

High-Throughput Molecular Imaging via Deep-Learning-Enabled Raman Spectroscopy

Conor C. Horgan, Magnus Jensen, Anika Nagelkerke, Jean-Philippe St-Pierre, Tom Vercauteren, Molly M. Stevens,* and Mads S. Bergholt*

Cite This: *Anal. Chem.* 2021, 93, 15850–15860

Read Online


ACCESS |

Metrics & More

Article Recommendations

Supporting Information

ABSTRACT: Raman spectroscopy enables nondestructive, label-free imaging with unprecedented molecular contrast, but is limited by slow data acquisition, largely preventing high-throughput imaging applications. Here, we present a comprehensive framework for higher-throughput molecular imaging via deep-learning-enabled Raman spectroscopy, termed DeepeR, trained on a large data set of hyperspectral Raman images, with over 1.5 million spectra (400 h of acquisition) in total. We first perform denoising and reconstruction of low signal-to-noise ratio Raman molecular signatures via deep learning, with a 10× improvement in the mean-squared error over common Raman filtering methods. Next, we develop a neural network for robust 2–4× spatial super-resolution of hyperspectral Raman images that preserve molecular cellular information. Combining these approaches, we achieve Raman imaging speed-ups of up to 40–90×, enabling good-quality cellular imaging with a high-resolution, high signal-to-noise ratio in under 1 min. We further demonstrate Raman imaging speed-up of 160×, useful for lower resolution imaging applications such as the rapid screening of large areas or for spectral pathology. Finally, transfer learning is applied to extend DeepeR from cell to tissue-scale imaging. DeepeR provides a foundation that will enable a host of higher-throughput Raman spectroscopy and molecular imaging applications across biomedicine.



INTRODUCTION

Raman spectroscopy has recently excelled as a highly complementary tool for biomedical research, providing non-destructive, label-free, molecular imaging with subcellular resolution. This has enabled a multitude of exciting biomedical applications from fundamental *in vitro* cellular studies^{1,2} to *ex vivo* spectral histopathology^{3,4} and *in vivo* fiber-optic endoscopy for optical biopsy at the molecular level.^{5,6} Despite its many advantages, Raman spectroscopy remains limited by the weakness of generated Raman signals, which necessitates spectral acquisition times on the order of 1 s per sampled point.⁷ As such, high-resolution Raman spectroscopic imaging of cells or tissues often requires multiple hours, which is prohibitive for high-throughput Raman spectroscopic imaging applications.^{8–10} To address these acquisition time and signal-to-noise ratio (SNR) challenges, advanced nonlinear Raman spectroscopy techniques, including coherent anti-Stokes Raman spectroscopy (CARS) and stimulated Raman spectroscopy (SRS) have been developed.^{11–13} However, while these advanced techniques have recently enabled rapid broadband Raman imaging in biomedicine, they do so through applications of pulsed laser systems that are technically demanding to operate and incur significant costs.¹⁴

A potential alternative, or complement, to hardware-based solutions lies in deep learning.¹⁵ Deep learning is a subset of machine learning capable of uncovering effective representa-

tions of data across multiple levels of abstraction and has demonstrated incredible results across several domains, including image classification and segmentation, natural language processing, and predictive modeling.^{16–19} Recently, the application of deep learning to point-based Raman spectroscopy has achieved promising results in the intra-operative diagnosis of brain tumors, the rapid identification of pathogenic bacteria, and the production of subcellular organelle segmentation maps.^{20–22} While such applications are likely to improve the Raman spectroscopic diagnostic accuracy,²³ even greater benefits lie in the potential for deep learning to improve Raman spectroscopic imaging by increasing signal acquisition speeds and enabling high-quality reconstruction from noisy, low-resolution input data. Hyperspectral Raman images, where each pixel contains a complete Raman spectrum, represent highly structured data with complex spatial and spectral correlations amenable to deep learning. This highly structured nature of the data is at present underutilized, with existing data processing and analysis techniques (e.g., chemometrics and

Received: May 24, 2021

Accepted: October 8, 2021

Published: November 19, 2021



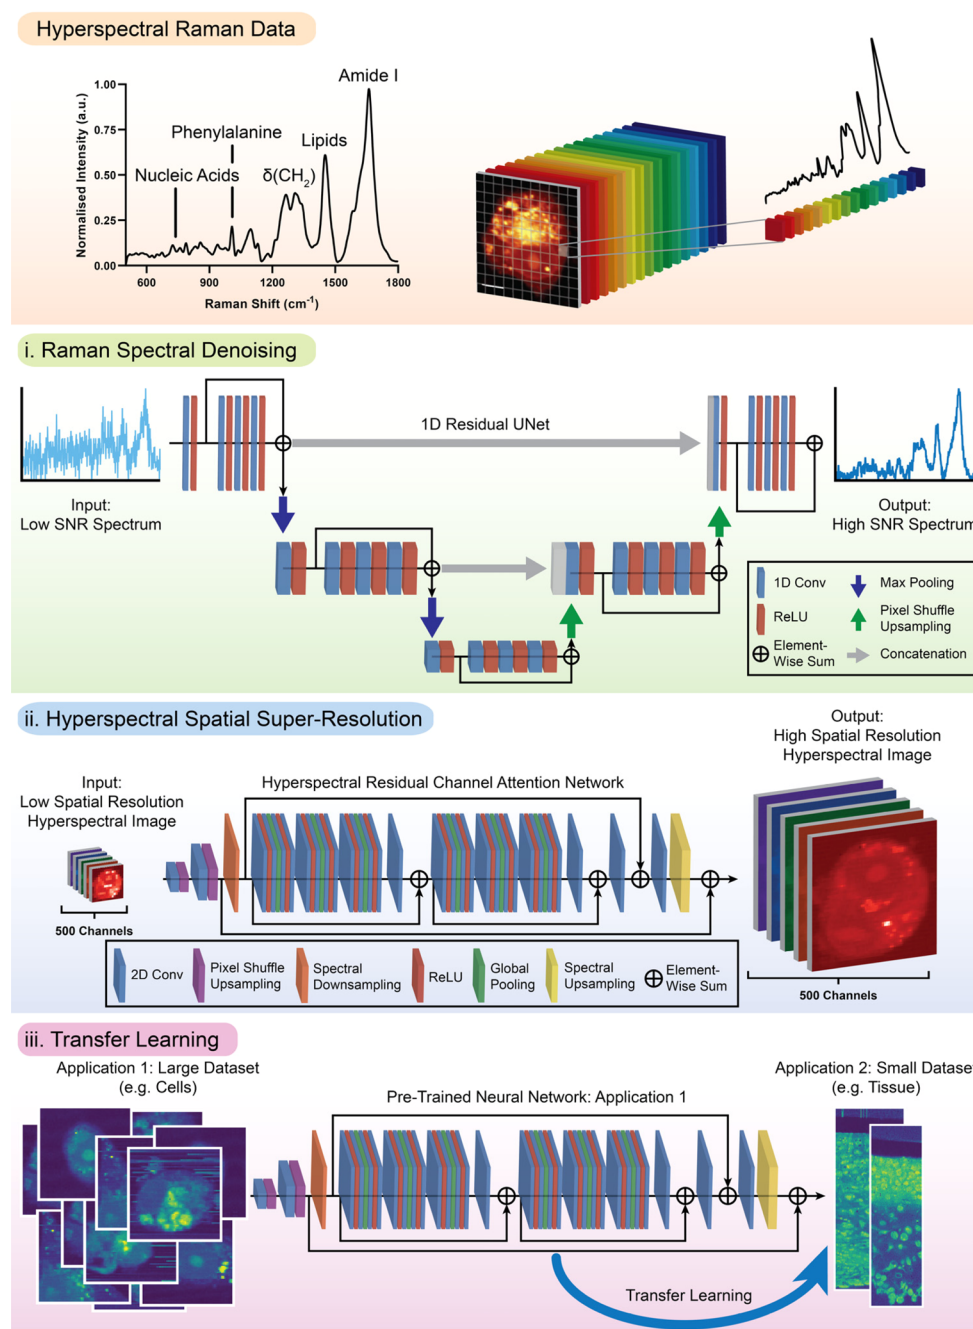


Figure 1. Deep-Learning-Enabled Raman Hyperspectral Super-Resolution Imaging. The deep-learning framework DeepeR is designed to operate on hyperspectral Raman images, where high information-content Raman spectra at each pixel provide detailed insight into the molecular composition of cells/tissues. To improve the speed of Raman spectroscopic imaging and enable high-throughput applications, we first (i) train a 1D ResUNet neural network for Raman spectral denoising to effectively reconstruct a high SNR Raman spectrum (long acquisition time) from a corresponding low SNR input spectrum (short acquisition time). Next, we (ii) train a hyperspectral residual channel attention neural network to accurately reconstruct high spatial resolution hyperspectral Raman images from corresponding low spatial resolution hyperspectral Raman images to significantly reduce imaging times. Then, by combining (i) and (ii), we achieve extreme speed-ups of up to 160 \times in Raman imaging time while maintaining high reconstruction fidelity. Finally, we (iii) demonstrate that transfer learning can be used to take our pretrained neural networks (trained on large datasets) to operate on an entirely unrelated hyperspectral data domain, for which there is only a limited data set (insufficient to effectively train a neural network from scratch).

multivariate analysis) failing to adequately exploit these complex correlations.¹⁵

Deep learning has made significant strides in signal reconstruction tasks, most notably for single-image super-resolution (SISR),^{24–26} image denoising,^{27,28} and signal denoising.^{29,30} In each of these domains, a neural network is trained on pairs of low-quality [e.g., low-resolution (LR)] and

high-quality [e.g. high-resolution (HR)] data, attempting to learn effective representations of low-quality inputs that reconstruct the corresponding high-quality outputs. Neural networks can thus be considered to learn prior information (e.g., shapes, sizes, and colors typical of different features) from the corpus of data in the training set in order to generate high-quality output data, given low-quality input data in the test set.

In the context of signal denoising, several groups have developed neural networks designed to reduce noise in electrocardiograms,^{29,30} while image denoising has been employed to improve image quality by removing noise generated by imaging hardware or compression artifacts.^{28,31} Similarly, much work has focused on SISR, with important potential life sciences applications already demonstrated for fluorescence microscopy, MRI, electron microscopy, and even endomicroscopy.^{32–36} SISR approaches have enabled HR fluorescence microscopy with $\sim 100\times$ lower light dose and $16\times$ higher frame rates for reduced photobleaching and phototoxicity.³⁴ Recently, SISR has been applied to line-scan Raman imaging, achieving a $5\times$ speed-up in line-scan imaging time.³⁷ This work applied a three-layer convolutional network to a narrow portion of the Raman spectrum, limiting potential wider applications.

Here, we present DeepeR, a comprehensive deep-learning framework for high-throughput molecular imaging via deep learning-enabled Raman spectroscopy. We first show, using a data set consisting of 172,312 pairs of low and high SNR Raman spectra of the entire fingerprint region, that deep learning significantly outperforms Savitzky–Golay (SG), wavelet, and principal component analysis (PCA) spectral smoothing algorithms $10\times$, enabling effective reconstruction of Raman signatures from low SNR Raman spectra. We next develop a convolutional neural network for hyperspectral Raman image super-resolution using an additional data set of 169 hyperspectral images representing 1.4 million Raman spectra (389 h of acquisition) in total. We achieve robust 2–4 \times spatial super-resolution image reconstruction, corresponding to 4–16 \times reduction in imaging time. Then, using a hybrid approach, we demonstrate Raman imaging with effective speed-ups of 40–160 \times while preserving molecular cellular information. Finally, we highlight the generalizability of our deep-learning framework, employing transfer learning to extend our pretrained neural networks from cells to tissues.

MATERIALS AND METHODS

Neural Network Architecture and Implementation: Raman Spectral Denoising. Denoising of Raman spectra was achieved via a one-dimensional (1D) ResUNet architecture (Figure 1). The network was trained for 500 epochs using the Adam optimizer,³⁸ with an L1-norm loss function and a one-cycle learning rate scheduler. Eleven independent models were trained, with the Raman spectra from a single hyperspectral Raman cell image used as the test set in each case, while the training and validation sets were formed from the Raman spectra of the remaining 10 hyperspectral Raman cell images. Evaluation on the validation set was used to prevent overfitting, while all results presented are the mean across the 11 test set folds. Full training details are provided in Supplementary Table 2. See Supporting Information for complete Python training scripts.

Neural Network Architecture and Implementation: Hyperspectral Super-Resolution. Hyperspectral Raman image spatial super-resolution was performed using a hyperspectral residual channel attention network (Supplementary Figure 1). The network was trained for 600 epochs using an Adam optimizer,³⁸ with an L1-norm loss function and a constant learning rate. Evaluation on the validation set was used to prevent overfitting, while all results presented are for the test set. Full training details are provided in Supplementary Table 3. See the Supporting Information for complete Python training scripts. Network performance was assessed using two common

image quality metrics, the peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM).

Data Augmentation. Data augmentation, essential for increasing the effective data set size, was performed using a custom PyTorch DataGenerator. Data augmentation included image subsampling, flipping, rotation, and mixup,³⁹ as well as spectral shifting, flipping, and background subtraction (Supplementary Figure 3).

Implementation. Complete implementation details are listed in the Supporting Information.

RESULTS AND DISCUSSION

Hyperspectral Raman Deep-Learning Framework.

DeepeR is designed to improve Raman spectroscopic acquisition times toward high-throughput Raman imaging applications. Working across hyperspectral Raman data, DeepeR performs (i) Raman spectral denoising, (ii) hyperspectral super-resolution, and (iii) transfer learning (Figure 1). Raman spectral denoising is performed using a 1D residual UNet (ResUNet),⁴⁰ which takes low SNR input spectra and reconstructs them to produce corresponding high SNR output spectra (Figure 1i). UNets have demonstrated excellent performance across a variety of applications in 1D and 2D, such as spectral artifact removal⁴¹ and image segmentation,⁴⁰ where the inputs and outputs have the same rank (shape). This is in part due to the UNet architecture, which enables it to learn from the data at multiple different feature scales. Thus, this makes UNets a suitable architectural choice for spectral denoising. Hyperspectral super-resolution is achieved using an adapted residual channel attention network (RCAN),⁴² a recent state-of-the-art neural network for spatial super-resolution of red-green-blue (RGB) images, to output an HR hyperspectral Raman image from an LR input (Figure 1ii, Supplementary Figure 1). The combination of (i) Raman spectral denoising and (ii) hyperspectral spatial super-resolution then enables significant Raman imaging speed-ups for high-throughput applications. Finally, DeepeR can be generalized to a wide range of Raman imaging applications through transfer learning, where neural networks pretrained on large hyperspectral datasets can be fine-tuned to operate effectively on small hyperspectral datasets (Figure 1iii).

Deep-Learning-Enabled Raman Denoising. We first developed a neural network training pipeline for Raman denoising—the reconstruction of a Raman spectrum with a high SNR from a corresponding low SNR Raman spectrum. We cultured MDA-MB-231 breast cancer cells, a widely studied cell line, and sequentially acquired low SNR (0.1 s integration time per spectrum) and high SNR (1 s integration time per spectrum) hyperspectral confocal Raman cell image pairs of varying size ($n = 11$ cells) using 532 nm laser excitation. This resulted in a large data set consisting of pairs of low and high SNR Raman spectra ($n = 172,312$ spectral pairs). Importantly, these Raman spectra contain an abundance of molecular information, including information about the relative concentrations and distributions of various nucleic acids, proteins, and lipids.¹ For instance, intense peaks can be seen near 795 cm^{-1} (DNA), 1004 cm^{-1} (phenylalanine), 1300 and 1440 cm^{-1} (lipids), and 1660 cm^{-1} (predominantly amide I of proteins). Successful denoising and reconstruction of low SNR Raman spectra require that this biochemical information be effectively preserved.

We applied a 1D ResUNet to this data set, performing 11-fold cross-validation by training 11 independent models on training/validation sets composed of the spectra from 10 hyperspectral Raman cell images, where the test sets in each case consisted of




Figure 2. Deep-Learning-Enabled Raman Denoising. (a) Exemplar test set pair of low SNR input Raman spectrum (light gray) and corresponding high SNR target Raman spectrum (dark gray) as well as SG (light blue), wavelet denoising (purple), PCA denoising (dark blue), and neural network (red) outputs for the given input spectrum (normalized to maximum peak intensity). (b) MSE (performed across all spectral channels and all image pixels) across all test set hyperspectral Raman cell images for raw input spectra, 1D ResUNet output spectra, PCA denoising output spectra, wavelet denoising output spectra, and SG output spectra (order x , frame width y) output spectra with respect to corresponding target spectra ($n = 11$) (error bars: mean \pm STD) (one-way ANOVA with Dunnett's multiple comparison test against raw input spectra, *** $P < 0.005$). (c) Exemplar 1450 cm^{-1} peak intensity heatmaps for low SNR input hyperspectral Raman image, PCA denoising of input hyperspectral Raman image, 1D ResUNet output, and target high SNR hyperspectral Raman image with corresponding imaging times shown in white (min:sec) (scale bar = $10\ \mu\text{m}$). (d) Exemplar vertex component analysis (VCA) performed on target high SNR hyperspectral Raman image identifies five key components (proteins/lipids [red], nucleic acids [blue], proteins [green], lipids [yellow], and background [black]), which are applied to low SNR input, PCA denoising output, and 1D ResUNet output images via non-negatively constrained least-squares regression, demonstrating that low SNR input and PCA denoising output data do not effectively identify different cell components. (e, f) Exemplar Raman spectra (white arrows in (c)) corresponding to (e) a lipid-rich cytoplasmic region and (f) the nucleus.

the spectra from the remaining hyperspectral Raman cell image (see [Materials and Methods](#) and Supplementary Table 1 for full implementation details). To increase the effective size of our data set and improve the robustness of the 1D ResUNet, we employed data augmentation including spectral flipping, spectral shifting, and background subtraction. Importantly, such augmentations are designed to maintain denoising

performance in the face of spectral changes (e.g., wavelength shifts) that occur across different Raman spectroscopy systems. The 1D ResUNet learned to produce high-quality output Raman spectra from low SNR input spectra that strongly aligned with the target (ground truth) high SNR spectra ([Figure 2](#)). To achieve this, the 1D ResUNet operates across multiple feature scales and spectral resolutions, learning spectral features (and




Figure 3. Deep-learning-enabled hyperspectral image super-resolution. (a) 2×, 3×, and 4× super-resolution of the example test set hyperspectral Raman image enables a significant reduction in imaging times (shown in white, min:sec) while recovering important spatial and spectral information (scale bars = 10 μm). Images shown are the result of a VCA performed on the target HR hyperspectral Raman image, which identified four key components (nucleic acids [blue], proteins [green], lipids [yellow], and background [black]). VCA components were applied to the nearest neighbor output, bicubic output, and HyRISR output images via non-negatively constrained least-squares regression. (b) Exemplar Raman spectrum with white arrow in (a) demonstrating that the neural network output (red) is more closely aligned to the target (ground truth) spectrum (dark gray). (c, d) Mean test set (c) PSNR, (d) SSIM, and (e) MSE values for nearest neighbor upsampling, bicubic upsampling, and HyRISR output for 2×, 3×, and 4× super-resolution ($n = 9$) (error bars: mean \pm STD) (One-way paired ANOVA with Geisser–Greenhouse correction and Tukey’s multiple comparisons test, * $P < 0.05$, ** $P < 0.01$, *** $P < 0.001$).

the molecular constituents they represent) typical of Raman spectra in the training set in order to identify a mapping between low SNR inputs and target high SNR outputs in the test set. Importantly, the neural network significantly outperformed PCA denoising, wavelet denoising, and SG filtering [assessed via a one-way analysis of variance (ANOVA) with Dunnett’s multiple comparison test against raw input spectra], three widely applied techniques for Raman spectral smoothing.^{43,44} Applications of the neural network to the test sets achieved a spectral mean-squared error (MSE) between the output and target spectra that were 10× lower than the next best-performing denoising technique (PCA denoising), with a mean MSE of 2.85×10^{-3} [95% confidence interval (CI): 2.55×10^{-3} , 3.15×10^{-3}] for the 1D ResUNet and 2.96×10^{-2} [95% CI: 2.55×10^{-2} , 3.38×10^{-2}] for PCA denoising (Figure 2b). This result demonstrates that the 1D ResUNet effectively learned the structure of Raman spectra, enabling it to discern true signals from the high level of background noise. Indeed, similar UNet architectures have previously been applied for artifact removal in

infrared spectroscopy and denoising of SRS microscopy images with impressive results.^{41,45} In contrast, PCA denoising, wavelet denoising, and SG filtering introduced spectral artifacts (likely due to the low SNR of the input data), which resulted in poor denoising performance even with the application of an additional asymmetric least-squares background subtraction and normalization steps post denoising.

We then examined whether the neural network-based Raman spectral reconstruction would result in significant loss of biochemical information or the introduction of “hallucinated” spatial or spectral features. To do this, we compared the quality of Raman hyperspectral images for each hyperspectral Raman cell image based on low SNR input data, PCA denoised low SNR data, neural network reconstructed data, and the ground truth high SNR data (Figure 2c). PCA denoising produced an image with enhanced contrast but amplified noise in different spatial and spectral regions. The 1D ResUNet produced a much sharper result with improved contrast of biomolecular features that closely align with the target high SNR image, suggesting the

preservation of biochemical information. Indeed, the 1D ResUNet significantly outperformed PCA denoising in terms of two commonly used image quality metrics, PSNR and SSIM. While the 1D ResUNet achieved a mean PSNR of 46.21 [95% CI: 45.76, 46.67] and a mean SSIM of 0.9532 [95% CI: 0.9154, 0.9910] across the 11 hyperspectral Raman cell images, PCA denoising resulted in the statistically significantly lower values (assessed via a two-tailed Wilcoxon paired signed rank test) of 39.36 [95% CI: 38.89, 39.83] and 0.8679 [95% CI: 0.8111, 0.9248], respectively.

To assess the preservation of biochemical information, we next applied a vertex component analysis (VCA) to an exemplar target high SNR hyperspectral Raman image (Figure 2d). VCA is a spectral unmixing technique designed for the unsupervised extraction of endmembers from hyperspectral data, enabling the identification of major constituent components in a hyperspectral Raman image (e.g. lipid-rich, nucleic-acid rich, background regions).⁴⁰ The VCA endmembers identified for the target high SNR hyperspectral Raman image (Supplementary Figure 2) were then applied to the input data, the PCA output data, and the 1D ResUNet output data via non-negatively constrained least-squares regression. Crucially, this analysis demonstrated that the 1D ResUNet output effectively identified and preserved key molecular species present in the hyperspectral image in line with the target high SNR hyperspectral image. In contrast, the PCA output failed to robustly distinguish the different Raman cellular signatures. Exemplar spectra from two different regions (nucleus and cytoplasm) of the hyperspectral images (Figure 2e,f) further demonstrated the superiority of the 1D ResUNet output for the accurate reconstruction of biochemical information contained in the Raman spectra. Importantly, applying the neural network on a per-hyperspectral pixel in this manner effectively enabled Raman spectroscopic imaging up to 10× faster than conventional Raman spectroscopy in this case, while preserving biochemical information.

Deep-Learning-Enabled Hyperspectral Image Spatial Super-Resolution. While the denoising results demonstrate a significant improvement in imaging times for conventional Raman spectroscopic imaging, the 1D ResUNet does not consider the high degree of molecular compositional correlation between adjacent pixels. We therefore sought to improve this and take spatial context into consideration by developing a 2D neural network for hyperspectral Raman image spatial super-resolution (HyRISR). To achieve this, we trained HyRISR to take a low spatial resolution (LR), high SNR hyperspectral image as input and output a corresponding high spatial resolution (HR), high SNR hyperspectral image. HyRISR learns to identify spatial and spectral correlations present in the training set in order to develop an accurate mapping between the LR inputs and target HR outputs in the test set. HyRISR follows a similar architecture to the RCAN, with the introduction of 2× spectral downsampling early in the network, followed by 2× spectral upsampling at the end of the network (Supplementary Figure 1). This use of spectral downsampling exploits the high spectral resolution (and hence high channel redundancy) of Raman spectra to reduce the computational load of HyRISR without sacrificing super-resolution performance.

We applied HyRISR to a data set of hyperspectral Raman images of MDA-MB-231 breast cancer cells ($n = 169$ Raman images) using a data split of 85:10:5 for training, validation, and test sets (see Materials and Methods and Supplementary Table 2 for full implementation details). To increase the effective size of our data set and improve the robustness of HyRISR, we

employed extensive data augmentation, including randomly applied image cropping, flipping, rotation, and mixup,³⁹ as well as randomly applied spectral flipping and shifting (Supplementary Figure 3). To generate LR input images, we applied 2×, 3×, or 4× spatial skip downsampling to corresponding HR images ($64 \times 64 \times 500$ or $63 \times 63 \times 500$, height \times width \times spectral channels) to reflect the raster scan nature of Raman imaging. Applications of HyRISR to the test set for 2×, 3×, and 4× super-resolution yielded superior performance, statistically significantly exceeding standard nearest neighbor and bicubic upsampling methods in terms of two image quality metrics, PSNR and SSIM, as well as in terms of MSE (Figure 3, Supplementary Figures 4–6, Supplementary Table 1). Importantly, these results, particularly the 2× and 3× super-resolution images, demonstrated good fidelity to the corresponding high-resolution target image, accurately reconstructing spectral features to correctly identify cellular components via VCA. This, combined with the reduced MSE values, demonstrates that the neural network can effectively preserve molecular information through accurate spectral reconstruction. Although at 2×, HyRISR produced minimal blurring, blur increased significantly at 3× and 4×, a well-known result of training our SISR neural network with an L1 loss function (see Results and Discussion for further details). Despite this blurring, the HyRISR output qualitatively better delineated cell boundaries, correctly identified subcellular features, and introduced fewer artifacts compared to bicubic upsampling. Importantly, these results were achieved with a training data set of just 144 Raman images (a considerably smaller data set than is typically employed for deep learning super-resolution of RGB images). The extension of this data set would likely yield significant improvements in super-resolution performance.

Notably, as with the denoising 1D ResUNet presented above, HyRISR enabled a significant reduction in imaging time, down from 68:15 (min:sec) to 17:03 in the case of 2× super-resolution and to 07:20 for 3× super-resolution. Importantly, this was achieved with only a limited loss of high-frequency details, with biochemical spectral information well maintained as evidenced by VCA (Figure 3a,b). In contrast, bicubic upsampling introduced numerous artifacts into the hyperspectral image. Although 4× super-resolution reduces imaging time further to 04:15, it does so with a much greater loss of fine details for both bicubic upsampling and HyRISR. While this might not be suitable for high-resolution cellular imaging, such 4× super-resolution Raman imaging could prove useful for other applications.

Hybrid Denoising and Super-Resolution Raman Spectroscopy for High-Throughput Molecular Imaging. While both our 1D ResUNet Raman spectral denoising and HyRISR neural networks enable significant speed-ups in Raman imaging time, single cell Raman imaging using either network remains on the order of minutes. Although this is considerably faster than conventional Raman imaging, it remains too slow for high-throughput applications such as cell imaging or automated spectral histopathology. To further improve the speed of Raman image acquisition, we next sequentially applied our two pretrained neural networks to perform Raman spectral denoising, followed by hyperspectral image super-resolution on a single hyperspectral Raman image. Sequential applications of the neural networks in this manner enabled the use of all data present in each data set (as opposed to the small subset of data present in both the denoising and HyRISR datasets that would enable training of a single network for end-to-end denoising and

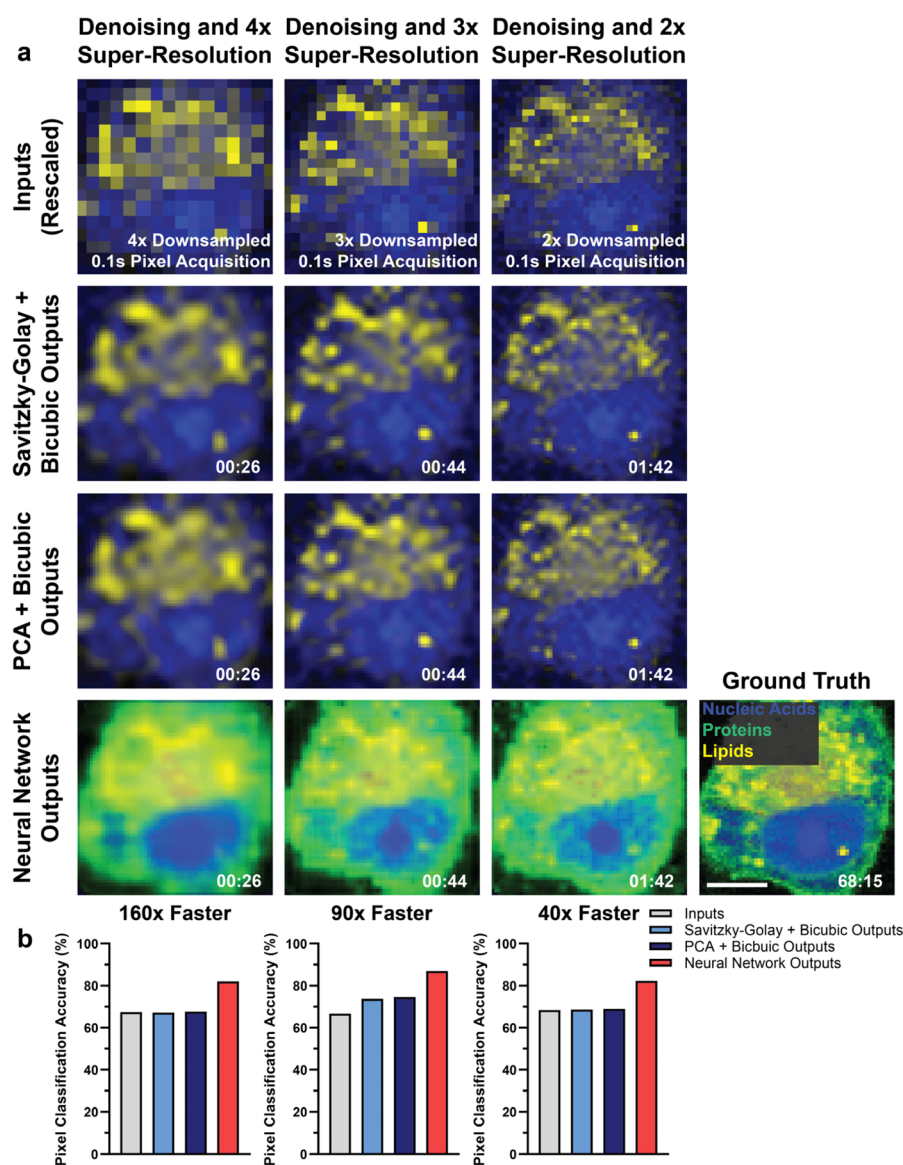


Figure 4. Combined Raman spectral denoising and hyperspectral image super-resolution enable extreme speed-ups in Raman imaging time. (a) Sequential application of Raman spectral denoising followed by hyperspectral image super-resolution enables extreme speed-ups in imaging time (shown in white) from 68:15 (min:sec) to 01:42 for 2× super-resolution, 00:44 for 3× super-resolution, and 00:26 for 4× super-resolution while largely preserving molecular information (scale bars = 10 μm). Images shown are the result of a VCA performed on the target HR, high SNR hyperspectral Raman image, which identified four key components (nucleic acids [blue], proteins [green], lipids [yellow], and background [black]). VCA components were applied to input, Savitzky-Golay plus bicubic upsampling, PCA plus bicubic upsampling, and neural network output images via non-negatively constrained least-squares regression. (b) Pixel classification accuracy for input, Savitzky-Golay filtering plus bicubic upsampling output, PCA denoising plus bicubic upsampling output, and neural network output images as compared to VCA pixel classification of target HR, high SNR hyperspectral Raman image.

super-resolution). Here, using the Raman spectra from a single cell (present in the test sets for both datasets), we achieved effective speed-ups of 40× (2× super-resolution), 90× (3× super-resolution), and 160× (4× super-resolution) while accurately reconstructing a high SNR, HR hyperspectral Raman image from a low SNR, and LR input hyperspectral Raman image (Figure 4, Supplementary Figure 7).

We again used VCA to identify key Raman spectral components in our ground truth image and employed non-negatively constrained least-squares regression to apply the identified VCA endmembers to the input images (nearest neighbor rescaled), Savitzky-Golay filtering plus bicubic upsampling output images, PCA denoising plus bicubic upsampling images, and neural network output images. Our

neural networks outperformed the combination of SG filtering and bicubic upsampling, accurately reconstructing both spatial and spectral information and maintaining robust VCA endmember identification (Figure 4a). In each case, this resulted in an improved pixel classification accuracy (as compared to pixel classification for the ground truth hyperspectral Raman image, determined as the VCA endmember with the maximum intensity value for each pixel as per non-negatively constrained least-squares regression) relative to the inputs and SG and PCA outputs (Figure 4b). Crucially, accurate spectral and spatial reconstruction was maintained even for 40× and 90× Raman imaging time speed-ups, enabling HR hyperspectral Raman cell imaging in under 2 min or under 1 min, respectively. While imaging time can be further reduced by employing 4×

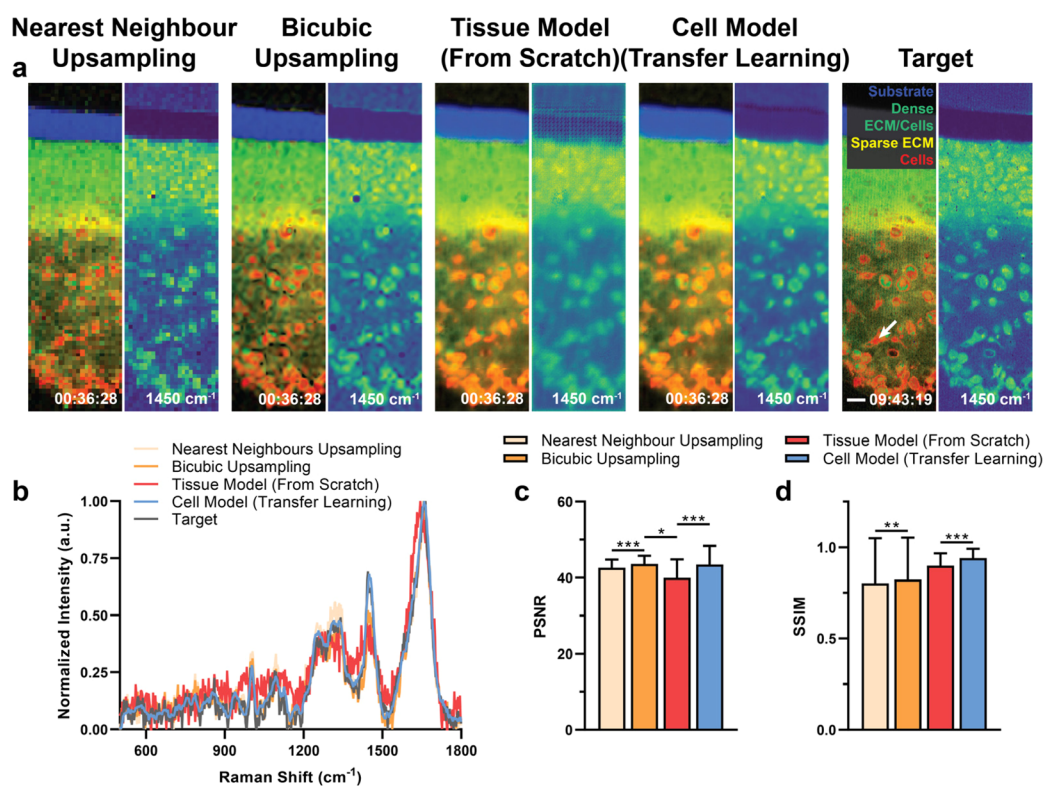


Figure 5. Transfer learning enables effective super-resolution for a small data set of tissue-engineered cartilage hyperspectral Raman images. (a) Transfer learning of our HISR neural network, trained only on MDA-MB-231 breast cancer cell images, enabled effective cross-domain 4 \times super-resolution of hyperspectral Raman images despite having only a very small data set of tissue-engineered cartilage for training. For each condition, images shown on the left are the result of VCA performed on the target HR and high SNR hyperspectral Raman image, which identified five key components (substrate [blue], dense ECM/cells [green], sparse ECM [yellow], cells [red], and background [black]). VCA components were applied to the nearest neighbor upsampling, bicubic upsampling, tissue model (from scratch), and cell model (transfer learning) images via non-negatively constrained least-squares regression. Images shown on the right for each condition are 1450 cm^{-1} peak intensity heatmaps. All images formed as the composition of overlapping 64×64 pixel image patches (scale bars = $10 \mu\text{m}$). (b) Exemplar Raman spectrum (white arrow in (a)) demonstrating that transfer learning achieves high accuracy reconstruction of the target spectra for each pixel. (c, d) Mean test set (c) PSNR and (d) SSIM values for the nearest neighbor upsampling, bicubic upsampling, and neural network outputs for 4 \times super-resolution, calculated on a per-image patch basis ($n = 12$ patches) (Error bars: mean \pm STD) (One-way paired analysis of variance ANOVA with Geisser–Greenhouse correction and Tukey’s multiple comparisons test, * $P < 0.05$, ** $P < 0.01$, *** $P < 0.001$).

super-resolution for a 160 \times Raman imaging time speed-up, reconstructed image quality continues to degrade and may produce undesirable artifacts at higher super-resolution scales. Despite this, image quality following such 160 \times speed-up is likely to be useful across many Raman imaging applications such as for the rapid screening of large areas containing multiple cells or for spectral pathology applications.

Generalized Hyperspectral Imaging from Cells to Tissues Using Transfer Learning. Finally, to highlight the generalizability and wide applicability of our DeepeR framework, we used transfer learning to apply HyRISR to a small data set of hyperspectral Raman images of unrelated origin. We thus extended DeepeR to the field of regenerative medicine for super-resolution hyperspectral Raman imaging of in vitro-formed cartilage constructs. Both the spatial and spectral information contained in the hyperspectral Raman images of tissue-engineered cartilage samples differ significantly from those of MDA-MB-231 breast cancer cells used to train HyRISR, representing an effective test of the transferability of DeepeR. Deep learning is a data-heavy approach that requires large, labeled datasets in order to be effective. For applications where such a large data set does not exist, data acquisition for deep learning can be prohibitively time-consuming and expensive. Here, we aimed at demonstrating that transfer learning, the

application of an existing neural network model trained on a large data set to a second, smaller data set, can achieve high-quality results (Figure 1 iii).

To do this, we used a small training data set consisting of 16 patches ($64 \times 64 \times 500$ each) from large HR Raman hyperspectral images of tissue-engineered cartilage. A separate test set of 12 overlapping patches ($64 \times 64 \times 500$ each) was extracted from a separate, single large Raman hyperspectral image ($100 \times 350 \times 500$) of a tissue-engineered cartilage sample. Here, transfer learning was performed by further training all neural network weights of a pretrained HyRISR model for 200 epochs on the small tissue-engineered cartilage training data set with a reduced learning rate. We then compared the super-resolution performance of this fine-tuned model, against both the nearest neighbor and bicubic upsampling as well as against HyRISR trained from scratch on the small training data set of hyperspectral Raman tissue-engineered cartilage image patches alone (Figure 5, Supplementary Figure 8). As expected, transfer learning of HyRISR achieved superior results to the nearest neighbor upsampling, bicubic upsampling, and HyRISR trained from scratch on the tissue-engineered cartilage data set alone in terms of PSNR and SSIM (Figure 5c,d). As with the super-resolution of hyperspectral Raman images of MDA-MB-231 breast cancer cells, the fine-tuned

neural network here produced a highly accurate reconstruction with a few introduced artifacts and a degree of over-smoothing. Meanwhile, bicubic upsampling resulted in an image that appears grossly similar to the target ground truth image, but suffered from the introduction of numerous artifacts, resulting in both spatial and spectral distortion (Figure 5b–d). However, as in the case of our super-resolution results, performance here is limited by the size of both the data set used for initial network training, as well as the size of the data set used for transfer learning.

DeepeR is a comprehensive deep-learning framework that offers a completely new approach to high-throughput Raman spectroscopic imaging. DeepeR can be applied online or offline to existing Raman spectroscopic systems without requiring any hardware modifications or imposing system limitations using transfer learning. Offline applications to existing hyperspectral Raman datasets could be used to develop custom models for specific applications, either from scratch or by transfer learning from our pretrained networks. Online application, with inference occurring in a matter of seconds for a GPU-equipped scientific computer, will deliver high-throughput imaging capabilities, transforming the potential range of applications for hyperspectral Raman imaging. DeepeR will thus help drive forward high-throughput hyperspectral Raman imaging, representing a major departure from existing chemometric and other multivariate statistical techniques.

Despite its significant advances, DeepeR does face a number of limitations that must be considered before application to additional hyperspectral Raman datasets. Most notably, our framework is unlikely to accurately reconstruct very fine (e.g., single pixel) details and so may not be suitable for HR imaging of small, complex specimens. Second, while transfer learning using our pretrained models will enable a much wider range of applications, hyperspectral Raman images with substantially different spatial and spectral features (e.g., different cells, tissues, etc.) will still require a sufficiently large data set for effective performance. Future work will seek to expand the depth and breadth of our hyperspectral Raman data set, encompassing spectra from a variety of instruments and samples. Lastly, before widespread applications to Raman spectroscopic imaging is possible, large-scale prospective validation will need to be performed, specific to each application, to ensure that diagnostic or scientific decisions match those made for corresponding HR, high SNR hyperspectral Raman data.

There remains scope for improvements in the performance of our deep-learning framework, most notably through the collection of larger training datasets from different biomedical applications and the development of more advanced neural network architectures.²³ The collection of a large data set of paired low SNR, LR and high SNR, HR hyperspectral Raman images would enable the training of a joint denoising and super-resolution neural network, which we anticipate would produce improved performance in line with existing studies on multitask neural networks.⁴⁷ Performance could potentially be further improved by implementing a generative adversarial network (GAN) architecture.⁴⁸ GANs have demonstrated an array of impressive results for the super-resolution of RGB images and medical images such as endomicroscopy images.^{26,49,50} However, GAN architectures pose particularly significant demands on computational resources in the context of hyperspectral image super-resolution.

Although here we demonstrated the application of DeepeR to Raman spectroscopy, equivalent neural network architectures

could be generalized to alternative techniques such as FT-IR imaging,⁵¹ hyperspectral imaging,⁵² or mass spectrometry imaging techniques.⁵³

In conclusion, DeepeR represents a comprehensive deep learning framework for high-throughput hyperspectral Raman imaging. This has the potential to transform the application of Raman spectroscopic imaging in the biomedical sciences, enabling a host of higher-throughput applications not previously possible. Crucially, the information and data we provide open source to the community, including our complete data set, pretrained models, and Python code, will enable rapid expansion and integration of our framework into existing Raman spectroscopy systems, driving forward high-throughput Raman imaging.

■ ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.analchem.1c02178>.

Datasets will be made available upon publication. All code and pretrained models are available online at <https://github.com/conor-horgan/>.

Additional materials and methods pertaining to the preparation of cell and tissue culture samples, Raman spectroscopic imaging setup and processing, details of the calculations of image quality metrics, neural network implementation details, and supplementary figures. (PDF)

■ AUTHOR INFORMATION

Corresponding Authors

Molly M. Stevens – Department of Materials, Department of Bioengineering, and Institute of Biomedical Engineering, Imperial College London, London SW7 2AZ, U.K.; orcid.org/0000-0002-7335-266X; Email: m.stevens@imperial.ac.uk

Mads S. Bergholt – Centre for Craniofacial and Regenerative Biology, King's College London, London SE1 9RT, U.K.; orcid.org/0000-0003-3986-8942; Email: mads.bergholt@kcl.ac.uk

Authors

Conor C. Horgan – Centre for Craniofacial and Regenerative Biology, King's College London, London SE1 9RT, U.K.; Department of Materials, Department of Bioengineering, and Institute of Biomedical Engineering, Imperial College London, London SW7 2AZ, U.K.

Magnus Jensen – Centre for Craniofacial and Regenerative Biology, King's College London, London SE1 9RT, U.K.

Anika Nagelkerke – Groningen Research Institute of Pharmacy, Pharmaceutical Analysis, University of Groningen, Groningen 9700 AD, The Netherlands; Department of Materials, Department of Bioengineering, and Institute of Biomedical Engineering, Imperial College London, London SW7 2AZ, U.K.

Jean-Philippe St-Pierre – Department of Chemical and Biological Engineering, University of Ottawa, Ottawa, Ontario K1N 6N5, Canada; Department of Materials, Department of Bioengineering, and Institute of Biomedical Engineering, Imperial College London, London SW7 2AZ, U.K.

Tom Vercauteren – School of Biomedical Engineering and Imaging Sciences, King's College London, London WC2R 2LS, U.K.

Complete contact information is available at:

<https://pubs.acs.org/10.1021/acs.analchem.1c02178>

Notes

The authors declare the following competing financial interest(s): Tom Vercauteren is founding director and shareholder of Hypervision Surgical Ltd and holds shares from Mauna Kea Technologies.

ACKNOWLEDGMENTS

This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 802778). This work is supported by the Wellcome/EPSRC Centre for Medical Engineering [WT 203148/Z/16/Z]. C.C.H. acknowledges funding from the NanoMed Marie Skłodowska-Curie ITN from the H2020 programme under grant number 676137. T.V. is supported by a Medtronic/Royal Academy of Engineering Research Chair [RCSR1819/7/34]. A.N. and M.M.S. acknowledge support from the GlaxoSmithKline Engineered Medicines Laboratory. M.M.S. acknowledges a Wellcome Trust Senior Investigator Award (098411/Z/12/Z). We acknowledge the use of the JADE HPC facility, which has received funding through the Engineering and Physical Sciences Research Council (EPSRC).

REFERENCES

- (1) Horgan, C. C.; Nagelkerke, A.; Whittaker, T. E.; Nele, V.; Massi, L.; Kauscher, U.; Penders, J.; Bergholt, M. S.; Hood, S. R.; Stevens, M. M. *J. Mater. Chem. B* **2020**, *8*, 4447–4459.
- (2) Ichimura, T.; Chiu, L. D.; Fujita, K.; Kawata, S.; Watanabe, T. M.; Yanagida, T.; Fujita, H. *PLoS One* **2014**, *9*, No. e84478.
- (3) Kong, K.; Rowlands, C. J.; Varma, S.; Perkins, W.; Leach, I. H.; Koloydenko, A. A.; Williams, H. C.; Notingher, I. *Proc. Natl. Acad. Sci. U. S. A.* **2013**, *110*, 15189–15194.
- (4) Takamori, S.; Kong, K.; Varma, S.; Leach, I.; Williams, H. C.; Notingher, I. *Biomed. Opt. Express* **2015**, *6*, No. 98.
- (5) Bergholt, M. S.; Zheng, W.; Ho, K. Y.; Teh, M.; Yeoh, K. G.; Yan So, J. B.; Shabbir, A.; Huang, Z. *Gastroenterology* **2014**, *146*, 27–32.
- (6) Bergholt, M. S.; Lin, K.; Wang, J.; Zheng, W.; Xu, H.; Huang, Q.; Ren, J. L.; Ho, K. Y.; Teh, M.; Srivastava, S.; Wong, B.; Yeoh, K. G.; Huang, Z. *J. Biophotonics* **2016**, *9*, 333–342.
- (7) Smith, G. P. S.; McGovern, C. M.; Fraser, S. J.; Gordon, K. C. *Adv. Drug Deliv. Rev.* **2015**, *89*, 21–41.
- (8) Krafft, C.; Schmitt, M.; Schie, I. W.; Cialla-May, D.; Matthäus, C.; Bocklitz, T.; Popp, J. *Angew. Chem., Int. Ed.* **2017**, *56*, 4392–4430.
- (9) Kong, K.; Zaabar, F.; Rakha, E.; Ellis, I.; Koloydenko, A.; Notingher, I. *Phys. Med. Biol.* **2014**, *59*, 6141–6152.
- (10) Rowlands, C. J.; Varma, S.; Perkins, W.; Leach, I.; Williams, H.; Notingher, I. *J. Biophotonics* **2012**, *5*, 220–229.
- (11) Evans, C. L.; Potma, E. O.; Puoris'haag, M.; Cote, D.; Lin, C. P.; Xie, X. S. *Proc. Natl. Acad. Sci. U. S. A.* **2005**, *102*, 16807–16812.
- (12) Saar, B. G.; Freudiger, C. W.; Reichman, J.; Stanley, C. M.; Holtom, G. R.; Xie, X. S. *Science* **2010**, *330*, 1368–1370.
- (13) Orringer, D. A.; Pandian, B.; Niknafs, Y. S.; Hollon, T. C.; Boyle, J.; Lewis, S.; Garrard, M.; Hervey-Jumper, S. L.; Garton, H. J. L.; Maher, C. O.; Heth, J. A.; Sagher, O.; Wilkinson, D. A.; Snuderl, M.; Venneti, S.; Ramkissoon, S. H.; McFadden, K. A.; Fisher-Hubbard, A.; Lieberman, A. P.; Johnson, T. D.; Xie, X. S.; Trautman, J. K.; Freudiger, C. W.; Camelo-Piragua, S. *Nat. Biomed. Eng.* **2017**, *1*, No. 0027.
- (14) Krafft, C.; Popp, J. *Anal. Bioanal. Chem.* **2015**, *407*, 699–717.
- (15) LeCun, Y.; Bengio, Y.; Hinton, G. *Nat. Methods* **2015**, *521*, 436–444.
- (16) Moen, E.; Bannon, D.; Kudo, T.; Graf, W.; Covert, M.; van Valen, D. *Nat. Methods* **2019**, *16*, 1233–1246.
- (17) Krizhevsky, A.; Sutskever, I.; Hinton, G. E. *ImageNet classification with deep convolutional neural networks*. *Adv. Neural Inf. Process. Syst.* **2012**.
- (18) Devlin, J.; Chang, M.-W.; Lee, K.; Toutanova, K. *BERT: pre-training of deep bidirectional transformers for language understanding*. arXiv Prepr. (2018).
- (19) Senior, A. W.; Evans, R.; Jumper, J.; Kirkpatrick, J.; Sifre, L.; Green, T.; Qin, C.; Židek, A.; Nelson, A. W. R.; Bridgland, A.; Penedones, H.; Petersen, S.; Simonyan, K.; Crossan, S.; Kohli, P.; Jones, D. T.; Silver, D.; Kavukcuoglu, K.; Hassabis, D. *Nature* **2020**, *577*, 706–710.
- (20) Hollon, T. C.; Pandian, B.; Adapa, A. R.; Urias, E.; Save, A. V.; Khalsa, S. S. S.; Eichberg, D. G.; D'Amico, R. S.; Farooq, Z. U.; Lewis, S.; Petridis, P. D.; Marie, T.; Shah, A. H.; Garton, H. J. L.; Maher, C. O.; Heth, J. A.; McKean, E. L.; Sullivan, S. E.; Hervey-Jumper, S. L.; Patil, P. G.; Thompson, B. G.; Sagher, O.; McKhann, G. M., II; Komotar, R. J.; Ivan, M. E.; Snuderl, M.; Otten, M. L.; Johnson, T. D.; Sisti, M. B.; Bruce, J. N.; Muraszko, K. M.; Trautman, J.; Freudiger, C. W.; Canoll, P.; Lee, H.; Camelo-Piragua, S.; Orringer, D. A. *Nat. Med.* **2020**, *26*, 52–58.
- (21) Ho, C.-S.; Jean, N.; Hogan, C. A.; Blackmon, L.; Jeffrey, S. S.; Holodniy, M.; Banaei, N.; Saleh, A. A. E.; Ermon, S.; Dionne, J. *Nat. Commun.* **2019**, *10*, 4927.
- (22) Zhang, J.; Zhao, J.; Lin, H.; Tan, Y.; Cheng, J. X. *J. Phys. Chem. Lett.* **2020**, *11*, 8573–8578.
- (23) Pradhan, P.; Guo, S.; Ryabchykov, O.; Popp, J.; Bocklitz, T. W. *J. Biophotonics* **2020**, *13*, 1–24.
- (24) Dong, C.; Change Loy, C.; He, K.; Tang, X. Learning a deep convolutional network for image super-resolution. in *Proceedings of the European Conference on Computer Vision* 184–199 (2014).
- (25) Kim, J.; Kwon Lee, J.; Mu Lee, K. Accurate image super-resolution using very deep convolutional networks. in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*; Springer International Publishing: Cham, 1646–1654, (2016).
- (26) Wang, X. et al. ESRGAN: Enhanced super-resolution generative adversarial networks. in *Proceedings of the European Conference on Computer Vision* (2018).
- (27) Yuan, Q.; Zhang, Q.; Li, J.; Shen, H.; Zhang, L. *IEEE Trans. Geosci. Remote Sens.* **2019**, *57*, 1205–1218.
- (28) Zhang, K.; Zuo, W.; Chen, Y.; Meng, D.; Zhang, L. *IEEE Trans. Image Process.* **2017**, *26*, 3142–3155.
- (29) Rahhal, M. M. A.; Bazi, Y.; AlHichri, H.; Alajlan, N.; Melgani, F.; Yager, R. R. *Inform. Sci.* **2016**, *345*, 340–354.
- (30) Chiang, H. T.; Hsieh, Y. Y.; Fu, S. W.; Hung, K. H.; Tsao, Y.; Chien, S. Y. *IEEE Access* **2019**, *7*, 60806–60813.
- (31) Liu, D.; Wen, B.; Liu, X.; Wang, Z.; Huang, T. S. *When image denoising meets high-level vision tasks: A deep learning approach*. arXiv Prepr. (2018).
- (32) Wang, H.; Rivenson, Y.; Jin, Y.; Wei, Z.; Gao, R.; Günaydin, H.; Bentolila, L. A.; Kural, C.; Ozcan, A. *Nat. Methods* **2019**, *16*, 103–110.
- (33) Chen, Y.; Xie, Y.; Zhou, Z.; Shi, F.; Christodoulou, A. G.; Li, D. *2018 IEEE 15th Int. Sympos. Biomed. Imaging* **2018**, 739–742.
- (34) Fang, L.; Monroe, F.; Novak, S. W.; Kirk, L.; Schiavon, C. R.; Yu, S. B.; Zhang, T.; Wu, M.; Kastner, K.; Latif, A. A.; Lin, Z.; Shaw, A.; Kubota, Y.; Mendenhall, J.; Zhang, Z.; Pekkurnaz, G.; Harris, K.; Howard, J.; Manor, U. *Nat. Methods* **2021**, *18*, 406–416.
- (35) Weigert, M.; Schmidt, U.; Boothe, T.; Müller, A.; Dibrov, A.; Jain, A.; Wilhelm, B.; Schmidt, D.; Broaddus, C.; Culley, S.; Rochamartins, M.; Segovia-Miranda, F.; Norden, C.; Henriques, R.; Zerial, M.; Solimena, M.; Rink, J.; Tomancak, P.; Royer, L.; Jug, F.; Myers, E. W. *Nat. Methods* **2018**, *15*, 1090–1097.
- (36) Ravi, D.; Szczotka, A. B.; Shakir, D. I.; Pereira, S. P.; Vercauteren, T. *Int. J. Comput. Assist. Radiol. Surg.* **2018**, *13*, 917–924.
- (37) He, H.; Xu, M.; Zong, C.; Zheng, P.; Luo, L.; Wang, L.; Ren, B. *Anal. Chem.* **2019**, *91*, 7070–7077.

- (38) Kingma, D. P.; Ba, J. L. Adam: A method for stochastic optimization. in *International Conference on Learning Representations*, (2015), 5.
- (39) Zhang, H.; Cisse, M.; Dauphin, Y. N.; Lopez-Paz, D. *mixup: beyond empirical risk minimization*; ICLR, (2018).
- (40) Ronneberger, O.; Fischer, P.; Brox, T. U-net: Convolutional networks for biomedical image segmentation. in *International Conference on Medical Image Computing and Computer-Aided Intervention*, (2015), 9351, 234–241.
- (41) Guo, S.; Mayerhöfer, T.; Pahlow, S.; Hübner, U.; Popp, J.; Bocklitz, T. *Analyst* **2020**, *145*, 5213–5220.
- (42) Zhang, Y. et al. Image super-resolution using very deep residual channel attention networks. in *Proceedings of the European Conference on Computer Vision*, (2018), 286–301.
- (43) Byrne, H. J.; Knief, P.; Keating, M. E.; Bonnier, F. *Chem. Soc. Rev.* **2016**, *45*, 1865–1878.
- (44) Butler, H. J.; Ashton, L.; Bird, B.; Cinque, G.; Curtis, K.; Dorney, J.; Esmonde-White, K.; Fullwood, N. J.; Gardner, B.; Martin-Hirsch, P. L.; Walsh, M. J.; McAinsh, M. R.; Stone, N.; Martin, F. L. *Nat. Protoc.* **2016**, *11*, 664–687.
- (45) Manifold, B.; Thomas, E.; Francis, A. T.; Hill, A. H.; Fu, D. *Biomed. Opt. Express* **2019**, *10*, No. 3860.
- (46) Nascimento, J. M. P.; Dias, J. M. B. *IEEE Trans. Geosci. Remote Sens.* **2005**, *43*, 898–910.
- (47) Zhou, R.; El Helou, M.; Sage, D.; Laroche, T.; Seitz, A.; Süsstrunk, S. *W2S: a joint denoising and super-resolution dataset*. arXiv Prepr, (2020).
- (48) Goodfellow, I. J.; et al. *Adv. Neural Inf. Process. Syst.* **2014**, 2672–2680.
- (49) Ledig, C.; Theis, L.; Huszar, F.; Caballero, J.; Cunningham, A.; Acosta, A.; Aitken, A.; Tejani, A.; Totz, J.; Wang, Z.; Shi, W. Photo-realistic single image super-resolution using a generative adversarial network. in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2017).
- (50) Ravi, D.; Szczotka, A. B.; Pereira, S. P.; Vercauteren, T. *Med. Image Anal.* **2019**, *53*, 123–131.
- (51) Berisha, S.; Lotfollahi, M.; Jahanipour, J.; Gurcan, I.; Walsh, M.; Bhargava, R.; van Nguyen, H.; Mayerich, D. *Analyst* **2019**, *144*, 1642–1653.
- (52) Halicek, M.; Dormer, J. D.; Little, J. V.; Chen, A. Y.; Myers, L.; Sumer, B. D.; Fei, B. *Cancers* **2019**, *11*, No. 1367.
- (53) Behrmann, J.; Etmann, C.; Boskamp, T.; Casadonte, R.; Kriegsmann, J.; Maaß, P. *Bioinformatics* **2018**, *34*, 1215–1223.