL-SR networks

For the signal length variation study, both the SRCNN and SRGAN were trained and evaluated. The training and validation data for SRCNN consisted of 5000 and 625 class-balanced signal present/absent images, respectively. For SRGAN training, due to more trainable parameters in the SRGAN, 20,000 images were used for training, and 2000 images were used for validation, respectively. Examples of HR, LR, and SR images produced by the networks are shown in Fig. 4(a).

For the architecture variation study, seven SRCNNs with varying numbers of convolutional layers ranging from 2 to 8 were employed. For all of the SRCNNs, the filter size in the first layer was fixed to 9×9 , whereas the filter size for the other layers was fixed to 5×5 . The number of filters in all layers was fixed to 32, except the last layer, in which the number of filters was fixed to 1. All SRCNNs were trained on 15,000 images and validated on 3000 images with class balance.

The SRCNN was trained with an MSE loss, and the SRGAN was trained using an MSE loss and an adversarial loss. All DL-SR networks to be evaluated in the Rayleigh detection task were trained on mini-batches at each iteration using the Adam optimizer.⁴¹ The DL-SR models that achieved the best performance on the validation set were used for evaluation. Both DL-SR networks were implemented under the TensorFlow 2.0 framework and trained on NVIDIA GPUs.

3.3 Microcalcification Cluster Detection Task with a Clustered Lumpy Background Model

Motivated by the clinical value of detecting MC clusters in mammograms that may be associated with malignancy in breast lesions,^{42,43} a stylized SKS/BKS binary signal detection task of




Fig. 4 Examples of (a) HR, LR, and SR images from SRCNN and SRGAN in Rayleigh detection task and (b) HR, LR, and SR images from SRCNN in MC cluster detection task.

identifying an image with or without an MC cluster present was studied. The objective of this study was to determine how the capacity of a NO affects observer performance on SR images. In essence, whether or not SR aids the performance of suboptimal observers was systematically studied.

3.3.1 Simulated image data for MC cluster detection task

The HR MC cluster dataset was created as follows. First, 128×128 CLB images were created to simulate the mammographic backgrounds, as described in Sec. 3.1. The signal-absent HR images \mathbf{f}_0 correspond to the case in which $\mathbf{f}_{s_0} = 0$ and, hence, were kept equal to the CLB images. The signal insertion pipeline employed to generate the signal-present HR image \mathbf{f}_1 is described as follows. A set of eleven 200×200 MC clusters segmented from digital mammograms acquired with the Selenia Dimensions system (Hologic, Inc.), available at <https://github.com/LAVI-USP/MCInsertionPackage>,⁴⁴ were employed to model the MC cluster signal. First, one out of the eleven segmented MC clusters was chosen at random, and a random rotation between 0 deg and 360 deg with zero padding was applied. Next, this rotated image \mathbf{s}_{MC} was cropped to a size of 128×128 and inserted into a CLB \mathbf{f}_b as⁴⁵

$$\mathbf{f}_1 = \mathbf{f}_b(c\mathbf{s}_{MC} + 1). \quad (22)$$

The scalar c represents a contrast factor uniformly sampled from the range [0.05, 0.06] that is chosen to visually match the contrast of real lesion.

Given the generated HR image, the corresponding LR image was simulated as follows, based on the degradation model described by You et al.:¹⁷

$$\mathbf{f}_{LR} = \mathbf{H}_{\text{blur},2}\mathbf{f}_i + \mathbf{n}, \quad i = 0,1. \quad (23)$$

Here $\mathbf{H}_{\text{blur},2}$ represents a Gaussian blurring operation with a standard deviation of 1.5 pixels, followed by downsampling by a factor of 2. Pixel-wise IID. Poisson noise with a standard deviation scaled by a factor $\sigma_p = 0.0001$ and IID. Gaussian noise with a standard deviation $\sigma_g = 0.001$ were added to both the HR and LR images. These noise values were chosen independently of the Rayleigh task so as to not saturate the observer performance on the LR images. To enable direct comparison with the HR and SR images, an additional operation \mathbf{U} representing upsampling by a factor of 2 was used on the LR images. Similar to the Rayleigh detection task, the MC cluster detection task was performed on the following datasets: (1) the HR dataset consisting of images of the type $\mathbf{f}_{HR} = \mathbf{f}_i + \mathbf{n}$, $i = 0,1$ is one of the MC cluster-absent/present hypotheses; (2) the LR dataset consisting of images of the type $\mathbf{f}_{LR} = \mathbf{H}_{\text{blur},2}\mathbf{f}_i + \mathbf{n}$, $i = 0,1$ along with the additional upsampling operation \mathbf{U} acting on \mathbf{f}_{LR} ; and (3) the SR datasets consisting of $\mathbf{f}_{SR} = S(\mathbf{U}\mathbf{f}_{LR})$, where S denotes the DL-SR operation performed by SRCNN.

3.3.2 Training details for DL-SR networks

The SRCNN employed in this study was trained on a dataset of 40,000 images and validated on a dataset of 4000 images, both with balanced classes. The network was trained with the Adam optimizer⁴¹ with a learning rate of 5×10^{-5} for 1000 epochs to minimize the MSE loss. The SRCNN model with the best validation performance was used. Examples of the SR images produced by the SRCNN along with the HR and the LR images are shown in Fig. 4(b).

3.4 Objective Evaluation of Deep Learning-Based Image Super-Resolution Networks

3.4.1 Objective evaluation metrics for the Rayleigh detection task

To evaluate the DL-SR networks with task-based metrics, three NOs, namely the RHO, Gabor CHO, and ResNet-IO, were employed. The test statistics for the three NOs were computed on the HR, LR, and SR images that were centrally cropped to a size of 64×64 . ROC curves were

computed, and the area under the ROC curve (AUC) was employed as a figure of merit. All evaluation metrics were computed on balanced test dataset of 40,000 images. Nonparametric estimation of the AUC confidence intervals was carried out using DeLong's algorithm,^{46,47} with the help of the pROC package in R.⁴⁸ Additionally, traditional IQ metrics such as PSNR and SSIM were computed on the LR and SR images.

To compute the RHO test statistic, 500,000 images containing two point objects and 500,000 images containing the line-shaped object were utilized to estimate the empirical covariance matrix $\mathbf{K}(\hat{\mathbf{f}})$. The threshold parameter λ in Eq. (9) was swept in from 10^{-9} to 10^{-4} , and the detection performance was evaluated on a validation set of 4000 class-balanced images. The value of λ that yielded the best RHO performance on the validation data was selected. This RHO with the selected parameter λ was applied to a test set consisting of 40,000 class-balanced images.

The channel matrix corresponding to the Gabor CHO comprised a set of 60 Gabor channels. Each Gabor channel was associated with one out of six passbands, one out of five orientations, and one out of two phases. The six passbands each have a spatial frequency bandwidth of 1 octave with a center frequency $\nu = 3/256, 3/128, 3/64, 3/32, 3/16$ and $3/8$ cycles/pixel. The five orientations were $0, 2\pi/5, 4\pi/5, 6\pi/5$, and $8\pi/5$, and the two phases were 0 and $\pi/2$. Examples of Gabor channel templates are shown in Fig. 5. The channelized covariance matrix was estimated using 100,000 images from each class with 500,000 noise realizations for each class.

The ResNet-IO, as shown in Fig. 6(a), was employed to approximate the IO test statistic. To obtain a good approximation of the IO using ResNets, the optimum network capacity needs to be determined empirically by sweeping the number of layers used in the ResNet architecture and choosing the configuration that gives the best detection performance. A large training dataset must be used to correctly represent the data distribution. Here the network was initialized with the help of the RHO template to give the best performance and to speed up convergence. A family of ResNets comprising various numbers of residual blocks were trained on a dataset consisting of 100,000 training images and validated on 4000 images from each of the two classes. The binary cross-entropy loss was minimized using Adam optimizer with a learning rate of 1×10^{-6} . Additionally, a "semionline learning" method in which the measurement noise was generated on-the-fly as described in Ref. 14 was utilized to mitigate the overfitting problem. The ResNet that had the best validation performance was chosen as the ResNet-IO.

3.4.2 Objective evaluation for the MC cluster detection task

As described previously, the objective of this study was to investigate the potential benefit of DL-SR as it relates to the capacity of an NO. A binary signal detection task was conducted to




Fig. 5 Examples of Gabor channel templates.

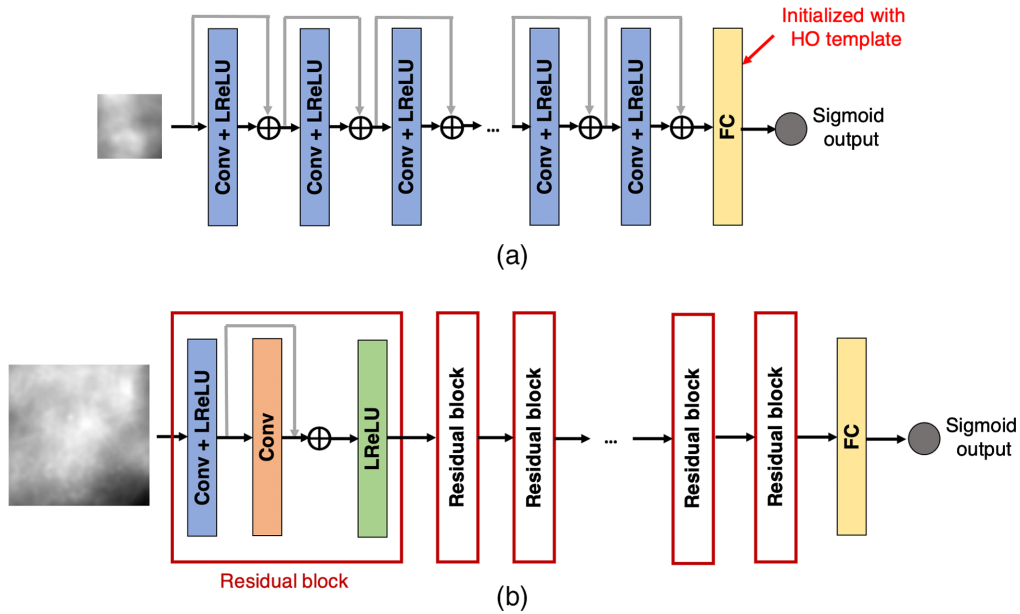


Fig. 6 Architectures of (a) ResNet-approximated IO for Rayleigh detection task and (b) ResNet-based observer for MC cluster detection task.

distinguish whether an image contains the MC cluster signal or not. To assess the task-based performance, a family of ResNet-based observers consisting of 2, 4, 6, or 8 residual blocks, respectively, were employed in the detection task. The architecture of the ResNet-based observers is shown in Fig. 6(b). Each of these observers was trained on class-balanced datasets of sizes 5000, 10,000, 20,000, 50,000, and 100,000 by minimizing the binary cross-entropy loss, until the detection capability of each observer was fulfilled. Each simulated MC cluster image in the training dataset was augmented four times by flipping. The AUC values produced by the trained ResNet-based observers on a held-out test set containing 20,000 images from each class were used to evaluate the signal detection performance. The ResNet-based observer that achieves the best test performance without further improvement with either a deeper network architecture or a larger training dataset could be considered an approximated IO.¹⁴

4 Results

4.1 Rayleigh Task

4.1.1 Impact of regularization on the Hotelling observer performance

In addition to introducing high-frequency features to an LR image, the DL-SR networks also suppress the per-pixel IID. noise added to the LR images. Due to this, the covariance matrix $\mathbf{K}(\hat{\mathbf{f}}_{\text{SR}})$ of the SR images is ill-conditioned. Hence, as mentioned in Sec. 2.2.2, regularization is needed to stably invert it to obtain the Hotelling template. Hence, the performance of the RHO depends upon the regularization parameter λ employed for truncating the singular values of \mathbf{K} . Figure 7 shows the Hotelling templates of the HR images, the LR images, and the images SR by the SRCNN and the SRGAN. It can be seen that, for low values of λ , the Hotelling template is noisy due to the unstable inversion of \mathbf{K} . On the other hand, for high values of λ , degradation of the signal specificity corresponding to the truncation of singular values can be seen.

4.1.2 Impact of signal length on observer performance

The traditional IQ metrics and AUC values for the signal length variation study computed on a class-balanced test set consisting of 40,000 images are plotted in Figs. 8 and 9, respectively.

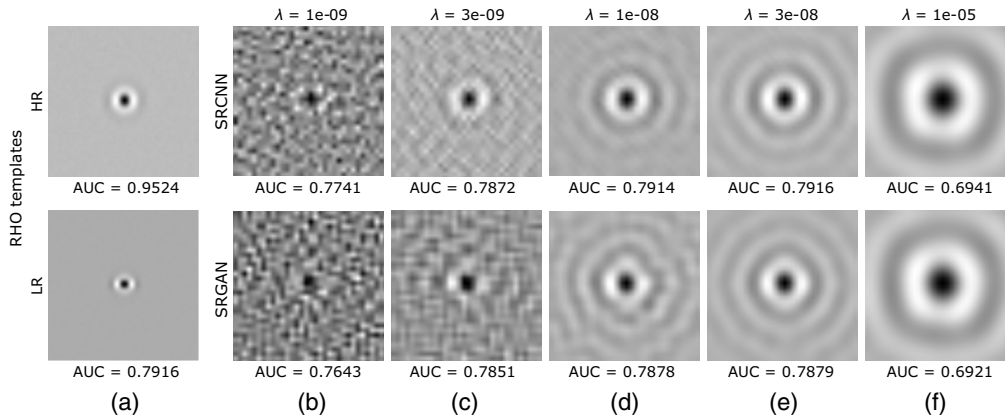


Fig. 7 RHO templates for the Rayleigh task of signal length 8 computed on (a) HR and LR images and (b)–(f) images from SRCNN and SRGAN resulting from sweeping the regularization parameter λ .

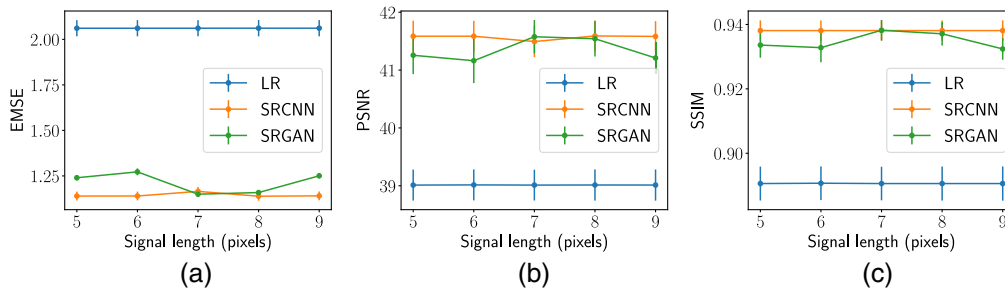


Fig. 8 Traditional IQ metrics including (a) ensemble MSE, (b) PSNR, and (c) SSIM of the HR, LR, and SR images. Both the SRCNN and SRGAN consistently and significantly improved the IQ across various signal lengths in terms of these traditional metrics.




Fig. 9 AUC values of the (a) RHO, (b) CHO, and (c) ResNet-IO for HR, LR, and SR images. It can be seen that the DL-SR resulted in a small improvement in the CHO performance, but no improvement in the RHO and ResNet-IO performance on the LR images. As such, the observer performance on the HR images is much higher than the performance on the LR and SR images.

As seen in Fig. 8, the SR images generated by the SRCNN and SRGAN show an improvement in IQ across various signal lengths compared with their LR counterparts in terms of the traditional IQ metrics. Moreover, no significant changes on traditional IQ metrics were observed among SR images when varying the signal length. This is due to the degradation model and DL-SR network architecture being consistent across different signal lengths and the physical difference among images with various signal lengths being minor.

However, as shown in Fig. 9, DL-SR performance as measured by NO performance provides different insights into the DL-SR behavior. First, it can be seen that AUC values corresponding to

all NOs increased consistently along with the increment of the signal length for the HR, LR, and both types of SR images. This is due to the detection task becoming easier with an increasing signal length. Second, the AUC values corresponding to HR images were significantly greater than those on LR images and SR images. This suggested that the second- and potentially higher-order statistical properties of the images may not be recovered by the DL-SR networks. Third, it is worth noting that, in some cases, there was a small improvement in the AUC values of RHO and a small but significant improvement in the AUC values of Gabor CHO corresponding to the SR images as compared with the LR images. This could be interpreted by both the linear observers, namely the RHO and the Gabor CHO acting on the SR images, having the benefit of a nonlinear preprocessing block in the form of the DL-SR network. Finally, as shown in Fig. 9(c), there was no improvement in the performance of the ResNet-IO as a result of the employed DL-SR networks, which is consistent with the data-processing inequality.²⁹

4.1.3 Impact of number of layers in DL-SR networks on observer performance

The traditional IQ metric MSE and the NO performance measured on the LR and SR images as the number of layers in SRCNN was varied are shown in Figs. 10 and 11, respectively. As shown

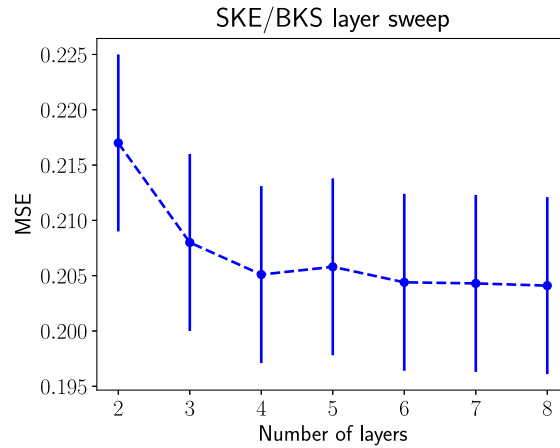


Fig. 10 Ensemble MSE between the SR and the HR images for SR networks with different numbers of layers. The LR images yield an MSE of 0.4369.

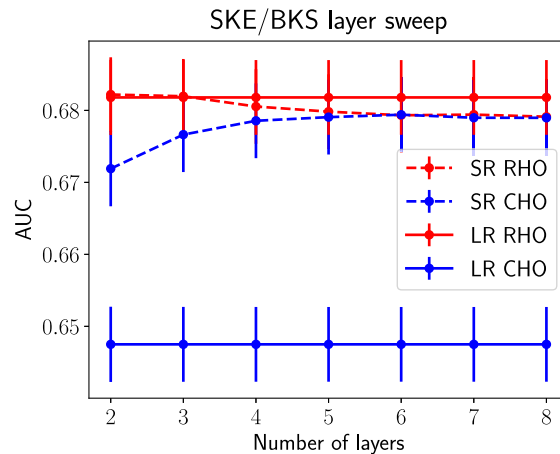


Fig. 11 RHO and CHO performance on SR images and LR images for SR networks with different numbers of layers.

in Fig. 10, the MSEs decreased when the number of layers in SRCNN increased, as expected. This indicates that the DL-SR networks improved certain first-order statistics of the images. However, this trend is not always consistent with the NO performance measured by AUC values. As shown in Fig. 11, it was observed that the AUC values for the RHO measured on SR images were no greater than those computed using the LR images. Also the RHO performance decreased as the number of DL-SR network layers increased. This suggests that the second-order statistical properties of the images were degraded by the DL-SR networks. To further analyze this, the singular values of the covariance matrix $\mathbf{K}(\hat{\mathbf{f}}_{\text{SR}})$ of the SRCNN-resolved images were computed for networks having different numbers of layers. As shown in Fig. 12, the singular values indicate that, as the number of layers in the DL-SR network increased, $\mathbf{K}(\hat{\mathbf{f}}_{\text{SR}})$ became increasingly ill-conditioned.

On the other hand, the AUC values for the Gabor CHO on SR images were greater than those measured on LR images, and the performance of Gabor CHO on SR images increased as the number of layer increased from 2 to 6, after which it saturated and reduced slightly for the SRCNN composed of 7 and 8 layers. This suggests that the second-order statistics of the Gabor channelized images were improved by the DL-SR networks but that this improvement reached a plateau as the number of layers increased. The singular values of the covariance matrix \mathbf{K}_v of the Gabor-channelized, SRCNN-resolved images were computed for the DL-SR networks with different numbers of layers. As shown in Fig. 13, the singular value decay of \mathbf{K}_v is faster for DL-SR networks with more layers, which is similar to the RHO.

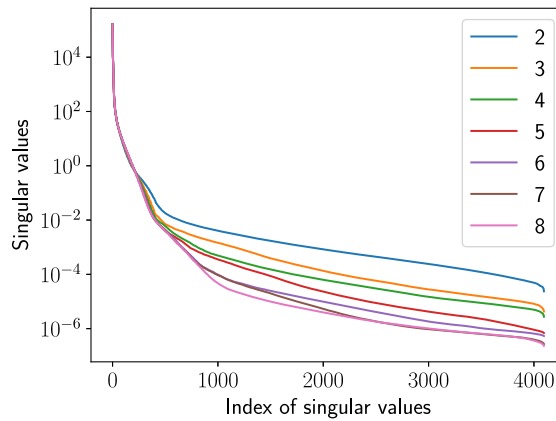


Fig. 12 Singular values of the empirical covariance matrix of the SR images from DL-SR networks of different numbers of layers.




Fig. 13 Singular values of the empirical covariance matrix of the Gabor channelized SR images from DL-SR networks of different numbers of layers.

4.2 Impact of Observer Capacity on Benefit of DL-SR for MC Cluster Detection Performance

The objective of this study is to determine how the capacity of a NO relates to its task performance on SR images. The traditional IQ metrics MSE, PSNR, and SSIM were computed for the LR and SR images generated by the SRCNN on the MC cluster dataset. As shown in Table 2, the IQ measured with these metrics improved for the SRCNN-resolved images compared with the LR counterparts.

The capacity of a ResNet-based observer was varied by varying the number of residual blocks that constitute the ResNet. Figure 14 shows the performance of ResNet-based observers consisting of 2, 4, 6, and 8 residual blocks trained on a dataset of 50,000 images (200,000 considering fourfold flip-augmentation). It was observed that ResNet-based observers of smaller capacity benefited from the particular DL-SR network employed. In this case, the DL-SR network can be interpreted as an additional preprocessing block for the ResNet observer that effectively increases the capacity of the observer. However, as the capacity of the observer was increased, the SR operation gave diminishing returns toward improving the task performance. As the NO performance plateaued with increasing capacity, it approached ResNet-IO, and the MC cluster detection performance on SR images was no greater than that in LR images. This behavior is consistent with the data processing inequality,²⁹ which suggests that postprocessing operations such as image super-resolution will not increase the information content in the image. As a result, the MC cluster detection performance of a ResNet-IO on SR images should not be expected to surpass that of the original LR images.

Next, ResNet-based observers of varying depths were trained on datasets consisting of different sizes to fulfill their corresponding capacity for each resolution. For each dataset, the optimal ResNet-based observer was identified based on the best performance on the validation dataset. The results in Fig. 15 show the performance of the optimal ResNet-based observer for each dataset size. It was observed that, as the amount of available training data increased, the MC cluster detection performance of the ResNet-based observers increased. More interestingly, given a small dataset with limited number of images such as 5000, 10,000, and 20,000, the DL-SR network indeed improved the detection performance on SR images compared with LR.

Table 2 Traditional IQ metrics computed on the LR and SR images in the MC cluster detection task.

Resolution	Ensemble MSE	PNSR	SSIM
LR	0.1580 ± 0.0104	50.1925 ± 0.5390	0.9942 ± 0.0006
SR	0.0486 ± 0.0021	55.2895 ± 0.3546	0.9973 ± 0.0002

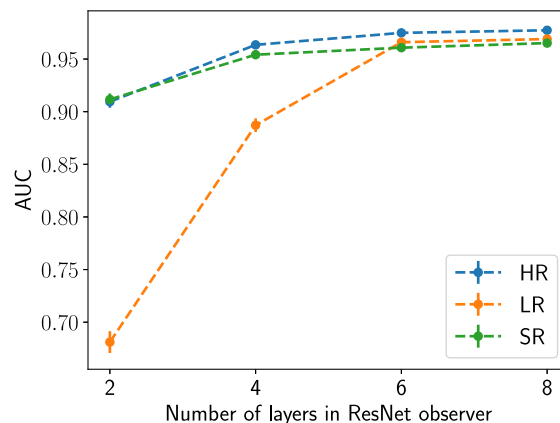


Fig. 14 Performance of ResNet-based observers of different numbers of layers trained on HR, LR, and SR datasets of size 50,000.




Fig. 15 Performance of the optimal ResNet-based observer for a particular dataset size trained on HR, LR, and SR images.

This demonstrates a situation in which the DL-SR operation aided the MC cluster detection performance. For training dataset sizes of 50,000 and beyond, the ResNet-based observer approached the ResNet-IO, and its performance on the images resolved by the DL-SR networks was no better than its performance on the LR images.

Both of the observations in Figs. 14 and 15 illustrate that, in the case of suboptimal neural-network (NN)-based observers, such as those with limited capacity or those trained on limited data, DL-SR networks may be employed to improve the detection performance compared with that achieved on the LR images. However, if the NN-based observer approximates IO, preprocessing the LR images using a DL-SR network will not improve the detection performance of the observer.

5 Discussion

Deep learning techniques have been adopted for a wide range of medical imaging applications, including image restoration. Despite the different traditional IQ metrics having been computed to assess the effect of these deep learning-based methods, a task-based evaluation of these approaches has been largely lacking. A recent study conducted by Li et al.¹⁵ demonstrated that deep neural network-based image denoising methods can result in a loss of task-relevant information, despite an improvement in several traditional IQ metrics. In a similar vein, this work studies the impact of DL-SR on binary signal detection tasks. It is important to reiterate that the main goal of this work is to comprehensively study the impact of DL-SR on task performance for known tasks under known statistical conditions. It is not to explore whether DL-SR can be a viable practical solution to a particular real problem. Such a systematic and comprehensive evaluation is not possible with common clinical datasets, which have several different and unknown sources of variability that may act as confounding factors in our analysis. Therefore, for the purposes of this work, the stylized setup presented is appropriate.

A Rayleigh detection task was employed to assess the impact of the design of the signal and the depth of the DL-SR network, and an MC cluster detection task was employed to study how DL-SR affects NN-based observers of different capacities. The numerical results for the SKE/BKS Rayleigh detection task revealed that the loss of task-relevant information in LR images cannot be recovered by the DL-SR operation, even though mild improvement of detection performance was observed with suboptimal observers. Furthermore, it was observed that, while increasing the depth of the DL-SR network improves the traditional IQ metrics, improved task performance does not always follow. This suggests that the mantra “deeper is better” while designing neural network architectures for image super-resolution is not necessarily applicable when task performance is considered. As such, seeking to minimize a loss function solely related to traditional IQ metrics may lead to a situation in which the image statistics important to the defined task are degraded.

Furthermore, it is of interest to investigate conditions under which the DL-SR improves the signal detection task performance. Using SRCNN as an example, an SKS/BKS MC cluster

detection task was conducted to investigate the capacity of the NN-based observers on SR images, as compared with that on LR and HR images. It was observed that DL-SR improved the signal detection performance of suboptimal observers that do not accurately approximate IOs due to either a limited amount of training data or the limited complexity of the observer. Given sufficient training data and an observer with sufficient complexity for the particular task considered, an IO can be approximated, and the benefit of DL-SR toward improving the task performance is lost. This suggests that the impact of DL-SR on a binary signal detection task depends on a combination of factors such as the DL-SR networks, the observers, and the defined task. Thus a task-based evaluation of DL-SR methods is essential to accurately quantify the benefit of DL-SR for clinical practice.

Some important topics remain to be investigated in the future. The binary signal detection tasks considered in this study are simplistic compared with real-world clinical tasks. Future work could investigate the performance of DL-SR methods as preprocessing blocks on tasks such as multi-class classification, lesion segmentation, and image registration. Since the introduction of SRCNN and SRGAN, several deep learning-based methods that improve the super-resolution performance have been proposed. The task-based evaluation pipeline presented in this study can readily be applied to the newer DL-SR methods in which different network architectures or loss functions are employed. It is known that deep learning-based methods may lead to hallucinations, especially when acting on data outside the training distribution.⁴⁹ Hence, an objective assessment of the robustness of DL-SR methods for distribution shifts is also an important topic for future investigation. Additionally, it will be important to conduct human reader studies to assess the performance of DL-SR methods for specific clinical tasks. The results demonstrated in our study will motivate the development of DL-SR methods in directions in which the loss of task-specific information can be mitigated by incorporating such information in designing the network architecture or the loss functions.⁵⁰

6 Conclusion

In this paper, we presented a task-based evaluation to assess the impact of DL-SR methods on binary signal detection. An SKE/BKS Rayleigh detection task and an SKS/BKS MC cluster detection task were conducted on simulated image datasets with a CLB. Our results verify that the performance of an IO cannot be improved via DL-SR methods, which is consistent with the data processing inequality. Also an improvement in traditional IQ metrics induced by DL-SR does not always correlate with the impact of DL-SR on observer performance. Despite this, the numerical experiments presented indicate that DL-SR methods improved the signal detection performance of suboptimal NOs in certain cases. The reported results emphasized the necessity of a task-based evaluation of DL-SR methods and suggest future avenues for developing effective DL-SR algorithms.

Disclosures

The authors declare no potential conflicts of interest.

Acknowledgments

This work was supported in part by the National Institutes for Health, Award Nos. R01EB020604, R01EB023045, R01NS102213, R01CA233873, and R21CA223799. The authors greatly appreciate Michael X. Wu for proofreading the manuscript carefully and thoughtfully. Preliminary results of this work were presented at SPIE Medical Imaging 2021 and published as an SPIE Proceedings paper.⁵¹

References

1. H. Chen et al., "Real-world single image super-resolution: a brief review," arXiv:2103.02368 (2021).

2. R. Keys, "Cubic convolution interpolation for digital image processing," *IEEE Trans. Acoust. Speech Signal Process.* **29**(6), 1153–1160 (1981).
3. S. Dai et al., "SoftCuts: a soft edge smoothness prior for color image super-resolution," *IEEE Trans. Image Process.* **18**(5), 969–981 (2009).
4. E. J. Candès and C. Fernandez-Granda, "Towards a mathematical theory of super-resolution," *Commun. Pure Appl. Math.* **67**(6), 906–956 (2014).
5. W. Yang et al., "Deep learning for single image super-resolution: a brief review," *IEEE Trans. Multimedia* **21**(12), 3106–3121 (2019).
6. C. Dong et al., "Image super-resolution using deep convolutional networks," *IEEE Trans. Pattern Anal. Mach. Intell.* **38**(2), 295–307 (2016).
7. W.-S. Lai et al., "Fast and accurate image super-resolution with deep Laplacian pyramid networks," *IEEE Trans. Pattern Anal. Mach. Intell.* **41**(11), 2599–2613 (2019).
8. C. Ledig et al., "Photo-realistic single image super-resolution using a generative adversarial network," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, pp. 4681–4690 (2017).
9. H. H. Barrett and K. J. Myers, *Foundations of Image Science*, John Wiley & Sons (2013).
10. X. He and S. Park, "Model observers in medical imaging research," *Theranostics* **3**(10), 774 (2013).
11. R. F. Wagner and D. G. Brown, "Unified SNR analysis of medical imaging systems," *Phys. Med. Biol.* **30**(6), 489 (1985).
12. W. Vennart, "ICRU report 54: medical imaging-the assessment of image quality: ISBN 0-913394-53-x. April 1996, Maryland, USA," *Radiography* **3**(3), 243–244 (1997).
13. C. E. Metz et al., "Toward consensus on quantitative assessment of medical imaging systems," *Med. Phys.* **22**(7), 1057–1061 (1995).
14. W. Zhou, H. Li, and M. A. Anastasio, "Approximating the ideal observer and Hotelling observer for binary signal detection tasks by use of supervised learning methods," *IEEE Trans. Med. Imaging* **38**(10), 2456–2468 (2019).
15. K. Li et al., "Assessing the impact of deep neural network-based image denoising on binary signal detection tasks," *IEEE Trans. Med. Imaging* **40**(9), 2295–2305 (2021).
16. K. Umehara, J. Ota, and T. Ishida, "Super-resolution imaging of mammograms based on the super-resolution convolutional neural network," *Open J. Med. Imaging* **7**(4), 180–195 (2017).
17. C. You et al., "Ct super-resolution GAN constrained by the identical, residual, and cycle learning ensemble (GAN-circle)," *IEEE Trans. Med. Imaging* **39**(1), 188–203 (2020).
18. Q. Lyu et al., "Super-resolution MRI and ct through GAN-circle," *Proc. SPIE* **11113**, 111130X (2019).
19. A. S. Chaudhari et al., "Super-resolution musculoskeletal MRI using deep learning," *Magn. Reson. Med.* **80**(5), 2139–2154 (2018).
20. B. Ma et al., "MRI image synthesis with dual discriminator adversarial learning and difficulty-aware attention mechanism for hippocampal subfields segmentation," *Comput. Med. Imaging Graph.* **86**, 101800 (2020).
21. S. Anandasabapathy et al., "An optical, endoscopic brush for high-yield diagnostics in esophageal cancer," *Proc. SPIE* **11620**, 116200B (2021).
22. O. Christianson et al., "An improved index of image quality for task-based performance of CT iterative reconstruction across three commercial implementations," *Radiology* **275**(3), 725–734 (2015).
23. H. H. Barrett et al., "Model observers for assessment of image quality," *Proc. Natl. Acad. Sci. U. S. A.* **90**(21), 9758–9765 (1993).
24. K. J. Myers et al., "Effect of noise correlation on detectability of disk signals in medical imaging," *J. Opt. Soc. Am. A* **2**(10), 1752–1759 (1985).
25. A. Badal et al., "Virtual clinical trial for task-based evaluation of a deep learning synthetic mammography algorithm," *Proc. SPIE* **10948**, 109480O (2019).
26. Z. Wang, J. Chen, and S. C. Hoi, "Deep learning for image super-resolution: a survey," *IEEE Trans. Pattern Anal. Mach. Intell.* **43**, 3365–3387 (2021).
27. D. Dai et al., "Is image super-resolution helpful for other vision tasks?" in *IEEE Winter Conf. Appl. Comput. Vision*, IEEE, pp. 1–9 (2016).

28. L. Jaffe, S. Sundram, and C. Martinez-Nieves, "Super-resolution to improve classification accuracy of low-resolution images," Tech. Rep., Tech. Rep. 19, Stanford University (2017).
29. N. J. Beaudry and R. Renner, "An intuitive proof of the data processing inequality," arXiv:1107.0740 (2011).
30. A. A. Sanchez, E. Y. Sidky, and X. Pan, "Task-based optimization of dedicated breast CT via Hotelling observer metrics," *Med. Phys.* **41**(10), 101917 (2014).
31. K. M. Hanson and K. J. Myers, "Rayleigh task performance as a method to evaluate image reconstruction algorithms," in *Maximum Entropy and Bayesian Methods*, W. T. Grandy and L. H. Schick, eds., pp. 303–312, Springer, Dordrecht (1991).
32. C. Zhang et al., "Understanding deep learning requires rethinking generalization," in *Proc. 5th Int. Conf. Learn. Represent.* (2017).
33. M. A. Kupinski et al., "Ideal-observer computation in medical imaging with use of Markov-chain Monte Carlo techniques," *J. Opt. Soc. Am. A* **20**(3), 430–438 (2003).
34. W. Zhou and M. A. Anastasio, "Markov-chain Monte Carlo approximation of the ideal observer using generative adversarial networks," *Proc. SPIE* **11316**, 113160D (2020).
35. Y. Zhang, B. T. Pham, and M. P. Eckstein, "The effect of nonlinear human visual system components on performance of a channelized Hotelling observer in structured backgrounds," *IEEE Trans. Med. Imaging* **25**(10), 1348–1362 (2006).
36. M. P. Eckstein, C. K. Abbey, and J. S. Whiting, "Human vs model observers in anatomic backgrounds," *Proc. SPIE* **3340**, 16–26 (1998).
37. L. Yu et al., "Prediction of human observer performance in a 2-alternative forced choice low-contrast detection task using channelized Hotelling observer: impact of radiation dose and reconstruction algorithms," *Med. Phys.* **40**(4), 041908 (2013).
38. J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," *Lect. Notes Comput. Sci.* **9906**, 694–711 (2016).
39. R. H. Hahnloser et al., "Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit," *Nature* **405**(6789), 947–951 (2000).
40. F. O. Bochud, C. K. Abbey, and M. P. Eckstein, "Statistical texture synthesis of mammographic images with clustered lumpy backgrounds," *Opt. Express* **4**(1), 33–43 (1999).
41. D. P. Kingma and J. Ba, "Adam: a method for stochastic optimization," arXiv:1412.6980 (2014).
42. T. Tot et al., "The clinical value of detecting microcalcifications on a mammogram," *Semin. Cancer Biol.* **72**, 165–174 (2021).
43. P. Timberg et al., "Visibility of microcalcification clusters and masses in breast tomosynthesis image volumes and digital mammography: a 4afc human observer study," *Med. Phys.* **39**(5), 2431–2437 (2012).
44. L. R. Borges, P. M. de Azevedo Marques, and M. A. Vieira, "A 2-AFC study to validate artificially inserted microcalcification clusters in digital mammography," *Proc. SPIE* **10952**, 109520R (2019).
45. M. Ruschin et al., "Using simple mathematical functions to simulate pathological structures-input for digital mammography clinical trial," *Radiat. Prot. Dosimetry* **114**(1-3), 424–431 (2005).
46. E. R. DeLong, D. M. DeLong, and D. L. Clarke-Pearson, "Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach," *Biometrics* **44**(3), 837–45 (1988).
47. X. Sun and W. Xu, "Fast implementation of delong's algorithm for comparing the areas under correlated receiver operating characteristic curves," *IEEE Signal Process. Lett.* **21**(11), 1389–1393 (2014).
48. X. Robin et al., "pROC: display and analyze ROC curves. R package version 1.10.0" (2017).
49. S. Bhadra et al., "On hallucinations in tomographic image reconstruction," arXiv:2012.00646 (2020).
50. J. Zhang et al., "Task-oriented low-dose CT image denoising," arXiv:2103.13557 (2021).
51. V. A. Kelkar et al., "Task-based evaluation of deep image super-resolution in medical imaging," *Proc. SPIE* **11599**, 115990X (2021).

Xiaohui Zhang received her BE degree in biomedical engineering from Beihang University, Beijing, China, in 2018. She is a PhD candidate in the Department of Bioengineering at the University of Illinois at Urbana–Champaign (UIUC). Her research interests include computational methods for neuroimaging and machine learning for medical imaging applications. She is also a member of SPIE.

Varun A. Kelkar received his MS degree in electrical and computer engineering from UIUC in 2019 and his BTech degree in engineering physics from the Indian Institute of Technology Madras, Tamil Nadu, India, in 2017. He is a PhD candidate in the Department of Electrical and Computer Engineering, UIUC, Illinois, USA. His research interests include computational imaging, inverse problems, signal processing, optics, and machine learning. He is a member of SPIE. He was a recipient of the 2019 SPIE Optics and Photonics Education Scholarship and the 2021 Oak Ridge Institute of Science and Education fellowship.

Jason Granstedt received his BS and MS degrees in electrical engineering from Virginia Polytechnic Institute and State University in 2015 and 2017, respectively. He is currently a PhD candidate in the Department of Computer Science at the UIUC. His research interests include task-based analysis of images and application of machine learning techniques to medical imaging. He is a member of SPIE.

Hua Li is a research associate professor in the Department of Bioengineering at the UIUC and a medical physicist at Carle Foundation Hospital, Urbana, Illinois, USA. Her research work focuses on developing innovative medical imaging and image analysis techniques to solve the challenges seen in clinical practice, toward improving personalized patient care. She serves as the deputy editor for the *Journal of Medical Physics* and a reviewer for a set of journals and NIH study sections.

Mark A. Anastasio is the Donald Biggar Willett Professor in Engineering and the head of the Department of Bioengineering at the UIUC. He is a fellow of SPIE, the American Institute for Medical and Biological Engineering, and the International Academy of Medical and Biological Engineering. His research addresses computational image science, inverse problems in imaging, and machine learning for imaging applications. He has contributed to emerging biomedical imaging technologies, including photoacoustic computed tomography and ultrasound computed tomography.