

University of Dayton

eCommons

---

Electrical and Computer Engineering Faculty  
Publications

Department of Electrical and Computer  
Engineering

---

10-29-2021


## Resampling and Super-Resolution of Hexagonally Sampled Images Using Deep Learning

Dylan Flaute  
*University of Dayton*

Russell C. Hardie  
*University of Dayton*, rhardie1@udayton.edu

Hamed Elwarfalli  
*University of Dayton*

Follow this and additional works at: [https://ecommons.udayton.edu/ece\\_fac\\_pub](https://ecommons.udayton.edu/ece_fac_pub)

 Part of the [Computer Engineering Commons](#), [Electrical and Electronics Commons](#), [Electromagnetics and Photonics Commons](#), [Optics Commons](#), [Other Electrical and Computer Engineering Commons](#), and the [Systems and Communications Commons](#)

---

### eCommons Citation

Flaute, Dylan; Hardie, Russell C.; and Elwarfalli, Hamed, "Resampling and Super-Resolution of Hexagonally Sampled Images Using Deep Learning" (2021). *Electrical and Computer Engineering Faculty Publications*. 438.

[https://ecommons.udayton.edu/ece\\_fac\\_pub/438](https://ecommons.udayton.edu/ece_fac_pub/438)

This Article is brought to you for free and open access by the Department of Electrical and Computer Engineering at eCommons. It has been accepted for inclusion in Electrical and Computer Engineering Faculty Publications by an authorized administrator of eCommons. For more information, please contact [mschlengen1@udayton.edu](mailto:mschlengen1@udayton.edu), [ecommons@udayton.edu](mailto:ecommons@udayton.edu).

# Resampling and super-resolution of hexagonally sampled images using deep learning

Dylan Flaute<sup>a\*</sup>, Russell C. Hardie<sup>a</sup>, Hamed Elwarfalli<sup>a</sup>

<sup>a</sup>University of Dayton, Department of Electrical and Computer Engineering, Dayton, Ohio, United States

**Abstract.** Super-resolution (SR) aims to increase the resolution of imagery. Applications include security, medical imaging, and object recognition. We propose a deep learning-based SR system that takes a hexagonally sampled low-resolution image as an input and generates a rectangularly sampled SR image as an output. For training and testing, we use a realistic observation model that includes optical degradation from diffraction and sensor degradation from detector integration. Our SR approach first uses non-uniform interpolation to partially upsample the observed hexagonal imagery and convert it to a rectangular grid. We then leverage a state-of-the-art convolutional neural network (CNN) architecture designed for SR known as Residual Channel Attention Network (RCAN). In particular, we use RCAN to further upsample and restore the imagery to produce the final SR image estimate. We demonstrate that this system is superior to applying RCAN directly to rectangularly sampled LR imagery with equivalent sample density. The theoretical advantages of hexagonal sampling are well known. However, to the best of our knowledge, the practical benefit of hexagonal sampling in light of modern processing techniques such as RCAN SR is heretofore untested. Our SR system demonstrates a notable advantage of hexagonally sampled imagery when employing a modified RCAN for hexagonal SR.

**Keywords:** hexagonal sampling, single-image super-resolution, convolutional neural network, image restoration.

\*Dylan Flaute, [dylan.flaute@gmail.com](mailto:dylan.flaute@gmail.com)

© 2021 Society of Photo-Optical Instrumentation Engineers (SPIE). One print or electronic copy may be made for personal use only. Systematic reproduction and distribution, duplication of any material in this publication for a fee or for commercial purposes, and modification of the contents of the publication are prohibited.

## 1 Introduction

Super-resolution (SR) is a term for the task of increasing the resolution of low-resolution (LR) images. The SR task has broad applications in security and surveillance,<sup>1,2</sup> medical imaging,<sup>3</sup> and object recognition.<sup>4</sup> Single-image SR (SISR) is fundamentally challenging because it generally requires the interpolation of undersampled imagery without access to additional spatial samples from other frames. When processing LR data acquired directly from a camera, the observation model includes diffraction blurring and detector integration effects, in addition to potential aliasing that results from undersampling.

Effectively addressing the ill-posed SISR problem requires the exploitation of prior statistical information about the images. Deep learning is a broad class of data-driven approaches to deriving statistical information. Recently, deep learning-based approaches have pushed the state-of-the-art for inverse problems in image processing, including but not limited to denoising, deconvolution, SR, and inpainting.<sup>5</sup> In the field of SISR, the use of neural networks<sup>6</sup> and convolutional neural networks (CNNs) has been widely successful. In CNN-based approaches, convolutional filters

---

Cite as: Dylan Flaute, Russell C. Hardie, Hamed Elwarfalli, "Resampling and super-resolution of hexagonally sampled images using deep learning," Opt. Eng. 60(10), 103105 (2021), <https://doi.org/10.1117/1.OE.60.10.103105>

are learned based on very large datasets of paired LR and high-resolution (HR) images. Examples of CNN-based approaches to SISR include SRCNN,<sup>7</sup> FSRCNN,<sup>8</sup> SCN,<sup>9</sup> VDSR,<sup>10</sup> DRRN,<sup>11</sup> LapSRN,<sup>12</sup> MSLapSRN,<sup>13</sup> ENet-PAT,<sup>4</sup> MemNet,<sup>14</sup> EDSR,<sup>15</sup> SRMDNF,<sup>16</sup> and RCAN.<sup>17</sup> A comprehensive survey on deep learning-based SISR approaches is provided in Ref. 18. Deep learning has also been successful in the field of deblurring. DeepDeblur,<sup>19</sup> ED-DSRN,<sup>20</sup> GFN,<sup>21</sup> and EDSR<sup>15</sup> all approach deblurring, mostly in addition to SR, with deep learning.

Although CNN-based approaches to learning and exploiting prior statistical information have been successful, the approach is still fundamentally bounded by the information present in the samples. Alternative approaches to advancing the state of SISR examine the sampling. For example, CAR<sup>22</sup> is a deep learning approach to find a more optimal down-sampling function that is easier for another network to undo. A simpler, theoretically-grounded sampling alternative is to sample on a hexagonal grid rather than a rectangular grid. This approach has the theoretical motivation that a hexagonal grid with the same density as the rectangular grid can support a higher cutoff frequency without introducing aliasing.

Some recent works demonstrate the benefits of hexagonal imagery for processing purposes by processing the hexagonal image directly. For example, the image classification system HexaConv<sup>23</sup> processes hexagonal images by adapting the Group-equivariant Convolutional Neural Network (G-CNN) architecture<sup>24</sup> to leverage the additional symmetries that are found on a hexagonal grid, yielding improved classification performance at a given parameter budget than baseline on the natural images in the CIFAR-10 dataset and on the aerial imagery, which is rotation-invariant by nature, in the Aerial Image Dataset (AID).<sup>25</sup> The image classification system HexagonNet<sup>26</sup> also processes hexagonal images directly and increments performance further on the AID dataset as well as achieving a new state-of-the-art on the ModelNet40 shape classification dataset. And, the authors of Ref. 27 use hexagonally-defined filters to process omnidirectional imagery for orientation-aware semantic segmentation. To facilitate future research into this type of processing, the Python library HexagDLY<sup>28</sup> extends the deep learning framework PyTorch<sup>29</sup> to include efficient hexagonal convolutions. The more abstract problem of developing convolution operators for any non-Cartesian grid is approached by the authors of Ref. 30, who use nearest-neighbor information to index the convolution operations. Aside from image processing, hexagonally gridded data is also useful for estimating the supply-demand gap for ride-sharing due to the increased symmetries of the grid.<sup>31</sup> Joint resampling and restoration of hexagonally sampled images using an Adaptive Wiener Filter has also been explored Ref. 32.

The theoretical advantages of hexagonal sampling are well known. However, hexagonal sampling comes at a cost. Most display devices and image processing algorithms and software today assume rectangular sample/pixel layouts. Furthermore, to the best of our knowledge, the practical benefit of hexagonal sampling in light of modern SR processing techniques is heretofore untested. In this paper, we demonstrate a system for SISR that takes hexagonally sampled LR imagery as inputs and generates rectangularly sampled SR images as outputs. To make the problem more realistic, our observation model includes the effects of diffraction and detector integration for a given camera system. Our SR approach first uses non-uniform interpolation to partially upsample the observed hexagonal imagery and convert it to a rectangular grid. We then leverage a state-of-the-art CNN architecture designed for SR known as Residual Channel Attention Network (RCAN).<sup>17</sup> In particular, we use RCAN to further upsample and restore the imagery to produce the final SR image estimate. To enrich the RCAN input, we also supply subpixel interpolation distances for

each pixel and feed into the RCAN architecture through a separate input head. This modification is inspired by the FIFNET architecture for multi-image SR.<sup>33</sup>

Since paired hexagonally sampled LR images and rectangularly sampled HR images are not readily available, we develop an observation model-based approach for generating realistic, full frame RGB, synthetic data. Our observation model effectively constitutes a “forward model” for the inverse problem we wish to solve. The observation model leverages a camera-specific optical transfer function (OTF) that models diffraction and detector integration (based on either hexagonal or rectangular detectors) for generating our data<sup>33,34</sup> as well as a resampling step. We apply the observation model to the DIV2K SISR dataset<sup>35</sup> to generate data to train our models so as to be ready to process real full frame RGB color camera outputs.

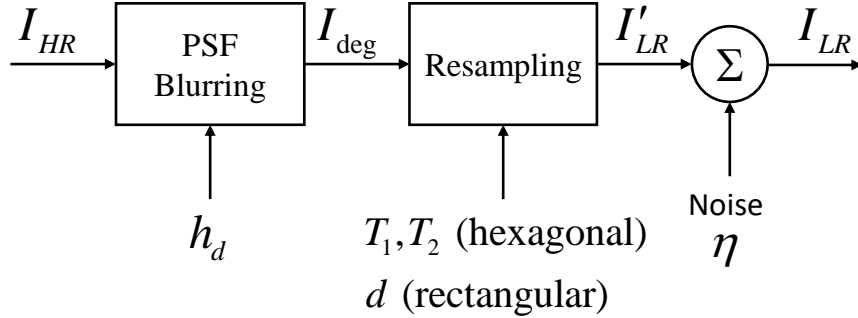
We demonstrate that our modified RCAN architecture processing data that was originally sampled hexagonally is superior to applying RCAN directly to rectangularly sampled LR imagery with equivalent sample density. To make this comparison, we choose a sample density such that no interpolation is necessary for generating the rectangularly sampled imagery (and thus, there is no interpolation error). We generate a rectangularly sampled color-imagery dataset from DIV2K with rectangular detector shape and a hexagonally sampled dataset with hexagonal detector shape. To restore the rectangular imagery, we train RCAN end-to-end and leverage a pre-trained RCAN model. Despite having no interpolation error in the input data, the rectangularly-based architecture still underperforms the hexagonally-based architecture. The implications of our results are two-fold: First, our system demonstrates the practical benefit of hexagonally sampled imagery for image processing, even in the previously untested context of deep learning-based post-processing. Second, our system reduces the practical drawback of using hexagonally sampled imagery – namely, that most display devices and image processing software assumes rectangular sample/pixel layouts – since our system can be used to resample hexagonal imagery to rectangular grids.

The remainder of this paper is laid out as follows. Section 2 describes the observation model that is used to generate training and testing data. Section 3 describes the architecture of our approach to resampling and restoring the imagery produced by that observation model. Section 4 describes the experiments we perform to compare hexagonally sampled imagery and rectangularly sampled imagery post-processed with those architectures. Section 5 describes the results of those experiments. Finally, Section 6 provides our discussion of those results.

## 2 Observation Model

A block diagram of our observation model is shown in Fig. 1. The model takes an ideal discrete-space HR image as input and outputs a simulated LR image that approximates a camera sensor output. The observation model consists of two main parts: an optical transfer function (OTF) to account for camera optics and detector integration, and a resampling step to simulate the LR sampling pattern of the sensor. We simulate both hexagonal and rectangular sampling by resampling the blurred HR image onto the appropriate grid. Finally, our model includes independent and additive white Gaussian noise with standard deviation  $\sigma_\eta$ . The key parameters of our observation model are listed in Table 1 for the experimental study presented here. The parameters in Table 1 have been selected to represent a typical acquisition scenario with an RGB camera, with a few liberties taken for numerical convenience. The wavelengths correspond to the Hamamatsu RGB





**Fig 1** Discrete forward observation model used to provide realistic degradation of ideal images for training and testing purposes. A pristine HR image is blurred with an impulse-invariant PSF and then resampled according the detector sampling pattern. Finally, noise is added to produce the LR image.

**Table 1** Summary of observation model parameters for the red, green, and blue color channels.

Parameter	Wavelength		
	$\lambda_R = 620.0 \text{ nm}$	$\lambda_G = 540.0 \text{ nm}$	$\lambda_B = 460.0 \text{ nm}$
F-number	$F = 4.000$	$F = 4.000$	$F = 4.000$
Cutoff frequency	$\rho_c = 0.4032 \text{ cyc}/\mu\text{m}$	$\rho_c = 0.4630 \text{ cyc}/\mu\text{m}$	$\rho_c = 0.5435 \text{ cyc}/\mu\text{m}$
Nyquist pitch	$p_N = 1.240 \mu\text{m}$	$p_N = 1.080 \mu\text{m}$	$p_N = 0.9200 \mu\text{m}$
HR pitch	$p = 1.000 \mu\text{m}$	$p = 1.000 \mu\text{m}$	$p = 1.000 \mu\text{m}$
Rectangular LR pitch	$d = 4.000 \mu\text{m}$	$d = 4.000 \mu\text{m}$	$d = 4.000 \mu\text{m}$
Hexagonal LR pitch 1	$T_1 = 4.298 \mu\text{m}$	$T_1 = 4.298 \mu\text{m}$	$T_1 = 4.298 \mu\text{m}$
Hexagonal LR pitch 2	$T_2 = 7.445 \mu\text{m}$	$T_2 = 7.445 \mu\text{m}$	$T_2 = 7.445 \mu\text{m}$
Noise std.	$\sigma = 1.000 \text{ DU}$	$\sigma = 1.000 \text{ DU}$	$\sigma = 1.000 \text{ DU}$

Color Sensor S9702. The f-number is moderate and commonly used for camera lenses. The simulated rectangular detector pitch is within the range of currently available commercial sensors. We selected the HR pitch so that the rectangular detector pitch would be an integer multiple, to allow for simple downsampling in modeling the rectangular detector scenario. The hexagonal pitches were selected to give the same sampling density as the rectangular sampling scenario. Finally the noise standard deviation has been selected to give a high, but realistic, signal-to-noise ratio (SNR) of approximately 36 dB for the simulated images. A high SNR was selected because the proposed study is intended to focus on detector sampling and not noise reduction. Note that we perform the simulation on three color channels for RGB images to model a full frame RGB color camera and some of the parameters in Table 1 are wavelength dependent. In the following two subsections we shall provide further details on the OTF and resampling blocks, respectively.

### 2.1 Optical transfer function

In this section we discuss the OTF model that accounts for diffraction-limited optics and blurring from spatial integration of the detector elements. Together, these two components can be combined

to provide the overall OTF model as

$$H(u, v) = H_{dif}(u, v)H_{det}(u, v), \quad (1)$$

where  $u$  and  $v$  are the horizontal and vertical spatial frequencies in cycles per unit distance. The first term,  $H_{dif}$ , is the OTF for diffraction-limited optics and the second term,  $H_{det}$ , accounts for detector integration.

The OTF component for diffraction assuming a circular pupil function<sup>36</sup> is given by

$$H_{dif}(\rho) = \begin{cases} \frac{2}{\pi} \left[ \cos^{-1}(\rho) - \rho\sqrt{1-\rho^2} \right] & \rho < 1, \\ 0 & \textit{else} \end{cases} \quad (2)$$

where  $\rho = \sqrt{u^2 + v^2}/\rho_c$ ,  $\rho_c = 1/(\lambda F)$  is the optical cutoff frequency,  $\lambda$  is the wavelength of light, and  $F$  is the f-number of the optics. Note that the f-number is defined as the ratio of the focal length of the optics to the effective aperture.<sup>37</sup>

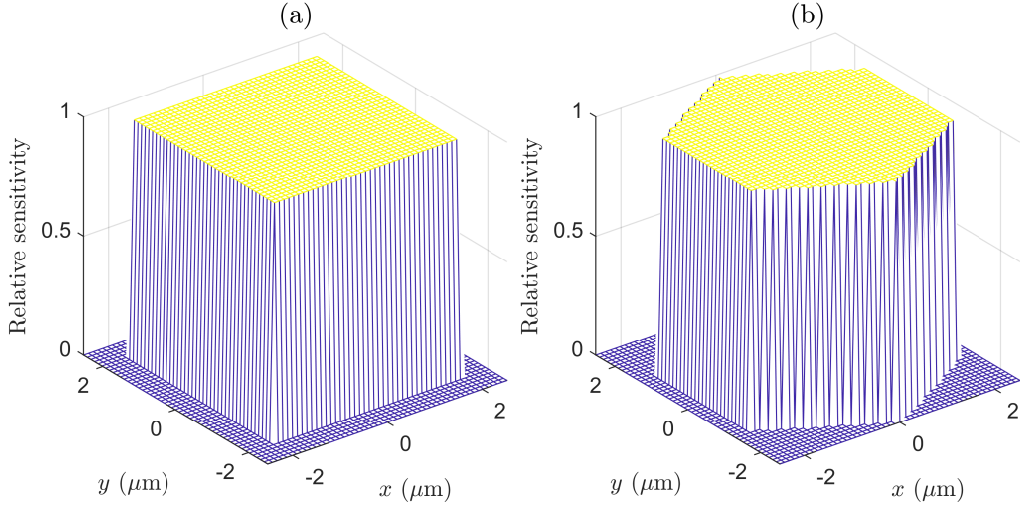
The detector component of the OTF accounts for the finite size of the detectors in the camera's focal plane array. At the native pixel scale, the detector integration plays only a small role in degrading the observed imagery. However, when upsampling an image for SR, the native detector active area will now span multiple HR pixels and play an important and potentially dominant role in the observation model. The detector OTF can be found by simply taking the 2D Fourier transform of the detector element shape as described by Hardie et al.<sup>34</sup> In our study, we wish to present a fair comparison between rectangular and hexagonal sampling. To do so, we have elected to make the active area of both types of detectors the same and have 100% fill factor. This will give rise to the same signal-to-noise ratio using comparable image acquisition parameters. Furthermore, this arrangement gives rise to the same sampling density in spatial samples per unit area. The spatial detector active areas used for the rectangular and hexagonal detectors are shown in Fig. 2. The specific dimensions of the detectors shown are based on the parameters in Table 1.

The diffraction OTF is shown in Fig. 3 for the parameters in Table 1 for the green channel. The detector OTFs along with the combined OTFs are shown in Fig. 4 for the green channel. Note that the detector OTF in Fig. 4(a) has a sinc form corresponding to the Fourier transform of the square detector shape in Fig. 2(a). and note that the OTF for the hexagonal detector in Fig. 4(b) has an approximately circularly symmetric form by virtue of the hexagonal detector shape shown in Fig. 2(b). The combined OTFs that include the diffraction and the detector are shown in Figs. 4(c) and (d) for rectangular and hexagonal, respectively. One can see from the OTFs that the low-pass effect of the two different detector types of the same area is quite similar. However, one might argue that the more radially symmetric OTF shown in Fig. 4(d) is somewhat advantageous: it gives an imager that employs hexagonal sampling a more angle-independent resolution performance.

The OTF model in Eq. (1) is for a continuous-space image. The continuous-space PSF is given by the inverse continuous-space Fourier transform (ICSFT) of the OTF in Eq. (1). This may be expressed as

$$h(x, y) = \text{ICSFT}\{H(u, v)\}. \quad (3)$$

To process the discrete high-resolution input frames, as shown in our observation model in Fig. 1, we employ the impulse-invariant discrete-space equivalent blurring. To do so, we assume our pristine input images are sampled at the Nyquist rate with a pitch of  $p$ . The Nyquist criterion dictates that  $1/p > 2\rho_c = 2/(\lambda F)$ . Equivalently, this means we must have  $p < \lambda F/2$  to guarantee



**Fig 2** Detector shapes based on the parameters in Table 1 for (a) rectangular detectors and (b) hexagonal detectors.

no aliasing.<sup>34</sup> An impulse-invariant equivalent discrete-space PSF,  $h_d(n_1, n_2)$ , may be found by sampling  $h(x, y)$  with a sample spacing of  $p$ . This impulse-invariant PSF is applied to produce


$$I_{deg}(n_1, n_2) = I_{HR}(n_1, n_2) * h_d(n_1, n_2), \quad (4)$$

where  $*$  is discrete 2D convolution.

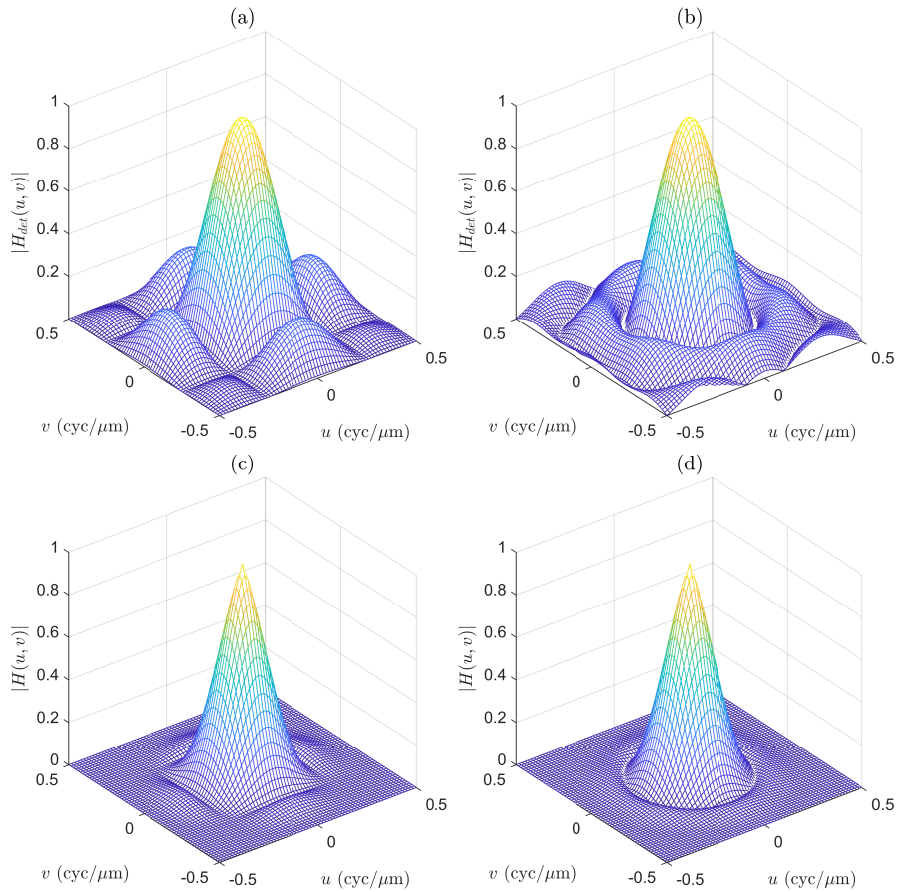
To accommodate full color RGB images, we compute the impulse-invariant PSF for the center wavelengths of red, green, and blue color channels. The assumed wavelengths are listed in Table 1. Note that for our simulation parameters the HR grid rate exceeds the Nyquist rate for the red and green channels, but is slightly below Nyquist for the blue channel. To address this, we truncate the theoretical blue OTF model to prevent any aliasing with regard to the impulse invariant filter. To justify this, note that there is an exceedingly small amount of signal energy beyond the folding frequency that is truncated. In particular, we computed the total volume of the OTF and the volume inside the folding frequency of the  $1 \mu\text{m}$  HR sample spacing for the blue channel. The OTF volume beyond the folding frequency ( $0.5 \text{ cyc}/\mu\text{m}$ ) represents only 0.0105% of the total OTF volume (or a fraction of 0.000105). Visually, note that the blue OTF is very similar to the green OTF shown in Fig. 4(c) and (d), and the OTF values in the figure are very small near the  $0.5 \text{ cyc}/\mu\text{m}$  folding frequency. Based on this, we believe this small adjustment to the blue OTF model is negligible and will not adversely impact the fidelity of the study.

## 2.2 Spatial sampling

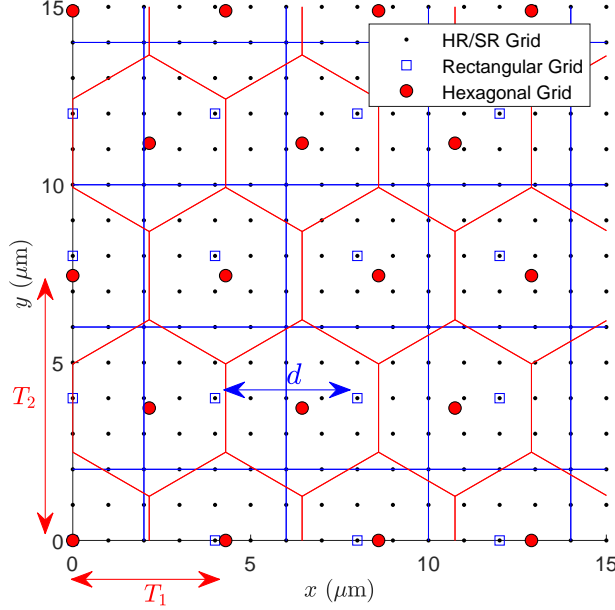
In the previous section, we outlined the PSF blurring model that simulates camera optics and detector integration on the HR image. The next step of the observation model is spatial resampling of the blurred input image  $I_{deg}$  according to the corresponding detector layout. The sampling grids using the parameters listed in Table 1 are shown in Fig. 5. The HR grid (black) is a dense rectangular Nyquist sampling configuration with spacing  $p = 1 \mu\text{m}$ . The pristine input image,  $I_{HR}$ , is



**Fig 3** Diffraction OTF model from Eq. (2) for the green channel parameters in Table 1.



**Fig 4** Detector OTFs from for (a) rectangular detectors and (b) hexagonal detectors. Combined OTFs from Eq. (1) for (c) rectangular detectors and (c) hexagonal detectors.

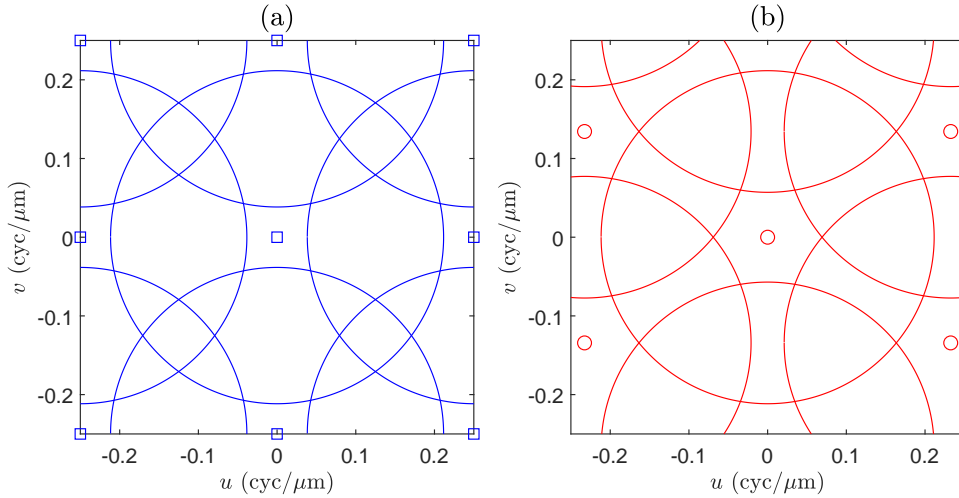


**Fig 5** Sampling grids used in the observation model. The HR grid (black) is a dense rectangular Nyquist sampling configuration. The LR rectangular sampling grid and associated detectors are shown in blue. The LR hexagonal sampling grid and associated detectors are shown in red.

assumed to be sampled on this grid, as is the SR output image. The LR rectangular sampling grid (blue) has a sample spacing of  $d$ . In our current experiment we have  $d = 4p$ . The hexagonal LR grid (red) can be thought of as two interlaced rectangular grids with spacing of  $T_1$  in the horizontal dimension and  $T_2$  in the vertical dimension. The interlacing offset is effectively a  $1/2$  sample spacing shift in the horizontal and vertical dimensions. To form an ideal hexagonal sampling pattern, we require that  $T_2 = \sqrt{3}T_1$ .<sup>38</sup> Approximate hexagonal grids are also sometimes employed<sup>32</sup> where the relationship between  $T_2$  and  $T_1$  is a rational fraction, or even simply a 2-to-1 grid where  $T_2 = 2T_1$ . These approximately hexagonal grids have less perfect frequency packing properties, but may provide other benefits with regard to fabrication or post processing.

Based on the LR sample positions, we use bicubic interpolation on the image  $I_{deg}$  to produce  $I'_{LR}$ . Since  $I_{deg}$  is on a Nyquist grid, the downsampling interpolation here is very accurate and unhindered by aliasing. We then add independent white Gaussian noise with a standard deviation of  $\sigma_\eta = 1$  digital unit (DU). Finally, we quantize the intensity to  $2^8$  levels to form the final LR image, denoted  $I_{LR}$ . This process is repeated for each of the color channels. Note that we are simulating a full frame RGB color camera, and not color filter mosaic based images. The spatial sampling density of the LR rectangular sampling grid is  $1/d^2$ . For the LR hexagonal grid the sampling density is  $2/(T_1T_2) = 2/(\sqrt{3}T_1^2)$ . By setting the sampling densities of the rectangular and hexagonal patterns equal, we obtain the relationship

$$T_1 = \sqrt{\frac{2}{\sqrt{3}}}d. \quad (5)$$



**Fig 6** Frequency domain spectral packing using the parameters in Table 1 for (a) rectangular sampling and (b) hexagonal sampling. The circles represent isocontours of the OTF at 10% of the peak value.

The relationship in Eq. (5) is used to set the LR hexagonal grid in relation to the LR rectangular grid in Table 1.

As we saw in Sec. 2.1, the detector OTF for a hexagonal detector is not that much different from a comparably sized rectangular detector. The real benefit of hexagonal sampling is the result of efficient frequency packing and reduced aliasing. Isotropic rectangular sampling causes the continuous frequency spectrum to repeat in a rectangular pattern every  $1/d$  in the horizontal and vertical frequency axis. With hexagonal sampling, the continuous frequency spectrum repeats in a hexagonal pattern. The repetition occurs every  $2/T_1$  along the horizontal frequency axis and every  $2/T_2$  along the vertical frequency axis. That pattern is itself repeated with an offset of  $1/T_1$  and  $1/T_2$  along the respective frequency axes. The frequency packing for the parameters in Table 1 is shown in Fig. 6. The circles represent the isocontours of the OTF at a level of 10% of the maximum. Clearly the overlapping circles show that there is a large amount of potential aliasing in both the rectangular and hexagonal sampling cases. However, the intrusion of the repeating spectra around DC frequency at  $u = v = 0$  cycles/ $\mu\text{m}$  is demonstrably less in the case of the hexagonal sampling in Fig. 6(b) compared with rectangular in Fig. 6(a).

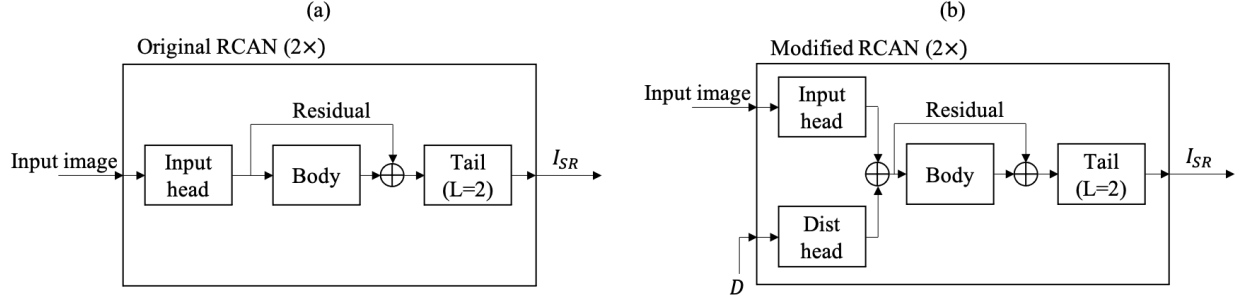
We can quantify the improved hexagonal frequency packing two different ways. First, consider that to sample exactly at the Nyquist rate with an isotropic rectangular sampling grid, we require  $d = 1/(2\rho_c)$ . For a hexagonal sensor, the equivalent requirement can be shown to be  $T_1 = 1/(\sqrt{3}\rho_c)$ , based on the frequency domain geometry. We can combine these relationships with those for sampling density. This tells us that the sampling density of a rectangular grid for Nyquist sampling is

$$S_{rect} = \frac{1}{d^2} = 4\rho_c^2. \quad (6)$$

The equivalent Nyquist sampling density for hexagonal sampling is given by

$$S_{hex} = \frac{2}{\sqrt{3}T_1^2} = \frac{6\rho_c^2}{\sqrt{3}}. \quad (7)$$





**Fig 7** The RCAN architecture for  $2\times$  upsampling. (a) The original architecture. (b) Modified architecture that adds a subpixel distance array to aid in hexagonal image SR.

Using these equations, we see that for the same  $\rho_c$ , the hexagonal spatial sampling grid can be 13.40% less dense. Another way to describe the potential benefits of hexagonal sampling is to consider the different cut-off frequency each sampling method can accommodate with no aliasing given the same sampling density. Performing this calculation using Eqs. (6) and (7), it can be shown that hexagonal sampling can accommodate a cut-off frequency that is 7.46% larger than rectangular. These theoretical advantages of hexagonal sampling are well known. However, to the best of our knowledge, the practical benefit of hexagonal sampling in light of modern processing techniques such as RCAN SR is heretofore untested. Thus, we believe it is very interesting to see if CNN based SR processing can exploit the advantages of hexagonal sampling in a significant way.

### 3 Super-Resolution for Hexagonally Sampled Images

In this section, we describe our approach for resampling and restoring rectangular and hexagonal imagery. Both approaches rely on the RCAN architecture<sup>17</sup> for SR. We briefly describe the RCAN architecture before defining the rectangular and hexagonal system.

#### 3.1 RCAN architecture

The RCAN architecture is a convolutional SR architecture for integer scale factors. It consists of stacks of learned convolutional filtering layers and upscaling. The architecture varies slightly depending on scale factor. When relevant, the architecture that performs  $L\times$  upsampling will be notated as RCAN ( $L\times$ ). A high-level diagram of RCAN ( $2\times$ ) is shown in Fig. 7(a). Our modified RCAN architecture is shown in Fig. 7(b), which we shall explain later in Sec. 3.2.1.

We denote the input of RCAN generally as  $I_{LR}$ . This input consists of the LR three-channel RGB color images. The standard input head of RCAN consists of a single convolutional layer that extracts feature channels from the input. That is,

$$F_0 = H_{\text{input}}(I_{LR}), \quad (8)$$

where  $H_{\text{input}}$  denotes the learned convolutional layer and  $F_0$  denotes the learned “shallow feature.” The body of RCAN uses Residual-In-Residual (RIR) deep feature extraction. The RIR body consists of  $G$  “residual groups” followed by a “long skip connection” (LSC). Denoting the  $g$ th residual group as  $H_g$  ( $1 \leq g \leq G$ ), we can formulate the RIR body as

$$F_{DF} = F_0 + \text{Conv}_{LSC} [H_G(H_{G-1}(\cdots H_1(F_0)\cdots))], \quad (9)$$

where  $\text{Conv}_{LSC}$  is a convolutional layer at the tail of the RIR.

Each residual group  $H_g$  is composed of  $N_{(g)}$  Residual Channel Attention Blocks (RCABs), a single convolutional layer, and a residual connection. Each RCAB first performs learned convolutions and then performs channel attention. Channel attention adaptively rescales channels in the feature map to exploit interdependencies across channels and emphasize important feature channels. The output of the channel attention mechanism is residually connected to the input of the RCAB. It is claimed that the channel attention mechanism allows the network to emphasize more useful features and makes the network more discriminative.<sup>17</sup>

The tail of the RCAN architecture consists of an upsampling module and a final restoration module. Note that the output is a three-channel rectangularly sampled RGB SR image. For upsampling, the RCAN architecture avoids the transposed convolution layers seen in encoder-decoder architectures<sup>8</sup> and opts for the light-weight ESPCN/PixelShuffle strategy.<sup>39</sup> An additional convolutional layer is used on the output. So, the SR image is formed as

$$I_{SR} = \text{Conv}_{out}(H_{up}(F_{DF})), \quad (10)$$

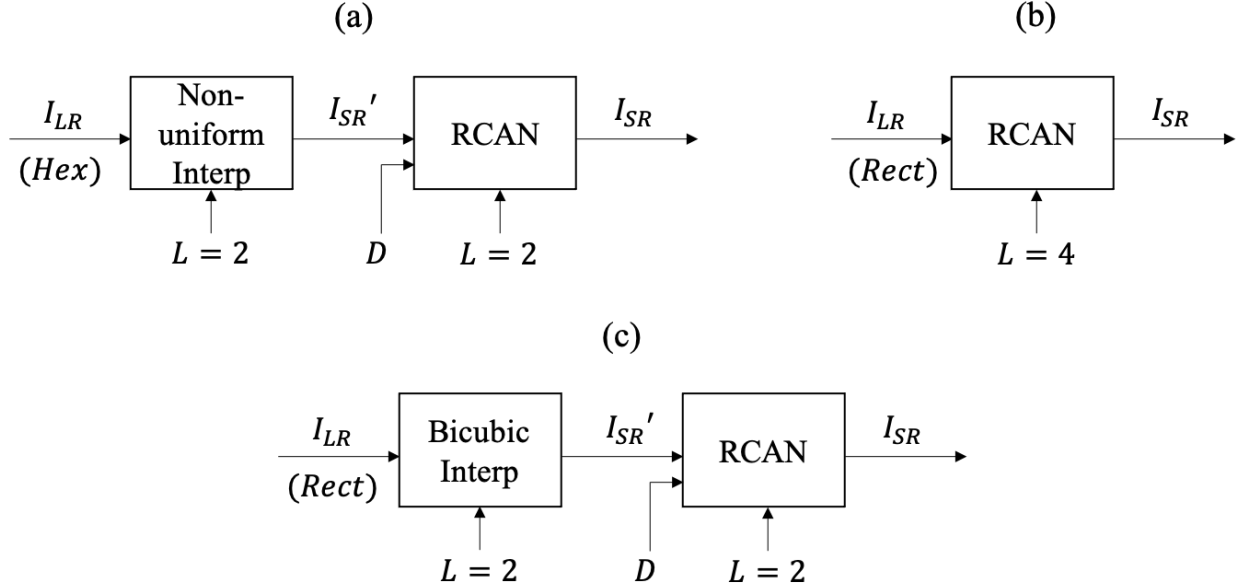
where  $H_{up}$  denotes the upsampling module and  $\text{Conv}_{out}$  denotes the output convolutional layer.

### 3.2 Hexagonal inverse model

For processing hexagonally sampled imagery, we develop a system that uses nonuniform interpolation to partially upsample the observed hexagonal imagery and convert it to a rectangular grid, and then feeds this imagery into an RCAN ( $2\times$ )-based network. The overall proposed system is shown in Fig. 8(a). As references, we also apply the standard RCAN ( $4\times$ ) architecture to rectangularly sampled imagery with the same sampling density as shown in Fig. 8(b) and we apply a directly comparable rectangular version of our system to the same rectangularly sampled imagery as shown in Fig. 8(c).

To motivate our system, ideally we would feed the LR hexagonal data directly into RCAN, and then RCAN would jointly resample and restore the image. Unfortunately, this is not practical with hexagonal imagery since the upsampling module in RCAN is designed to input and output images with the same aspect ratio and images on a hexagonal grid do not have the same ‘‘aspect ratio’’ as images on a rectangular grid (in terms of pixel counts). It is non-trivial to modify the upsampling module in RCAN to handle different aspect ratios and we are not aware of any existing methods to this end. Thus, we cannot simply feed the hexagonal image into the network and learn the inverse function. A naive solution would be to resample the hexagonal imagery onto a rectangular grid of the same sample density as the input. However, this solution would lose the extra spatial frequency content in the hexagonal samples.

Our system addresses both of these problems by using a two-step approach. In the first step, we use nonuniform interpolation to upsample  $I_{LR}$  onto a rectangular grid with twice the input image’s sample density. We call the upsampled image  $I'_{SR}$ . The sample spacing of  $I'_{SR}$  is  $d/2$ . Increasing the sampling density in this step increases the cutoff frequency of the grid and is sufficient to encode the frequency content of the hexagonal image (as limited by the hexagonal grid structure’s cutoff frequency). In the second step,  $I'_{SR}$  is fed into the trained RCAN ( $2\times$ ) network described in Sec. 3.1 and illustrated in Fig. 7(b). We train the network on  $(I'_{SR}, I_{SR})$  pairs to upsample  $I'_{SR}$  onto a rectangular grid with twice the sample density of  $I'_{SR}$  (i.e., a sample spacing of  $d/4$ ) while simultaneously correcting for interpolation error introduced in the first step of our system. The



**Fig 8** (a) The two-stage hexagonal SR system with LR hexagonal input and SR output. (b) Standard RCAN ( $4\times$ ) processing used for LR rectangular imagery. (c) The two-stage rectangular SR system.

training process is described in detail later, in Sec. 4. The output of the system overall is  $I_{SR}$ , which is the output of the RCAN as defined by Eq. (10).

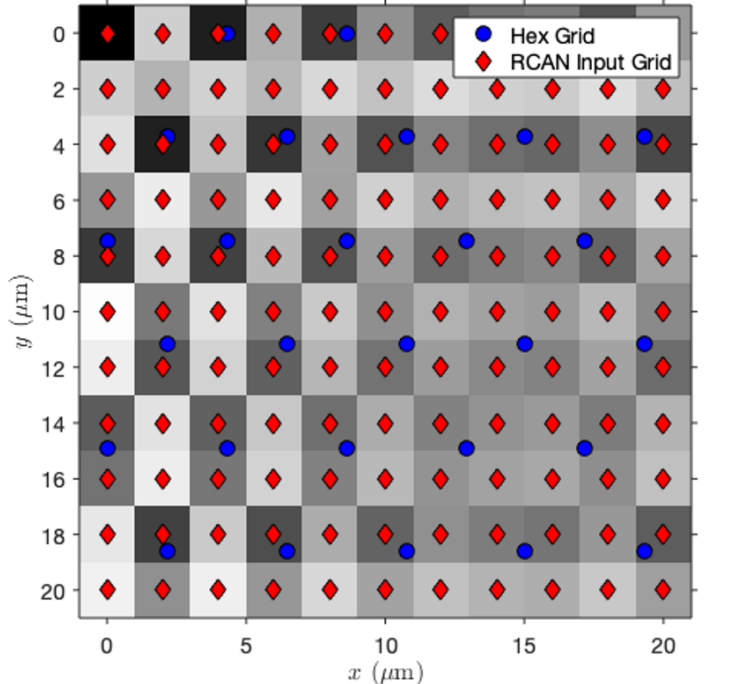
### 3.2.1 Sub-pixel distance matrix

The nonuniform interpolation step used by the method in Fig. 8(a) creates interpolation error that varies spatially. Some interpolated pixel locations are located very near the original hexagonal data and others are much farther. Inspired by the FIFNET sub-pixel registration channel,<sup>33</sup> we incorporate a sub-pixel distance matrix input to the network to quantify the alignment between the hexagonally sampled output of the observation model and the interpolated image that is fed into our modified RCAN shown in Figs. 8(a) and 7(b). This aids our modified RCAN system in dealing with the spatially varying interpolation error.

A formalization of the distance matrix follows. Let  $X$  be the set of spatial coordinates on the hexagonal lattice with elements  $\mathbf{x}_{i,j} \in \mathbb{R}^2$ . Similarly, let the spatial coordinates of the samples in the  $M \times N$  non-uniformly interpolated image be denoted  $\mathbf{r}_{m,n} \in \mathbb{R}^2$ . Then the distance matrix  $D \in \mathbb{R}^{M \times N}$  is defined such that

$$D_{m,n} = \min_{\mathbf{x}_{i,j} \in X} \|\mathbf{r}_{m,n} - \mathbf{x}_{i,j}\|_2. \quad (11)$$

In other words,  $D_{m,n}$  encodes how close the pixel  $m, n$  in the non-uniformly interpolated image is to the nearest sample in the hexagonally sampled input. A sample of the subpixel distance matrix is shown in Fig. 9. The hexagonal sampling grid,  $X$ , is shown in blue. The interpolated sampling grid is shown in red. The distance value for every interpolated sample,  $D_{m,n}$ , is indicated by the grayscale of the pixel behind the red sample. Note that red samples that are located near a blue sample have a small distance value (i.e., dark pixel). We expect these red samples to have less interpolation error. The opposite is true for red samples located far from the nearest blue sample.



**Fig 9** Sample subpixel distance matrix with the hexagonal grid and interpolation grid overlaid. The hexagonal sampling grid,  $X$ , is shown in blue. The rectangular sampling grid is shown in red. The distance array,  $D$ , is represented by the pixel grayscale values.

The distance matrix  $D$  is fed into RCAN via a separate input head. This separate input head consists of a stack of convolutional layers with a large receptive field. The output of these convolutional layers is connected to the output of the traditional input head such that Eq. (8) becomes

$$F_0 = H_{\text{input}}(I_{LR}) + H_{\text{dist}}(D), \quad (12)$$

where  $H_{\text{dist}}$  denotes the head for the sub-pixel distance matrix input. The precise architecture of  $H_{\text{dist}}$  is described in Table 2. The convolutions are padded and biases is used. Note that  $D$  is shown as an input to the system in the system flowchart in Fig. 8(a), and the distance head is shown in our modified RCAN architecture illustrated in Fig. 7(b).

**Table 2** Distance matrix head architecture.

Channels In	Channels Out	Filter Size
1	64	$7 \times 7$
64	64	$7 \times 7$
64	64	$7 \times 7$

## 4 Experimental Setup

In this section, we describe our methodology for comparing hexagonally sampled imagery and rectangularly sampled imagery with deep learning-based post-processing for super-resolution. We are interested in comparing our proposed system for processing hexagonally sampled imagery with systems for processing rectangularly sampled imagery for super-resolution in order to demonstrate the benefits of hexagonal sampling. We also consider several baseline systems.

- **Hex + NI (4×).** For hexagonally sampled data, we use nonuniform interpolation (NI) to upscale by a factor of 4 (with respect to the input sample density) onto a rectangular grid.
- **Rect + Bic (4×).** For rectangularly sampled data, we use bicubic interpolation (Bic) to upscale by a factor of 4.
- **Hex + NI (4×) + Wiener.** To the Hex + NI (x4) system, we add Wiener filtering<sup>40</sup> for restoration.
- **Rect + Bic (4×) + Wiener.** To the Rect + Bic (4×) system, we add Wiener filtering for restoration.
- **Hex + NI (2×) + RCAN (2×).** As proposed in Sec. 3, for hexagonally sampled data, we use NI to upscale by a factor of 2 (with respect to sample density) onto a rectangular grid and then feed the imagery into an RCAN (2×) network. (See Fig. 8(a).)
- **Rect + Bic (2×) + RCAN (2×).** To mirror Hex + NI (2×) + RCAN (2×) with rectangularly sampled data, we use bic to upscale by a factor of 2 and then feed the imagery into an RCAN (×2) network. (See Fig. 8(c).)
- **Rect + RCAN (4×).** For rectangularly sampled data, we use an RCAN (4×) network. (See Fig. 8(b).)

Nonuniform interpolation is implemented with MATLAB’s `scatteredInterpolant` function and bicubic interpolation is implemented with MATLAB’s `interp2` function. For the systems that use Wiener filtering, one Wiener filter<sup>40</sup> is used for each RGB channel using the observation model’s PSF for that channel. The noise-to-signal ratio (NSR) parameter of the Wiener filters are fit by finding an NSR parameter between  $10^{-5}$  and 1 that minimizes mean squared error (MSE) between the output of the Wiener filter and the desired image for every image in the training data (independently for each channel). The per-channel NSR values are averaged across the training data. The training methodology for RCAN networks and the evaluation protocol are described in the rest of this section.

#### 4.1 Data and preprocessing

For data, we require a high-resolution dataset with which to simulate realistic camera imagery using the observation model described in Sec. 2. We choose the DIV2K dataset,<sup>35</sup> which consists of 2K resolution natural imagery. The full DIV2K dataset consists of 800 training images, 100 validation images, and 100 testing images. However, SR truth for the testing images has not been publicly released. Following Ref. 15, we report out testing performance using a portion of the DIV2K validation set. In particular, our train split consists of DIV2K images 1-800, our validation split consists of DIV2K images 801-810, and our test split consists of DIV2K images 811-900. This approach ensures a clean separation of data for training, validation, and testing.

During preprocessing, the `scatteredInterpolant` MATLAB function is used to resample the hexagonally imagery to the rectangular grid in preprocessing for NI systems and the `interp2` MATLAB function is used to resample rectangularly sampled imagery for bic systems. We replace any NaN values that arise during this step with “0.” For data augmentation, random left-right flip, random up-down flip, and transposing are used to generate all eight possible combinations of flipping

and rotation. Data augmentation is applied *before* applying the observation model. That is, for each image in the original dataset, we transform it into eight images. The observation model is applied to each transformed image separately. As a result of this, each epoch of training for experiments with data augmentation consists of a pass through *all of the augmentation sets*, meaning every image in the DIV2K training images is seen 8 times per epoch.

#### 4.2 Training hyperparameters

Our RCAN networks are trained with default settings from the original RCAN network: we also use 10 residual groups and 10 residual blocks. We consider both randomly initializing our networks and initializing our networks with publicly available pre-trained RCAN (2×) and RCAN (4×) network parameters. These pre-trained networks were originally trained to upsample images that were downsampled with bicubic interpolation. It should be noted that it is not a priori obvious that using pre-trained weights is effective. We chose to consider pre-training because the original RCAN networks found that initializing the networks RCAN (3×) and RCAN (4×) with the pre-trained parameters from RCAN (2×) is effective, showing that the weights from one model are broadly applicable to slightly different scales. We consider whether a similar intuition applies when the scales are held constant but the task is slightly changed (namely, realistic degradations are applied to the imagery). This choice is considered in the ablation study in Sec. 5.3. The sub-pixel distance matrix input head described in Table 2 is randomly initialized in all cases.

A patch size of  $96 \times 96$  pixels is used for training. The batch size is set to 16 (patches). The Adam optimizer<sup>41</sup> is used with a learning rate of  $10^{-4}$ . The hexagonal model is trained for 200 epochs (1600 passes through the data). The learning rate is decayed by a factor of 0.5 at steps 120 and 160. It is found that a much shorter training schedule is appropriate for the rectangular model. By adjusting the schedule based on the validation learning curve, we arrived at training length of 23 epochs with decay at steps 8 and 15. Presumably, this is because the rectangular model, initialized with pre-trained weights, needs only to learn to adapt to the optical degradation whereas the hexagonal model needs to also undo resampling noise from the process interpolating rectangularly gridded samples.

Network training and testing is performed with an Ubuntu 18.04 workstation with an AMD Ryzen 7 2700X Eight-Core Processor with a maximum clock speed of 4 GHz. It is equipped with two NVIDIA GeForce RTX 2080 Ti Graphics Processing Units (GPUs). Both GPUs are used during training. Training the RCAN (2×) models as described takes approximately one day. Training the RCAN (4×) model takes approximately four hours. Testing on our 90 image test set takes approximately one minute.

#### 4.3 Loss function

We consider three loss functions: MSE, L1 loss, and the Charbonnier loss function. For a loss function, the traditional choice is MSE, since this effectively seeks to maximize the signal-to-noise ratio (SNR). The formula for MSE is given by

$$L(y, \hat{y}) = \text{mean}((y - \hat{y})^2), \quad (13)$$

where  $y$  is the true image,  $\hat{y}$  is the estimated image, the squaring is performed elementwise, and mean is taken across all pixels in each batch. While MSE is theoretically sound, empirical studies



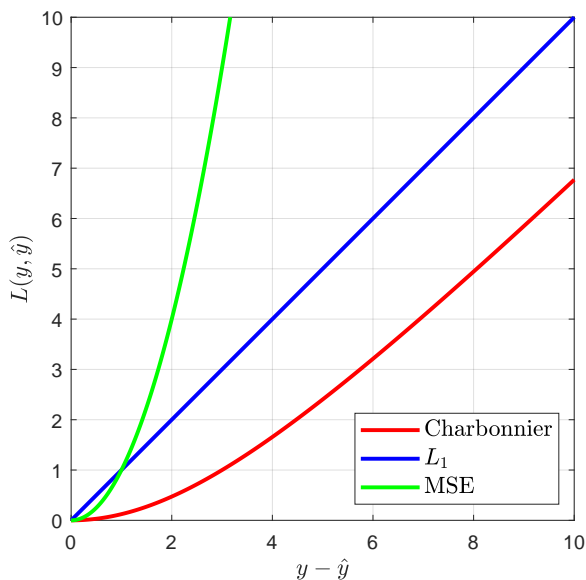
in Ref. 15 found the L1 loss to be superior to MSE for training EDSR and the original RCAN networks were trained with an L1 loss.<sup>17</sup> The formula for L1 loss is given by

$$L(y, \hat{y}) = \text{mean}(|y - \hat{y}|). \quad (14)$$

While MSE and L1 are common, we are inspired to consider a Charbonnier loss function used in the joint deblurring and super-resolution model ED-DSRN.<sup>20</sup> This loss function has been shown to perform well in a task that is similar to the one at hand here. The formula for the Charbonnier loss function is

$$L(y, \hat{y}) = \text{mean}(\sqrt{(y - \hat{y})^2 + \varepsilon^2} - \varepsilon). \quad (15)$$

A visual comparison of the Charbonnier loss with the L1 loss function and the MSE loss function is given in Fig. 10. The advantage of the Charbonnier loss function is that the function is smooth when error is close to 0 like MSE, but it grows linearly when the error is large like L1. We believe this improves global convergence during training. In our application, the intensity values of images  $y$  and  $\hat{y}$  are scaled to an approximate range of 0 to 255. For this range we have observed good results using  $\varepsilon = 4$ , and this value is used in all of our experiments.



**Fig 10** Plot of the Charbonnier loss as a function of the error term  $y - \hat{y}$  with  $\varepsilon = 4$  (red), the L1 loss function (blue), and the MSE loss function (green).

#### 4.4 Evaluation criteria

We consider two evaluation metrics: peak signal-to-noise ratio (PSNR) and Structural Similarity (SSIM) index.<sup>42</sup> We consider both metrics in two settings: First, we convert the image to grayscale by evaluating the Y channel of the YCbCr color space and evaluate our metrics. Second, we use full-color RGB imagery. In all cases, we shave 6 pixels from each edge of the image to avoid

evaluating border effects. Note that our results are not directly comparable with other prior results in the literature because of the use of the camera-parameter based observation model employed here.<sup>33</sup> We believe the observation model provides more realistic data, but it is not yet widely used in the literature. Note we report PSNR values in dB.

## 5 Experimental Results

In this section, we describe the results of the hexagonal and rectangular image processing systems on our test split (DIV2K images 811-900) as well as provide an ablation study of our architectural choices. We provide both quantitative and qualitative analysis. We consider the saved checkpoint with the best validation PSNR on the Y channel.

### 5.1 Quantitative Results

We calculate the PSNR and SSIM as described in Sec. 4.4. The results are given in Table 3.

**Table 3** Results on our test split (DIV2K images 811-900). “NI” refers to nonuniform interpolation and “Bic” refers to bicubic interpolation. The bold face entries represent the best results for the given metric.

	Y channel		RGB	
	PSNR (dB)	SSIM	PSNR (dB)	SSIM
Hex + NI (4×)	25.424	0.70895	23.944	0.69016
Hex + NI (4×) + Wiener	26.460	0.75731	24.958	0.73770
Hex + NI (2×) + RCAN (2×)	<b>29.204</b>	<b>0.84425</b>	<b>27.692</b>	<b>0.82957</b>
Rect + Bic (4×)	25.826	0.72621	24.348	0.70796
Rect + Bic (4×) + Wiener	26.060	0.73410	24.587	0.71642
Rect + Bic (2×) + RCAN (2×)	29.004	0.83735	27.494	0.82254
Rect + RCAN (4×)	29.073	0.83939	27.556	0.82442

On our test split, our proposed Hex + NI (2×) + RCAN (2×) shown in Fig. 8(a) achieves a PSNR of 29.204 and an SSIM of 0.84425 when evaluated on the Y channel and 27.692 and 0.82957, respectively, when evaluated on full RGB. The equivalent system using rectangularly sampled imagery, Rect + Bic (2×) + RCAN (2×), achieves a PSNR of 29.004 and an SSIM of 0.83735 on the Y channel and 27.494 and 0.82254, respectively, when evaluated on full RGB. The purely RCAN-based Rect + RCAN (4×) system shown in Fig. 8(b) achieves a PSNR of 29.073 and an SSIM of 0.83939 on the Y channel and 27.556 and 0.82442, respectively, on full RGB. This means that the hexagonal sampled imagery outperforms the rectangular sampled imagery, whether the rectangular imagery is processed by the end-to-end deep learning Rect + RCAN (4×) system (+.131 PSNR and +.00486 SSIM on Y channel, and +.136 and +.00515, respectively, on RGB) or by the two-stage Rect + Bic (2×) + RCAN (2×) system that mirrors the hexagonal processing (+0.200 PSNR and +.00690 SSIM on Y channel, and +.198 and +.00703, respectively, on RGB). Moreover, this shows that for rectangular imagery, using only RCAN as in the Rect + RCAN (4×) system is superior to using bicubic interpolation with RCAN as in the Rect + Bic (2×) + RCAN (2×) system (+0.069 PSNR and +.00204 SSIM on Y channel, and +0.062 and +.00188, respectively, on RGB). To ground these differences, note that the qualitative results shown in Fig. 11 represent a difference of 0.281 PSNR and 0.008896 SSIM between the center and right columns, which are similar magnitudes as these averaged results. As can be seen in Table 3,

the trends observed in the quantitative Y channel results are very similar to the trends in the RGB results, so we focus on Y channel results for the remainder of this section.

Comparing our RCAN systems against the baseline interpolation and interpolation + Wiener systems in Table 3, our RCAN-based post-processing achieves higher performance. Every RCAN-based system outperforms every baseline system by at least +2.544 PSNR and .08004 SSIM. Comparing baseline systems across sampling type in Table 3, we find that  $4\times$  interpolation of rectangularly sampled imagery is better than  $4\times$  interpolation of hexagonally sampled imagery. To be precise, Rect + Bic ( $4\times$ ) outperforms Hex + NI ( $4\times$ ) by .402 PSNR and 0.01726 SSIM. However, post-processing with a Wiener filter shows the benefits of hexagonally sampled imagery: Hex + NI ( $4\times$ ) + Wiener outperforms Rect + Bic ( $4\times$ ) + Wiener by .400 PSNR and .02321 SSIM.

## 5.2 Qualitative Results

For a qualitative comparison between our proposed Hex + NI ( $2\times$ ) + RCAN ( $2\times$ ) system and the Rect + RCAN ( $4\times$ ) system, close-ups of images from our test split are provided in Figs. 11-16. In all of these figures the truth is shown in the left column, the output of Hex + NI ( $2\times$ ) + RCAN ( $2\times$ ) processing is shown in the middle column, and Rect + RCAN ( $4\times$ ) outputs are shown in the rightmost column. The full images are shown in the top row and regions of interest (ROIs) are highlighted in subsequent detail rows. Themes throughout these images include more faithful rendering of text (Figs. 12, 15), improved rendering of support beams in small windows and doors (Figs. 11, 13, 14), and more faithfully capturing textures with high-spatial-frequency lines (Figs. 15, 16).

## 5.3 Ablation Study

In this section, we present results from our ablation study. Since these experiments were actually carried out during the model selection phase, only Y channel scores on our validation split (DIV2K images 801-810) are reported.

### 5.3.1 Hexagonal system

For the hexagonal model’s ablation, we are interested in the effects of the loss function, the sub-pixel distance matrix, the augmentation scheme, and the choice of pre-trained weights (scratch vs RCAN ( $2\times$ ) vs RCAN ( $4\times$ )). The results of these experiments are described in Table 4.

Regarding choice of the loss functions described in Sec. 4.3, we find that the Charbonnier loss function is the most effective. The L1 loss function imposes a penalty of 0.004 in PSNR for a score of 29.578. The MSE loss function is found to be much worse, imposing a penalty of 0.031 in PSNR for a score of 29.551. For the remainder of the experiments, we consider the Charbonnier loss function.

Going down the rows of Table 4, in the final block we successively remove features from the architecture. First, we remove the distance head described in Sec. 3.2.1. A penalty of 0.003 in PSNR is observed for a score of 29.579. Next, we do not use data augmentation and adjust the schedule somewhat to account for the fact that each epoch now contains fewer images (see Sec. 4). A further penalty of 0.132 in PSNR is observed for a score of 29.447.

Next, we consider the effects of using pre-trained weights as discussed in Sec. 4. First, we consider using the pre-trained RCAN ( $4\times$ ) weights instead of the pre-trained RCAN ( $2\times$ ) weights. This tests the effect of using *some* pre-trained weights, even if the weights are sub-optimal. Note



**Fig 11** DIV2K Image 823. The truth is shown in the left column, the output of Hex + NI ( $2\times$ ) + RCAN ( $2\times$ ) processing is shown in the middle column, and Rect + RCAN ( $4\times$ ) outputs are shown in the rightmost column. On the first detail row, the hexagonal model captures some notion of a horizontal line corresponding to the decks of the ship, which is absent from the rectangular model. On the second detail row, the hexagonal model renders the brown door better than the rectangular model.




**Fig 12** DIV2K Image 825. On the first detail row, the “D” and “E” are rendered more faithfully by the hexagonal model than the rectangular. On the second detail row, the yellow band in the flag is rendered more faithfully by the hexagonal system. On the final detail row, the top row of letters (particularly the “P,” “D,” “R,” and “O”) are more faithfully rendered by the hexagonal system.

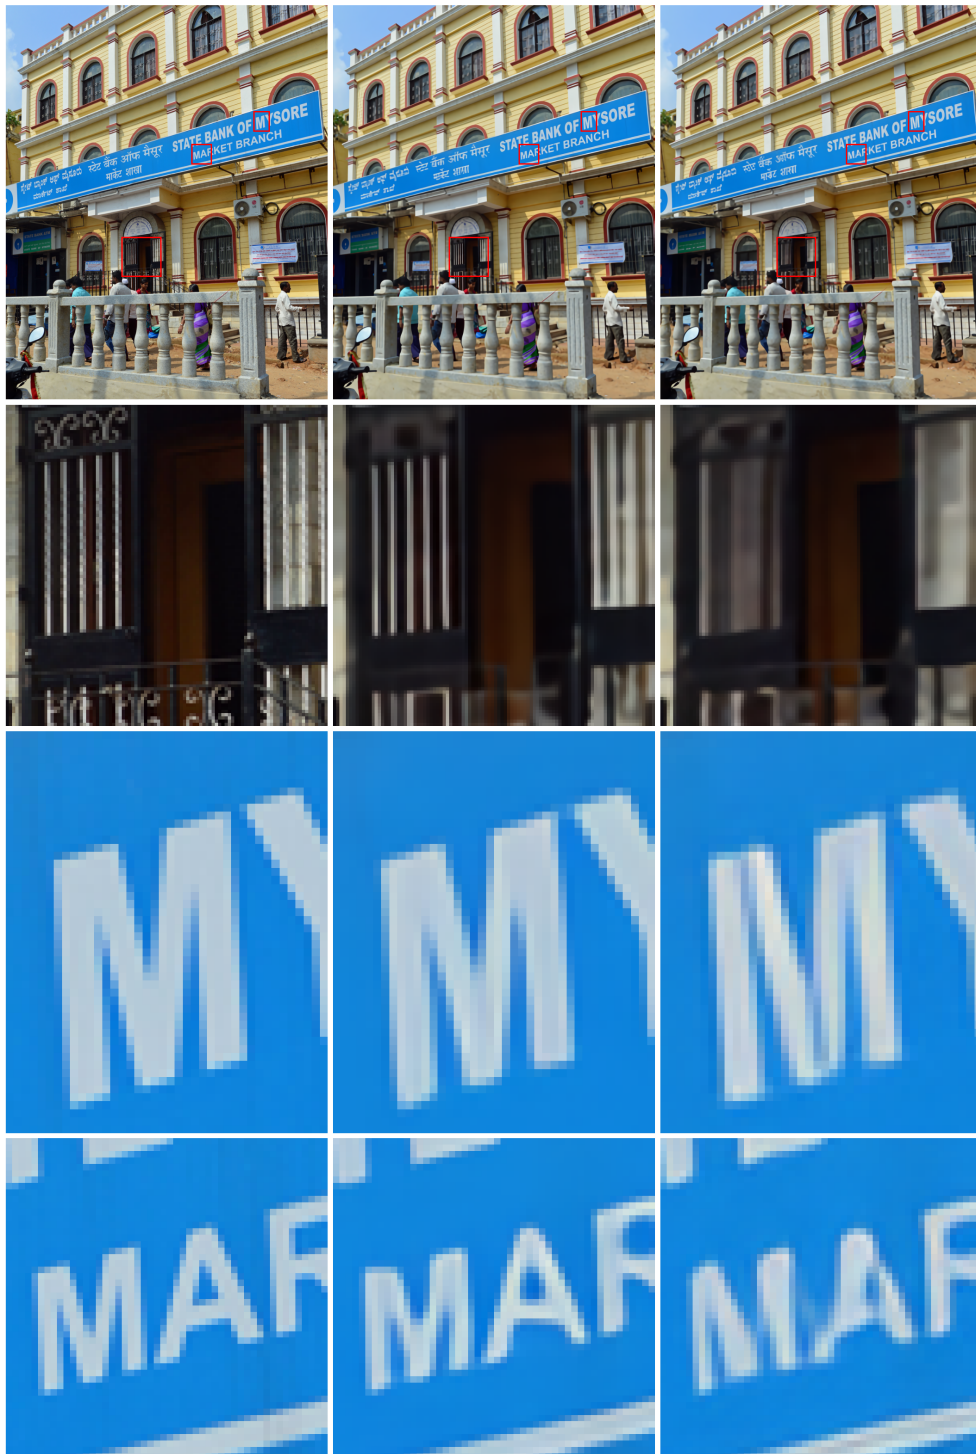


**Fig 13** DIV2K Image 830. On the first detail row, the hexagonal model captures all vertical supports while the rectangular does not. On the second detail row, the rectangular model mis-colors the band between the lines on the building while the hexagonal model does not. On the third detail row, the hexagonal model captures the vertical support in the window while the rectangular model renders a blur. On the fourth detail row, the hexagonal model generates interference patterns in the water whereas the rectangular model is dominated by one direction of line. On the fifth detail row, the hexagonal rendering is clear whereas the rectangular rendering is blurry. On the sixth detail row, the hexagonal model captures more detail between the orange and the gray areas than the rectangular model.






**Fig 14** DIV2K Image 836. On both rows, the hexagonal model captures more detail of the vertical and horizontal supports of windows.



**Fig 15** DIV2K Image 891. In the first detail row, the hexagonal model captures the detail of vertical supports in a window more faithfully than the rectangular model. In the second and third detail rows, the hexagonal model more faithfully renders letters in text.



**Fig 16** DIV2K Image 900. The hexagonal system faithfully renders the vertical lines in the grill of the car whereas the rectangular image is clearly aliased.

**Table 4** Hexagonal system ablation study.

Name	Schedule	PT	Aug	Dist	Loss	PSNR (dB)
Hex	{120, 160, 200}	x2	Yes	Yes	Charb	29.582
Hex L1	{120, 160, 200}	x2	Yes	Yes	L1	29.578
Hex MSE	{120, 160, 200}	x2	Yes	Yes	MSE	29.551
-Dist	{120, 160, 200}	x2	Yes	No	Charb	29.579
-Aug	{300, 400, 501}	x2	No	No	Charb	29.447
-x2+x4	{300, 400, 501}	x4	No	No	Charb	29.432
-PT	{300, 400, 501}	No	No	No	Charb	29.211

the pixel shuffle layer is randomly initialized for this experiment since the  $4\times$  pixel shuffle layer would return the wrong image size. A penalty of 0.015 in PSNR is observed from the last step in the ablation for a score of 29.432. This stands in contrast to no pre-training at all, where a further penalty of 0.221 in PSNR is observed for a score of 29.211. So, using sub-optimal pretrained weights is found to be better than training from scratch. This is consistent with the result from Ref. 17, where using one scale of RCAN weights was found to be effective for initializing a training for another scale.

To recap the results of the hexagonal ablation study, we find performance drops the most when we remove augmentation ( $-0.132$  in PSNR) and when we train from scratch ( $-0.221$  in PSNR). Consequently, we believe that using pre-trained weights (despite the weights not representing how

**Table 5** Rectangular model ablation study.

Name	Schedule	PT	Aug	Loss	PSNR (dB)
Rect	{8, 15, 23}	x4	Yes	Charb	29.556
+Long	{120, 160, 200}	x4	Yes	Charb	29.543
-x4+x2	{8, 15, 23}	x2	Yes	Charb	29.446
-PT	{120, 160, 200}	No	Yes	Charb	29.306
-Aug	{300, 400, 501}	No	No	Charb	29.010

to denoise the image) and using data augmentation are especially important to our success. We note, however, that our adaptation of the training schedule for data augmentation is imperfect, which makes it more difficult to directly compare experiments before and after removing data augmentation. A better schedule for non-augmented training would be {960, 1280, 1600} to match the number of examples seen in total during training with augmentation. However, since less data is seen overall, such a long schedule is not necessarily required in practice and we do note that the learning curve still appears to be mainly converged.

### 5.3.2 Rectangular system

Next, we perform an ablation study for our rectangular system. Note that we only perform ablation for Rect + RCAN (4×) and not Rect + Bic (2×) + RCAN (2×) because it provided the best performance. Here, there are fewer design decisions to be made. In particular, there is no use for a sub-pixel distance matrix since every sample is perfectly aligned with the sample grid (i.e., there is no interpolation noise). On the other hand, we consider the effects of shortening the training schedule (see Sec. 4). The ablation study is summarized in Table 5.

Going down the rows of Table 5, we first consider the effect of training schedule. When we train for the same length as the hexagonal model, a penalty of 0.013 in PSNR is observed for a score of 29.543. Next, we consider the effects of pre-training in the same manner as with the hexagonal model. When we swap to sub-optimal weights, a penalty of 0.097 in PSNR is incurred for a score of 29.446. When we train from scratch, a further penalty of 0.140 in PSNR is incurred for a score of 29.306. Finally, when we remove augmentation, we incur another penalty of 0.296 in PSNR for a score of 29.010. Again, we find that using data augmentation and using some form of pre-training are especially significant to our performance.

## 6 Discussion

We demonstrate an approach to resampling hexagonal imagery to a rectangular grid that retains the benefits of hexagonally sampled imagery. This approach uses non-uniform interpolation to resample the hexagonal data onto an upsampled rectangular grid. We then apply RCAN, a state-of-the-art deep learning-based super-resolution system, for restoration and further SR processing. To aid the RCAN processing, we also code the distance of each interpolated sample to the nearest original hexagonal sample and provide this as auxiliary input features to RCAN. To demonstrate the approach, we use a camera degradation model that simulates diffraction and detector integration of a hexagonal camera given a high-resolution rectangularly sampled "truth" image.

On the DIV2K SR dataset, we find that the hexagonal system outperforms applying RCAN directly on rectangularly sampled imagery with the same sample density, no interpolation noise, and modeled rectangular camera degradations in terms of PSNR and SSIM. Images from the holdout



set are provided that show clearly visible differences between the outputs generated by the hexagonal system and by the rectangular system. We also find that the hexagonal system outperforms an identical two-stage system that ingests and is trained on realistic rectangular imagery.

Theoretically, hexagonal sampling encodes more frequency information. We posit that this theoretical benefit of hexagonal sampling is significant for our success in achieving superior SR performance. We would expect to see some amount of benefit to hexagonal sampling for a wide range of sensor parameters, provided that the corresponding rectangular camera system is under-sampled. However, we live in a predominantly rectangularly sampled world: computer monitors, cell phone screens, and printers all usually use rectangularly sampled images. As a result, the typical tools for processing imagery assume rectangular sampling grids. We believe that our system provides a bridge between hexagonal samples and rectangular processing that reduces the loss incurred by converting to the rectangular sampling scheme.

Future work should consider directly learning to jointly resample, restore, and increase the resolution of the imagery within one end-to-end deep learning system. Since we report that the rectangular system based on RCAN (4×) outperforms the two-stage rectangular system significantly, we posit that an end-to-end system for hexagonal imagery will be similarly more performant. Developing such a system is a challenging problem because there is no clear-cut method for resampling the hexagonal imagery to a rectangular grid with convolutions or a PixelShuffle layer. Theoretically, however, such a method is advantageous because it avoids introducing interpolation noise in the resampling step. Hopefully the availability of libraries like HexagDLy<sup>28</sup> will inspire future researchers to pursue this problem.

### *Disclosures*

That authors declare that there are no relevant financial interests in the manuscript and no other potential conflicts of interest to disclose.

### *Code, Data, and Materials Availability*

The truth dataset used in our study is publicly available and described in Ref. 35. The author’s code and degraded image are not publicly available at this time, but are available from the authors upon reasonable request.

### *References*

- 1 W. W. W. Zou and P. C. Yuen, “Very low resolution face recognition problem,” *IEEE Transactions on Image Processing* **21**(1), 327–340 (2012).
- 2 P. Li, L. Prieto, D. Mery, *et al.*, “Face recognition in low quality images: A survey,” *ArXiv abs/1805.11519* (2018).
- 3 W. Shi, J. Caballero, C. Ledig, *et al.*, “Cardiac image super-resolution with global correspondence using multi-atlas patchmatch,” in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2013*, K. Mori, I. Sakuma, Y. Sato, *et al.*, Eds., 9–16, Springer Berlin Heidelberg, (Berlin, Heidelberg) (2013).
- 4 M. S. M. Sajjadi, B. Schölkopf, and M. Hirsch, “Enhancenet: Single image super-resolution through automated texture synthesis,” in *2017 IEEE International Conference on Computer Vision (ICCV)*, 4501–4510 (2017).

- 5 G. Ongie, A. Jalal, C. A. Metzler, *et al.*, “Deep learning techniques for inverse problems in imaging,” *IEEE Journal on Selected Areas in Information Theory* **1**(1), 39–56 (2020).
- 6 C. A. Dávila and B. R. Hunt, “Training of a neural network for image superresolution based on a nonlinear interpolative vector quantizer,” *Appl. Opt.* **39**, 3473–3485 (2000).
- 7 C. Dong, C. C. Loy, K. He, *et al.*, “Image super-resolution using deep convolutional networks,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **38**, 295–307 (2016).
- 8 C. Dong, C. C. Loy, and X. Tang, “Accelerating the super-resolution convolutional neural network,” in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, *et al.*, Eds., 391–407, Springer International Publishing, (Cham) (2016).
- 9 Z. Wang, D. Liu, J. Yang, *et al.*, “Deep networks for image super-resolution with sparse prior,” in *2015 IEEE International Conference on Computer Vision (ICCV)*, 370–378 (2015).
- 10 J. Kim, J. K. Lee, and K. M. Lee, “Accurate image super-resolution using very deep convolutional networks,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1646–1654 (2016).
- 11 Y. Tai, J. Yang, and X. Liu, “Image super-resolution via deep recursive residual network,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2790–2798 (2017).
- 12 W.-S. Lai, J.-B. Huang, N. Ahuja, *et al.*, “Deep laplacian pyramid networks for fast and accurate super-resolution,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5835–5843 (2017).
- 13 W.-S. Lai, J.-B. Huang, N. Ahuja, *et al.*, “Fast and accurate image super-resolution with deep laplacian pyramid networks,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **41**(11), 2599–2613 (2019).
- 14 Y. Tai, J. Yang, X. Liu, *et al.*, “Memnet: A persistent memory network for image restoration,” in *2017 IEEE International Conference on Computer Vision (ICCV)*, 4549–4557 (2017).
- 15 B. Lim, S. Son, H. Kim, *et al.*, “Enhanced deep residual networks for single image super-resolution,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 1132–1140 (2017).
- 16 K. Zhang, W. Zuo, and L. Zhang, “Learning a single convolutional super-resolution network for multiple degradations,” in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3262–3271 (2018).
- 17 Y. Zhang, K. Li, K. Li, *et al.*, “Image super-resolution using very deep residual channel attention networks,” in *Computer Vision – ECCV 2018*, V. Ferrari, M. Hebert, C. Sminchisescu, *et al.*, Eds., 294–310, Springer International Publishing, (Cham) (2018).
- 18 Z. Wang, J. Chen, and S. C. Hoi, “Deep learning for image super-resolution: A survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* , 1–1 (2020).
- 19 S. Nah, T. H. Kim, and K. M. Lee, “Deep multi-scale convolutional neural network for dynamic scene deblurring,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 257–265 (2017).
- 20 X. Zhang, F. Wang, H. Dong, *et al.*, “A deep encoder-decoder networks for joint deblurring and super-resolution,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1448–1452 (2018).



- 21 X. Zhang, H. Dong, Z. Hu, *et al.*, “Gated fusion network for degraded image super resolution,” *International Journal of Computer Vision* **128**, 1699–1721 (2020).
- 22 W. Sun and Z. Chen, “Learned image downscaling for upscaling using content adaptive resampler,” *IEEE Transactions on Image Processing* **29**, 4027–4040 (2020).
- 23 E. Hoogeboom, J. W. Peters, T. S. Cohen, *et al.*, “Hexaconv,” in *International Conference on Learning Representations*, (2018).
- 24 T. S. Cohen and M. Welling, “Group equivariant convolutional networks,” in *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML’16*, 2990–2999, JMLR.org (2016).
- 25 G.-S. Xia, J. Hu, F. Hu, *et al.*, “Aid: A benchmark data set for performance evaluation of aerial scene classification,” *IEEE Transactions on Geoscience and Remote Sensing* **55**(7), 3965–3981 (2017).
- 26 J. Luo, W. Zhang, J. Su, *et al.*, “Hexagonal convolutional neural networks for hexagonal grids,” *IEEE Access* **7**, 142738–142749 (2019).
- 27 C. Zhang, S. Liwicki, W. Smith, *et al.*, “Orientation-aware semantic segmentation on icosahedron spheres,” in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 3532–3540 (2019).
- 28 C. Steppa and T. L. Holch, “Hexagdly—processing hexagonally sampled data with cnns in pytorch,” *SoftwareX* **9**, 193 – 198 (2019).
- 29 A. Paszke, S. Gross, F. Massa, *et al.*, “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, *et al.*, Eds., 8024–8035, Curran Associates, Inc. (2019).
- 30 M. Jacquemont., L. Antiga., T. Vuillaume., *et al.*, “Indexed operations for non-rectangular lattices applied to convolutional neural networks,” in *Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications - Volume 5: VISAPP*, 362–371, INSTICC, SciTePress (2019).
- 31 J. Ke, H. Yang, H. Zheng, *et al.*, “Hexagon-based convolutional neural network for supply-demand forecasting of ride-sourcing services,” *IEEE Transactions on Intelligent Transportation Systems* **20**(11), 4160–4173 (2019).
- 32 R. Burada, “Joint resampling and restoration of hexagonally sampled images using adaptive wiener filter,” Master’s thesis, University of Dayton (2015).
- 33 H. Elwarfalli and R. Hardie, “Fifnet: A convolutional neural network for motion-based multiframe super-resolution using fusion of interpolated frames,” *Computer Vision and Image Understanding* **202**, 103097 (2021).
- 34 R. Hardie, D. R. Droege, A. J. Dapore, *et al.*, “Impact of detector-element active-area shape and fill factor on super-resolution,” *Frontiers in Physics* **3**, 31 (2015).
- 35 E. Agustsson and R. Timofte, “Ntire 2017 challenge on single image super-resolution: Dataset and study,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, (2017).
- 36 J. W. Goodman, *Introduction to Fourier optics*, Roberts and Company Publishers (2005).
- 37 G. D. Boreman, *Basic electro-optics for electrical engineers*, vol. 31, SPIE Press (1998).
- 38 L. Middleton and J. Sivaswamy, *Hexagonal image processing: A practical approach*, Springer Science & Business Media (2006).

- 39 W. Shi, J. Caballero, F. Huszár, *et al.*, “Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1874–1883 (2016).
- 40 R. C. Gonzalez and R. E. Woods, *Digital Image Processing (3rd Edition)*, Prentice-Hall, Inc., Upper Saddle River, NJ, USA (2006).
- 41 D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” (2017).
- 42 Z. Wang, A. Bovik, H. Sheikh, *et al.*, “Image quality assessment: from error visibility to structural similarity,” *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004).

**Dylan Flaute** graduated *magna cum laude* with a Bachelor’s degree in electrical engineering and a Bachelor’s degree in mathematics from the University of Dayton in December 2020. There, he was awarded The Thomas R. Armstrong, ’38 Award of Excellence for Outstanding Electrical Engineering Achievement from the School of Engineering and The Senior Award for Academic Excellence in Mathematics from the Department of Mathematics. He is currently a Junior Autonomy Engineer at Jacobs, where he works with the Air Force Research Laboratory’s ACT3 team. His research interests include digital signal and image processing and machine learning.

**Russell C. Hardie** is a Full Professor in Department of Electrical and Computer Engineering at the University of Dayton, and holds a joint appointment in the Department of Electro-Optics and Photonics. Dr. Hardie received the University of Dayton’s top university-wide teaching award in 2006 (the Alumni Award in Teaching). He also received the Rudolf Kingslake Medal and Prize from SPIE in 1998 for work on super-resolution imaging. His research interests include a wide range of topics in signal and image processing, image restoration, medical image analysis, and machine learning.

**Hamed Elwarfalli** is a Ph.D. student in the Department of Electrical and Computer Engineering at the University of Dayton, where he received his Master of Science degree in Electrical Engineering in 2017. He received a Bachelor’s degree in Electrical Engineering from the College of Electronics Technology, Libya, in 2010. His research interests include image processing and machine learning.

## List of Figures

- 1 Discrete forward observation model used to provide realistic degradation of ideal images for training and testing purposes. A pristine HR image is blurred with an impulse-invariant PSF and then resampled according the detector sampling pattern. Finally, noise is added to produce the LR image.
- 2 Detector shapes based on the parameters in Table 1 for (a) rectangular detectors and (b) hexagonal detectors.
- 3 Diffraction OTF model from Eq. (2) for the green channel parameters in Table 1.
- 4 Detector OTFs from for (a) rectangular detectors and (b) hexagonal detectors. Combined OTFs from Eq. (1) for (c) rectangular detectors and (c) hexagonal detectors.

- 5 Sampling grids used in the observation model. The HR grid (black) is a dense rectangular Nyquist sampling configuration. The LR rectangular sampling grid and associated detectors are shown in blue. The LR hexagonal sampling grid and associated detectors are shown in red.
- 6 Frequency domain spectral packing using the parameters in Table 1 for (a) rectangular sampling and (b) hexagonal sampling. The circles represent isocontours of the OTF at 10% of the peak value.
- 7 The RCAN architecture for  $2\times$  upsampling. (a) The original architecture. (b) Modified architecture that adds a subpixel distance array to aid in hexagonal image SR.
- 8 (a) The two-stage hexagonal SR system with LR hexagonal input and SR output. (b) Standard RCAN ( $4\times$ ) processing used for LR rectangular imagery. (c) The two-stage rectangular SR system.
- 9 Sample subpixel distance matrix with the hexagonal grid and interpolation grid overlaid. The hexagonal sampling grid,  $X$ , is shown in blue. The rectangular sampling grid is shown in red. The distance array,  $D$ , is represented by the pixel grayscale values.
- 10 Plot of the Charbonnier loss as a function of the error term  $y - \hat{y}$  with  $\varepsilon = 4$  (red), the L1 loss function (blue), and the MSE loss function (green).
- 11 DIV2K Image 823. The truth is shown in the left column, the output of Hex + NI ( $2\times$ ) + RCAN ( $2\times$ ) processing is shown in the middle column, and Rect + RCAN ( $4\times$ ) outputs are shown in the rightmost column. On the first detail row, the hexagonal model captures some notion of a horizontal line corresponding to the decks of the ship, which is absent from the rectangular model. On the second detail row, the hexagonal model renders the brown door better than the rectangular model.
- 12 DIV2K Image 825. On the first detail row, the “D” and “E” are rendered more faithfully by the hexagonal model than the rectangular. On the second detail row, the yellow band in the flag is rendered more faithfully by the hexagonal system. On the final detail row, the top row of letters (particularly the “P,” “D,” “R,” and “O”) are more faithfully rendered by the hexagonal system.
- 13 DIV2K Image 830. On the first detail row, the hexagonal model captures all vertical supports while the rectangular does not. On the second detail row, the rectangular model mis-colors the band between the lines on the building while the hexagonal model does not. On the third detail row, the hexagonal model captures the vertical support in the window while the rectangular model renders a blur. On the fourth detail row, the hexagonal model generates interference patterns in the water whereas the rectangular model is dominated by one direction of line. On the fifth detail row, the hexagonal rendering is clear whereas the rectangular rendering is blurry. On the sixth detail row, the hexagonal model captures more detail between the orange and the gray areas than the rectangular model.
- 14 DIV2K Image 836. On both rows, the hexagonal model captures more detail of the vertical and horizontal supports of windows.

- 15 DIV2K Image 891. In the first detail row, the hexagonal model captures the detail of vertical supports in a window more faithfully than the rectangular model. In the second and third detail rows, the hexagonal model more faithfully renders letters in text.
- 16 DIV2K Image 900. The hexagonal system faithfully renders the vertical lines in the grill of the car whereas the rectangular image is clearly aliased.

## List of Tables

- 1 Summary of observation model parameters for the red, green, and blue color channels.
- 2 Distance matrix head architecture.
- 3 Results on our test split (DIV2K images 811-900). “NI” refers to nonuniform interpolation and “Bic” refers to bicubic interpolation. The bold face entries represent the best results for the given metric.
- 4 Hexagonal system ablation study.
- 5 Rectangular model ablation study.