Step-Calibrated Diffusion for Biomedical Optical Image Restoration

Yiwei Lyu^{*1}, Sung Jik Cha^{*2}, and Cheng Jiang¹, Asadur Chowdury¹, Xinhai Hou¹, Edward Harake¹, Akhil Kondepudi¹, Christian Freudiger³, Honglak Lee^{1,4}, Todd Hollon¹

¹ University of Michigan, Ann Arbor, MI, 48109

² Western Michigan University, Kalamazoo, MI, 49008

³ Invenio Imaging, Santa Clara, CA, 95051

⁴ LG AI Research, Ann Arbor, MI, 48106

Abstract. High-quality, high-resolution medical imaging is essential for clinical care. Raman-based biomedical optical imaging uses non-ionizing infrared radiation to evaluate human tissues in real time and is used for early cancer detection, brain tumor diagnosis, and intraoperative tissue analysis. Unfortunately, optical imaging is vulnerable to image degradation due to laser scattering and absorption, which can result in diagnostic errors and misguided treatment. Restoration of optical images is a challenging computer vision task because the sources of image degradation are multi-factorial, stochastic, and tissue-dependent, preventing a straightforward method to obtain paired low-quality/high-quality data. Here, we present Restorative Step-Calibrated Diffusion (RSCD), an unpaired image restoration method that views the image restoration problem as completing the finishing steps of a diffusion-based image generation task. RSCD uses a step calibrator model to dynamically determine the severity of image degradation and the number of steps required to complete the reverse diffusion process for image restoration. RSCD outperforms other widely used unpaired image restoration methods on both image quality and perceptual evaluation metrics for restoring optical images. Medical imaging experts consistently prefer images restored using RSCD in blinded comparison experiments and report minimal to no hallucinations. Finally, we show that RSCD improves performance on downstream clinical imaging tasks, including automated brain tumor diagnosis and deep tissue imaging. Our code is available at https://github.com/MLNeurosurg/truncated_diffusion.

Keywords: Unpaired Image Restoration · Biomedical Optical Imaging · Microscopy · Diffusion

1 Introduction

Medical imaging plays a major role in clinical medicine. Computed tomography, radiography, magnetic resonance imaging, and optical imaging are examples of common and indispensable medical imaging modalities used for diagnosis, guiding treatment decisions, and monitoring treatment response. Ramanbased biomedical optical imaging uses Raman scattering to non-invasively evaluate human tissues for diagnostic purposes. As an advanced medical imaging modality, it has an increasing role in patient care and is now being used for non-invasive cancer detection [29,41], brain tumor diagnosis [17,20], and surgical specimen analysis [16,31,34]. Raman-based optical imaging has several advantages over conventional medical imaging as it does not require ionizing radiation, does not cause tissue damage, and can acquire images rapidly (seconds) at the patient's bedside. Optical imaging uses light in the infrared electromagnetic spectrum to visualize biological tissues. Unfortunately, imaging within this spectral region causes optical imaging to be vulnerable to image degradation. Moreover, the sources of image degradation are multifaceted [30,41] (see Figure 1):

Tissue scattering: scattering of the incident light can result in noisy images, especially when imaging at depth.

^{*}Equal Contribution



Fig. 1: Examples of biomedical optical images restored using our proposed method, **RSCD**. Across a range of known and unknown sources of image degradation, RSCD provides high-quality image restoration of fresh, surgical specimens imaged during brain tumor surgery. Our method can restore optical images with severe image degradation such that, after restoration, they can be used for downstream clinical tasks, including automated brain tumor diagnosis and deep tissue imaging during surgery. (Near-registered images are approximate ground truth high-quality images, see Section 4.3)

- *Tissue absorption*: biologically tissues contain chromophores that absorb light at specific wavelengths, reducing signal strength.
- Auto-fluorescence: some biological tissues have intrinsic fluorescence that degrades optical images by reducing signal-to-noise ratios.
- Instrument noise: image noise can be introduced from the detectors, electronic components, and external sources of interference.

The frequency and degree of image degradation are inherently unpredictable, making widespread clinical integration of Raman-based optical imaging challenging [18, 41]. Moreover, while small paired low-quality/high-quality image data has been generated in a controlled laboratory setting [30, 42], it is impossible to collect a large scale paired dataset of perfectly aligned low-quality/high-quality clinical optical imaging that includes the full range of possible image degradations.

The majority of state-of-the-art image restoration methods require either paired image data or prior knowledge about the corruption process [25, 27, 42]. Unpaired restoration is a more challenging computer vision task and includes a diverse set of proposed methods. CycleGAN is a popular generative method for unpaired image-to-image translation [47]. Noise-to-noise methods involve mapping low-quality images to synthetically degraded paired images [26]. Deep image prior demonstrates that deep neural network architectures are sufficient to capture low-level image statistics prior to any model training [38]. Unfortunately, previous methods applied to biomedical imaging can suffer from hallucinations and perceptually poor reconstructions [1]. Importantly, poor reconstructions and hallucinations can have detrimental downstream effects in medical imaging, increasing the risk of nondiagnostic images or diagnostic errors. An ideal method for Raman-based optical image restoration would (1) restore images degraded from a wide range of

3



Fig. 2: An overview of *Restorative Step-Calibrated Diffusion (RSCD)*. We view the low-quality image as the output of an incomplete diffusion generation process that starts from Gaussian noise (t = T) and performs T steps of denoising (reverse diffusion) to generate a restored image at t = 0. We use a step calibrator model to predict t_{pred} , the number of steps of diffusion model denoising needed for image restoration, and we perform the reverse diffusion starting from t_{pred} . In addition, we use dynamic recalibration to dynamically adjust the number of steps required for optimal image restoration, t'_{pred} . The dynamic recalibration process and subsequent d steps of reverse diffusion denoising can be repeated until the restoration process runs to completion, obtaining the restored image at t = 0.

corruption sources, (2) only require unpaired data, (3) avoid hallucinations or perceptual artifacts, and (4) be time-efficient to allow for real-time, intraoperative image restoration.

We present **Restorative Step-Calibrated Diffusion (RSCD)**, a novel diffusion-based unpaired image restoration method that efficiently restores low-quality Raman-based optical images with minimal to no perceptual artifacts or hallucinations. The intuition behind RSCD is that low-quality/corrupted/degraded images are images generated from an incomplete generative diffusion process, where the diffusion model is trained to generate high-quality images in T steps but only performs T - t steps, resulting in a lower-quality image. The objective, then, is to complete the generative diffusion process by (1) determining the number of steps t via a step calibrator model and (2) completing the remaining t steps via a generative diffusion model.

Both models can be trained using unpaired high-quality images. RSCD is hallucination-resistant because it only involves editing the image via noise removal rather than full image generation from a random prior as is conventionally done in diffusion-based image restoration methods [25]. Moreover, RSCD is time-efficient because empirically t is small compared to T, avoiding the need to complete a full reverse diffusion process. Since optical image degradation is unpredictable, we designed a dynamic recalibration process that evaluates the quality of the image *during* diffusion to reset t as needed. The well-trained step calibrator model and dynamic recalibration are the key novelties of RSCD that enable it to consistently restore biomedical optical images that have various degrees and distributions of image degradation, as shown in Figure 1.

We summarize our contributions as follows:

- We propose a novel unpaired image restoration method, RSCD, that is fast, reliable, and ideally suited for Raman-based optical imaging where the noise is unpredictable with varying strength and pattern across and within images.
- RSCD is evaluated against other image restoration baselines, and outperforms them on image quality metrics. RSCD also achieves state-of-the-art performance on various unpaired perceptual metrics.
- Optical imaging experts consistently prefer images restored via RSCD over other methods, and report minimal to no hallucinations during human evaluations.

- 4 Y. Lyu et al.
- RSCD can improve performance on downstream clinical computer vision tasks, including automated brain tumor diagnosis and deep tissue imaging.

2 Background

2.1 Intraoperative Raman-based optical imaging

In this paper, we focus on the restoration of stimulated Raman histology (SRH), a rapid and label-free optical imaging method based on Raman spectroscopy [12]. SRH is used for a wide range of biomedical imaging tasks and has been clinically validated for imaging fresh, unprocessed human tissues and surgical specimens [34]. SRH begins with imaging a tissue specimen at two Raman shifts, 2845cm⁻¹ and 2930cm⁻¹, which highlights the optical image features generated by the lipid and protein concentrations, respectively, to generate image contrast. SRH can capture high-resolution, diagnostic-quality images across multiple organs and tissues [19, 34]. A major advantage of SRH is that it can be performed rapidly (~1 minute) without the need for tissue processing or staining, making it ideally suited for intraoperative tissue evaluation and diagnosis during surgery.

2.2 Generative diffusion models

Denoising diffusion probabilistic models are used to generate high-quality and diverse images [8,15]. Training diffusion models consist of two processes. The first is the forward diffusion process, which gradually adds Gaussian noise to an image x_0 over T steps to obtain x_1, x_2, \ldots, x_T , where $x_t \sim N(\sqrt{1-\beta_t}x_{t-1};\beta_t I)$ for each $1 \leq t \leq T$ and β_1, \ldots, β_T follows a noise schedule. The total noise added over the T steps should be strong enough to reduce the image to Gaussian noise, such that $x_T \sim N(0, I)$. Since combining multiple steps of Gaussian noise results in Gaussian noise, we see that for each $t, x_t \sim N(\sqrt{\alpha_t}x_0, (1-\alpha_t)I)$ where $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$. The second phase is the restoration phase, where we train a model ϵ_{θ} that takes in a noised image $(x_t = \sqrt{\alpha_t}x_0 + \sqrt{1-\alpha_t}\epsilon$ where $\epsilon \sim N(0, I)$) and the step number t, and tries to predict the noise (ϵ) , and the training loss is $||\epsilon - \epsilon_{\theta}(x_t, t)||$. When we use the model to generate an image, we start by randomly sampling $x_T \sim N(0, 1)$, and we repeat this T times until we generate a high-quality clean image x_0 .

3 Methodology

3.1 Overview

The intuition behind RSCD is to view the observed low-quality image as a diffusion model's output after an incomplete reverse diffusion process, where the number of incomplete steps varies depending on the noise level of the image. Specifically, the low-quality image x_t is from a diffusion model's generative process that starts from $x_T \sim N(0, 1)$ and is intended to reach x_0 , a synthetic high-quality image, through step-wise noise removal. In this view, completing the remaining t steps of the reverse diffusion process will result in a high-quality output image.

RSCD uses a trainable step calibrator model that predicts the value of t for each low-quality image. The step calibration is crucial for restoration quality because, as illustrated in Figure 3, if t is too small, then the reverse diffusion process cannot complete a sufficient number of steps to restore the low-quality image. If t is too large, then the reverse diffusion process generates excessively smooth and homogenized images with hallucinations. Moreover, because the degree and type of image degradation varies between low-quality images, optimal t values for each image will vary and manually setting a fixed t value as a hyperparameter results in suboptimal image quality and perceptual features. To further improve restoration quality, we dynamically recalibrate the number of steps needed to restore the image during the reverse diffusion process. An overview of our method is shown in Figure 2.



Fig. 3: Importance of the Step calibrator. If we perform less than the optimal number of diffusion steps $(t < t_{pred})$, the image remains degraded. If we perform more than the needed steps $(t > t_{pred})$, the output image is excessively smooth, fine details are removed, and contains hallucinations (yellow arrow). $(t_{pred} = 34$ for this image.)

3.2 Training Data

Training data was generated from approximately 2500 patients who underwent intraoperative SRH imaging to evaluate tissue during surgery [34]. Whole slide SRH images are approximately 6000×6000 pixels, which are then divided into 256×256 pixel patches, resulting in ~1 million total patches. To obtain high-quality SRH images, optical imaging experts manually selected 4.5K high-quality patches. Similar to large language model data curation strategies [2], we used a classifier-based filter method to obtain the full high-quality diffusion model training set. A binary classifier was trained on the 4.5K annotated high-quality patches as positive examples. Negative examples were randomly selected from the remaining uncurated SRH patch dataset. The trained classifier was then used to split the full SRH dataset into low-quality/high-quality images, resulting in a total of ~840K high-quality SRH images for training.

3.3 Step calibrator

We used a ResNet-50 model with MLP prediction head as the step calibrator. During training, a step number $t \sim \mathcal{U}(0, T)$ is sampled and t steps of Gaussian noise are added to a high-quality image according to a cosine schedule [15]. The calibrator is trained to predict t using an L_2 loss between the prediction t_{pred} and t. A challenge we identified with this naive implementation is that the severity of image degradation varies within an optical image. We found that the trained model tends to predict steps based on the region with the least amount of degradation. To address this issue, we perform the following augmentation when training the step calibrator: we randomly divide the image into two regions, and we add t steps of noise to one of the regions; we then sample a second, smaller t' such that $0 \le t' \le t$, and add t' steps of noise to the other region. We train the model to predict the larger noise t to ensure the step calibrator favors calibrations that will restore the most degraded regions in the image. We show an example of this augmentation in Figure 11 and the full algorithm in Algorithm 2, both in Appendix C.1.

In practice, we set T = 1000. After training, the step calibrator was used to predict noise level on approximately 63K low-quality images that had the lowest score from section 3.2. The distribution of the predicted t_{pred} is shown in Figure 12 in Appendix C.1. We found that the majority of low-quality images require less than 100 steps to restore, and the number of low-quality images decreases exponentially as the number of steps increases. No low-quality image required more than 200 steps for image restoration.

3.4 Diffusion model

The diffusion model is trained as a generative diffusion model that generates high-quality SRH images unconditionally with cosine noise schedule. We follow a similar training objective and procedure as described

Algorithm 1 Restorative Step-Calibrated Diffusion with Step Calibration and Dynamic Recalibration: Sampling

1: Requires: Low-quality image x, step calibrator S, diffusion model ϵ_{θ} , recalibration interval d, total steps T, hyperparameters based on noise schedule $\alpha_1, ..., \alpha_T, \bar{\alpha}_1, ..., \bar{\alpha}_T, \beta_1, ..., \beta_T$ 2: $t \leftarrow S(x) \#$ Initial step calibration 3: $t_{end} = 0$ 4: while t > 0 do $x_t \leftarrow x$ 5: $t_{end} = max(t - d, 0) \ \# \ d$ -steps of reverse diffusion process 6: for k in $[t, t - 1, ..., t_{end} + 1]$ do 7: $z \sim N(0,1)$ $x_{k-1} \leftarrow \frac{1}{\sqrt{\alpha_k}} (x_k - \frac{1 - \alpha_k}{\sqrt{1 - \alpha_k}} \epsilon_{\theta}(x_k,k)) + \sqrt{\beta_k} z$ 8: 9: 10: end for $x \leftarrow x_{t_{end}}$ 11: $t \leftarrow S(x) \#$ Dynamic recalibration 12:13: end while 14: return x

in [15], except that we use a shortcut for training efficiency: RSCD only requires training the model with noise level randomly sampled between 1 and T' steps (where T' just need to be larger than the maximum t_{pred} of low-quality images) instead of sampling from the full range 1 to T when training diffusion-based generation models, as shown in Algorithm 3 in Appendix C.2. In practice, we set T' = 200 due to the distribution in of t_{pred} as discussed in the previous section. This allows us to train the model more efficiently than conventional diffusion-based image restoration methods [25]. We first train the model for one pass through all 840K images, then fine-tune the model on the 4.5K high-quality images for 20 epochs.

3.5 Dynamic Recalibration

The step calibrator is trained to predict Gaussian noise levels; however, Raman-based optical image degradation is not limited to Gaussian noise. We found that, while empirical Gaussian noise levels are a good indicator of the degree of overall SRH image degradation, step calibrations that under or overestimate the required number of diffusion steps can result in poor image restorations. Therefore, to better calibrate the number of steps during the image restoration process, we perform dynamic recalibration: after we predict t_{pred} for an input low-quality image, instead of directly performing all t_{pred} steps of denoising, we only apply d steps of denoising, i.e. $x_{t_{pred}} \rightarrow x_{t_{pred}-1} \rightarrow \dots \rightarrow x_{t_{pred}-d}$. Then, we will use the step calibrator to predict the remaining steps of denoising needed for $x_{t_{pred}-d}$, and continue denoising starting from the updated predicted number of steps. Another d steps of denoising is performed before additional calibration. We repeat the process until the predicted number of steps remaining is less than d, and then we complete the remaining steps without additional recalibration. The process is illustrated in Algorithm 1.

4 Experiments

4.1 Baselines and Ablations

In the following experiments, we compare RSCD to the several unpaired image restoration baselines: CycleGAN [46], synthetic noise/noise2noise [26, 30], conditional diffusion [25, 36], regularized reverse diffusion (RRD) [5], deep image prior [38], and median blur. In addition, we also conduct ablation studies over the following design choices: dynamic recalibration, step calibrator, and cosine noise schedule. We compare our method against no dynamic recalibration, no step calibrator (either using a non-parametric



Fig. 4: Visual comparison of unpaired image restoration methods. CycleGAN hallucinates/inpaints nuclei within non-cellular structures (yellow arrow). Synthetic noise and Deep Image Prior (DIP) produce overly smoothed, unrealistic images. Conditional diffusion and Regularized Reverse Diffusion (RRD) generally perform insufficient image restoration.

noise estimator [3] as replacement or denoise for a fixed number of steps), evaluating the commonly-used linear noise schedule versus the cosine schedule, and training our step calibrator without augmentation. Details of each baseline and ablation can be found in Appendix D.1.

4.2 Evaluation on Unpaired Images

We evaluate the perceptual quality of image restoration on 12K unpaired low-quality images sampled from the largest public SRH dataset, OpenSRH [22]. We evaluate RSCD against the above baseline methods and ablations. We use Frechet Inception Distance (FID) between the restored images and the 4.5K expert-selected high-quality images as the perceptual quality metric [14].

We report all results in Table 1. RSCD outperformed all baselines and ablations, indicating that our method produces high-quality and realistic image restorations of low-quality Raman-based optical imaging. We provide additional details on training/inference efficiency and compute resources in Appendix D.5.

		FID \downarrow
	Original unrestored LQ images	53.22
Baselines	Median Blur	95.28
	Deep Image Prior [38]	57.56
	Synthetic noise [26, 30]	58.24
	CycleGAN [46]	37.26
	Conditional Diffusion [25, 36]	47.66
	Regularized Reverse Diffusion [5]	38.37
Ours	RSCD	32.02
Ablations	No Dynamic Recalibration	34.62
	Nonparametric Noise Estimator	35.53
	No Step Calibration, 10 steps fixed	36.41
	No Step Calibration, 50 steps fixed	42.06
	Linear noise schedule	38.71
	No Augmentation during Step Calibrator Training	33.14

Table 1: Results restoring low-quality open-source SRH images from OpenSRH [22]. FID score is measured between all restored images and the 4.5K expert-selected high-quality SRH images. Our method outperforms all baselines and ablations.

		$ $ FID \downarrow	$\mathrm{PSNR}\uparrow$	SSIM \uparrow	$\mathrm{LPIPS}\downarrow$
	Paired HQ Images	0.00	100.00	1.000	0.000
	Paired LQ Images	47.64	26.97	0.604	0.598
Baselines	Median Blur	72.71	28.87	0.791	0.336
	Deep Image Prior [38]	41.51	29.27	0.806	0.284
	Synthetic noise [26, 30]	77.43	22.94	0.662	0.464
	CycleGAN [46]	22.03	28.39	0.779	0.667
	Cond Diffusion [25, 36]	34.58	28.07	0.740	0.377
	Regularized Reverse Diffusion [5]	23.37	27.89	0.682	0.459
Ours	RSCD	21.05	29.06	0.791	0.280
Ablations	No Dynamic Recalibration	21.44	28.99	0.781	0.281
	Nonparametric Noise Estimator	22.77	27.90	0.691	0.418
	No Step Calibrator, 10 steps fixed	22.58	28.29	0.721	0.357
	No Step Calibrator, 50 steps fixed	29.72	29.05	0.788	0.335
	Linear noise schedule	32.71	29.06	0.789	0.283
	No Augmentation during Step Calibrator Training	21.15	29.03	0.786	0.284

Table 2: Results of restoring paired low-quality/near-registered high-quality SRH images. FID score is computed between all restored low-quality images and all near-registered images. Bolded numbers are best in each metric, and underlined are second best. RSCD achieved the best FID and LPIPS scores as well as the second-best PSNR and SSIM across all baselines and ablations. PSNR, peak signal-to-noise ratio; SSIM, structural similarity index measure; LPIPS, learned perceptual image patch similarity.

4.3 Evaluation on Near-registered Images

Next, we aim to evaluate the quality and fidelity of image restoration using an approximately paired low-quality/high-quality SRH imaging dataset. As mentioned before, perfectly paired images are challenging to obtain because optical image degradation is conditional on several (stochastic) factors. However, it is possible to obtain noisy/clean images that are nearly paired, by scanning the same specimen with a cold and warm laser (cold laser is known to produce more noisy images) and then re-align the two images. Through this process, we obtained 2,135 pairs of near-registered low-quality/high-quality SRH patches. These pairs are called 'near-registered' because non-affine deformation and optical tissue sectioning make perfect spatial alignment impossible, therefore these images cannot be used as paired training data. Nevertheless, the near-registered SRH dataset serves as a useful evaluation benchmark for restoration quality and fidelity. Details of the collection process of near-registered images are in Appendix D.3.

We report image quality and restoration metrics comparing RSCD and above baselines in Table 2. Our method achieves the best FID, which means our restored images are perceptually the closest to real high-quality images, while most baselines produce blurry or unrealistic images that result in significantly



Fig. 5: Examples of paired low-quality/near-registered SRH images, and the restored image via RSCD. Even in the setting of severe image degradation, RSCD produces high-quality, high-fidelity restorations with no clinically significant hallucinations. See Section 4.3 for details.



Fig. 6: Results of human expert evaluations. (a) The human expert preferences between RSCD and baselines. (b) Percentage of restored images by each method that the experts indicated hallucinated. As shown, the experts overall preferred the restorations from RSCD over the baselines.

worse FID scores. RSCD also has the second highest PSNR and SSIM scores, only behind deep image prior. This indicates that the images restored by our method are close to the ground truth high-quality images while avoiding being excessively smooth and unrealistic (excessively smooth or blurry images, such as the ones produced by median blur or deep image prior, are known to have inflated scores on metrics based on pixel-wise MSE [40], including PSNR and SSIM). RSCD was the only method to achieve good metrics on both pixel-level and perceptual metrics. Moreover, deep image prior inference is over $30 \times$ slower compared to RSCD, making it difficult to apply in the intraoperative setting. Examples of near-registered image pairs restored using RSCD are shown in Figure 5.

We also conducted ablation studies on the near-registered images, and our method achieved top performance among all 4 metrics compared to ablations, justifying our design choices of using the step calibrator, cosine noise schedule, and dynamic recalibration.

4.4 Human Expert Preference

The primary application of RSCD is to restore low-quality SRH images for intraoperative surgical specimen analysis and diagnosis by clinicians and domain experts, so their opinions are **the most important metric** for evaluating SRH image restoration quality. We recruited three clinicians and three optical imaging experts to evaluate the quality of the restored images and provide preference ratings. Each rater was asked to give their preferences between RSCD restoration and the five baseline methods' restorations. For every preference task, the experts are given the original low-quality image and the two restored images, one from our method and one from the baselines, in random and blinded order. The rater then selected their preferred restored image based on restoration quality and fidelity. In addition, raters were asked if either restoration contained hallucinations. We obtained expert preferences on restorations of 100 randomly selected low-quality images, and we gave every preference task to two raters to measure inter-rater agreement. We include more details about the expert preference collection process in Appendix D.6.

We report the results in Figure 6. Human experts preferred RSCD restoration more often than baselines. Raters also report the least number of hallucinations in our method, less than deterministic methods such as median blur that can create artifacts with severely degraded images. Inter-rater agreement was over 80% for both quality and hallucination assessments. Our method generates high-quality and reliable restorations of low-quality SRH images according to clinicians and domain experts.

5 Downstream Clinical Tasks

In the previous section, we showed that RSCD can restore low-quality SRH images to higher quality, more realistic, and more preferable images compared to other existing unpaired image restoration methods. In



Fig. 7: Classification accuracy of the SRH tumor classifier from [20] on low-quality images and restored images by RSCD. RSCD restored the majority of nondiagnostic images to diagnostic quality for tumor classification. Our method also significantly increased the classification accuracy and lowered the risk of brain tumor diagnostic errors.

this section, we demonstrate the downstream potential of RSCD in restoring SRH images on two clinical tasks: automated deep learning-based tumor diagnostics and z-stack restoration in deep tissue imaging.

5.1 Deep learning-based tumor diagnostics

Image degradation can decrease the performance of computer vision systems and result in decreased diagnostic accuracy for automated brain tumor classification [20]. In this experiment, we test whether RSCD can be used to facilitate more accurate deep learning-based diagnosis of SRH images by restoring model inputs.

We use the SRH tumor classification model from [20], the most widely accepted study on deep learningbased brain tumor diagnosis using SRH. The model classifies SRH images into one of 3 classes: normal tissue, tumor tissue, or non-diagnostic. The model classifies images as non-diagnostic if the image quality is sufficiently poor that human experts are not able to determine the underlying tissue diagnosis. The unpaired low-quality images from section 4.2 were used for evaluation. The original low-quality SRH image and the restored image underwent a feedforward pass through the SRH tumor classification model. We report the classification results in Figure 7.

Without restoration, over 50% of the SRH images were classified as non-diagnostic and another 26% were incorrectly classified. After restoration with RSCD, the classifier correctly diagnosed over 75% low-quality SRH images, and only miss-classified 6% of images. Restoring SRH images via RSCD can significantly improve deep learning-based automated diagnostic accuracy, and can drastically reduce non-diagnostic predictions caused by low image quality. Importantly, SRH restoration significantly reduces the risk of a wrong diagnosis, which can have a severe detrimental effect on patient care and surgical treatment. RSCD can make existing automated diagnostic tools safer and more reliable. We include more details and examples in Appendix E.1.

5.2 Z-stack Image Restoration

A known limitation of Raman-based optical imaging is that signal-to-noise ratios decrease as imaging depth increases due to laser scattering and absorption by tissue above the scanned depth. In this experiment, we show that RSCD can restore z-stack data, which contains SRH images acquired at sequentially deeper spatial locations in the tissue. z-stack images are volumetric and capture 3-D biological structures by imaging in all three spatial dimensions: height, width, and depth.

RSCD was used to restore z-stacked SRH images taken from 425 surgical specimens. Additional information regarding the dataset can be found in Appendix E.2. FID scores were computed at each z-depth level for original and restored images with respect to the 4.5K high-quality SRH images. We show the results in Figure 9. FID scores for the original low-quality images steadily increase as imaging depth increases, indicating worse image quality with imaging depth. RSCD consistently reduces the FID score for deep SRH images, cumulatively improving the overall quality of the z-stacked volume, shown in Figure 8.



Step-Calibrated Diffusion for Biomedical Optical Image Restoration 11

Fig. 8: Depth z-stack restoration results. The raw 3D z-stack volume (left) and the restored volume (right) are shown. Depth z-images in the center show image restoration using RSCD. The top and bottom plots show the pixel intensities along the row indicated in orange and green in the adjacent images. As depth increases from 22 microns to 60 microns, image noise increases as shown in the orange, low-quality pixel intensities. However, images restored by RSCD show little to no additional noise with increasing depth.

6 Related works

Several previous works share the similarity of using only the last part of the reverse diffusion generative process for image generation or restoration purposes (instead of requiring the full reverse diffusion) with our work. TDPM [45] increases the efficiency of diffusion-based image generation by training a GAN to replace the majority of the reverse diffusion steps and only perform the last part of the reverse diffusion process. SDEdit [32] and CCDF [6] perform image editing and restoration by first adding a strong Gaussian noise to the input image, then performing a few reverse diffusion steps. Regularized Reverse Diffusion (RRD) [5] builds upon CCDF: it does not require noise addition, uses a non-parametric Gaussian noise level estimator [3] to determine the number of reverse diffusion steps to perform, and regularizes the restoration process with Fourier features. In comparison, RSCD has the following advantages: (1) we designed and trained a highly effective step calibrator to determine the number of steps required, while TDPM, SDEdit,



Fig. 9: FID scores of original and restored z-stack images. Lower FID scores indicate better perceptual quality. RSCD consistently improves image quality at all depths.

and CCDF set this as a fixed hyperparameter and RRD determines this with a non-parametric Gaussian noise level estimator, all of which does not work well when the input image has non-uniform noise of varying strengths (as evident by our ablation studies on fixed steps and nonparametric noise estimator); (2) Adding a large amount of noise at once to the input image (like SDEdit and CCDF) increases the risk of hallucination, which is unacceptable for medical images; instead, RSCD gradually adds back a small amount of Gaussian noise during each DDPM step to mitigate the noise distribution difference between real noise and Gaussian noise; (3) RSCD uses dynamic recalibration to ensure consistency and stability of our restoration process, instead of Fourier feature-based regularization terms used by RRD. Since RRD was originally designed specifically for restoring MRI noises, its Fourier feature-based regularization terms do not work well for SRH image noises, as evident by our experiment results.

We include a more comprehensive discussion of related works on image restoration and deep learning applications in biomedical optical imaging in Appendix A.

7 Limitations

Our work is motivated by a biomedical imaging problem and intentionally focuses on biomedical optical imaging in the clinical setting. Our strategy is motivated by the problem of restoring medical images corrupted through non-reproducible, stochastic degradation sources and must be resistant to hallucinations. While this problem is not limited to biomedical optical imaging, other open-source datasets with non-reproducible low-quality images are not available. Additional studies are needed to confirm whether RSCD generalizes to other settings as additional related datasets are released. Another limitation is that the current version of RSCD does not take into account neighboring spatial regions when restoring a specific field of view. SRH image patches are sampled from megapixel images and consistent restoration should be enforced across the full image. We are currently exploring strategies to incorporate neighboring patches into the restoration process to improve restoration consistency, such as sliding windows attention [10] or unified conditional diffusion [43].

8 Conclusion

We present Restorative Step-Calibrated Diffusion (RSCD), a reliable and efficient method to restore biomedical optical imaging without requiring paired high-quality data. RSCD outperforms other widely used unpaired image restoration methods on both image quality and perceptual evaluation metrics. More importantly, experts in biomedical optical imaging consistently prefer images restored using RSCD in blinded comparison experiments and report minimal to no hallucinations. RSCD can improve model performance on downstream clinical tasks, including automated brain tumor diagnosis and deep tissue imaging. Our method reduces diagnostic errors that can have detrimental impacts on clinical care. This study demonstrates how computer vision has the potential to improve automated clinical diagnostics and patient care in today's precision medicine landscape.

9 Acknowledgements

We would like to thank Karen Eddy, Lin Wang, Hubert Zhang for their administrative support and data collection.

This work was supported by the following National Institute of Health (NIH) funding sources: K12NS080223 (T.C.H.), T32GM141746 (C.J.), F31NS135973 (C.J.). This work was support by the Chan Zuckerberg Foundation (CZI) Advancing Imaging Through Collaborative Project grant, Cook Family Brain Tumor Research Fund (T.C.H), the Mark Trauner Brain Research Fund, the Zenkel Family Foundation (T.C.H.),

Ian's Friends Foundation (T.C.H.) and the UM Precision Health Investigators Awards grant program (T.C.H.).

This research was also supported, in part, through computational resources and services provided by Advanced Research Computing, a division of Information and Technology Services at the University of Michigan.

T.C.H. is a shareholder of Invenio Imaging, Inc., a company developing SRH microscopy systems.

References

- Belthangady, C., Royer, L.A.: Applications, promises, and pitfalls of deep learning for fluorescence image reconstruction. Nat. Methods 16(12), 1215–1225 (Dec 2019) 2
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Others: Language models are few-shot learners. Adv. Neural Inf. Process. Syst. 33, 1877–1901 (2020) 5
- Chen, G., Zhu, F., Heng, P.A.: An efficient statistical method for image noise level estimation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 477–485 (2015) 7, 11, 19, 20
- Choi, J., Kim, S., Jeong, Y., Gwon, Y., Yoon, S.: ILVR: Conditioning method for denoising diffusion probabilistic models (Aug 2021) 16
- Chung, H., Lee, E.S., Ye, J.C.: MR image denoising and Super-Resolution using regularized reverse diffusion. IEEE Trans. Med. Imaging 42(4), 922–934 (Apr 2023) 6, 7, 8, 11, 16, 19, 20
- Chung, H., Sim, B., Ye, J.C.: Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12413–12422 (2022) 11, 16
- 7. Derpanis, K.G.: Overview of the ransac algorithm. Image Rochester NY 4(1), 2–3 (2010) 20
- 8. Dhariwal, P., Nichol, A.: Diffusion models beat GANs on image synthesis (May 2021) 4
- Du, W., Chen, H., Yang, H.: Learning invariant representation for unsupervised image restoration. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 14471–14480. IEEE (Jun 2020) 16
- 10. Esser, P., Rombach, R., Ommer, B.: Taming transformers for High-Resolution image synthesis (Dec 2020) 12
- Fei, B., Lyu, Z., Pan, L., Zhang, J., Yang, W., Luo, T., Zhang, B., Dai, B.: Generative diffusion prior for unified image restoration and enhancement. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9935–9946 (2023) 16
- Freudiger, C.W., Min, W., Saar, B.G., Lu, S., Holtom, G.R., He, C., Tsai, J.C., Kang, J.X., Xie, X.S.: Labelfree biomedical imaging with high sensitivity by stimulated raman scattering microscopy. Science **322**(5909), 1857–1861 (Dec 2008) 4
- He, J., Li, Y., Yuan, Q., et al.: Tdiffde: A truncated diffusion model for remote sensing hyperspectral image denoising. arXiv preprint arXiv:2311.13622 (2023) 16
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: GANs trained by a two time-scale update rule converge to a local nash equilibrium. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. pp. 6629–6640. NIPS'17, Curran Associates Inc., Red Hook, NY, USA (Dec 2017) 7
- Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: Proceedings of the 34th International Conference on Neural Information Processing Systems. pp. 6840–6851. No. Article 574 in NIPS'20, Curran Associates Inc., Red Hook, NY, USA (Dec 2020) 4, 5, 6, 17, 20
- Hoesli, R.C., Orringer, D.A., McHugh, J.B., Spector, M.E.: Coherent raman scattering microscopy for evaluation of head and neck carcinoma. Otolaryngol. Head Neck Surg. 157(3), 448–453 (Sep 2017) 1
- 17. Hollon, T., Jiang, C., Chowdury, A., Nasir-Moin, M., Kondepudi, A., Aabedi, A., Adapa, A., Al-Holou, W., Heth, J., Sagher, O., Lowenstein, P., Castro, M., Wadiura, L.I., Widhalm, G., Neuschmelting, V., Reinecke, D., von Spreckelsen, N., Berger, M.S., Hervey-Jumper, S.L., Golfinos, J.G., Snuderl, M., Camelo-Piragua, S., Freudiger, C., Lee, H., Orringer, D.A.: Artificial-intelligence-based molecular classification of diffuse gliomas using rapid, label-free optical imaging. Nat. Med. (Mar 2023) 1, 16

- 14 Y. Lyu et al.
- Hollon, T., Lewis, S., Freudiger, C.W., Sunney Xie, X., Orringer, D.A.: Improving the accuracy of brain tumor surgery via raman-based technology. Neurosurg. Focus 40(3), E9 (Mar 2016) 2
- Hollon, T.C., Lewis, S., Pandian, B., Niknafs, Y.S., Garrard, M.R., Garton, H., Maher, C.O., McFadden, K., Snuderl, M., Lieberman, A.P., Muraszko, K., Camelo-Piragua, S., Orringer, D.A.: Rapid intraoperative diagnosis of pediatric brain tumors using stimulated raman histology. Cancer Res. 78(1), 278–289 (Jan 2018) 4
- 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, 2nd, G.M., 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.: Near real-time intraoperative brain tumor diagnosis using stimulated raman histology and deep neural networks. Nat. Med. (Jan 2020) 1, 10, 16, 22, 27
- Hou, X., Jiang, C., Kondepudi, A., Lyu, Y., Chowdury, A.Z., Lee, H., Hollon, T.C.: A self-supervised framework for learning whole slide representations (2024) 16
- Jiang, C., Chowdury, A., Hou, X., Kondepudi, A., Freudiger, C.W., Conway, K., Camelo-Piragua, S., Orringer, D.A., Lee, H., Hollon, T.C.: OpenSRH: optimizing brain tumor surgery using intraoperative stimulated raman histology. Adv. Neural Inf. Process. Syst. 35(DB), 28502–28516 (Dec 2022) 7, 16
- Jiang, C., Hou, X., Kondepudi, A., Chowdury, A., Freudiger, C.W., Orringer, D.A., Lee, H., Hollon, T.C.: Hierarchical discriminative learning improves visual representations of biomedical microscopy. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 19798–19808 (Mar 2023) 16
- Kadkhodaie, Z., Simoncelli, E.: Stochastic solutions for linear inverse problems using the prior implicit in a denoiser. Adv. Neural Inf. Process. Syst. 34, 13242–13254 (2021) 16
- 25. Kawar, B., Elad, M., Ermon, S., Song, J.: Denoising diffusion restoration models (Jan 2022) 2, 3, 6, 7, 8, 16, 19
- Lehtinen, J., Munkberg, J., Hasselgren, J., Laine, S., Karras, T., Aittala, M., Aila, T.: Noise2Noise: Learning image restoration without clean data. In: Dy, J., Krause, A. (eds.) Proceedings of the 35th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 80, pp. 2965–2974. PMLR (2018) 2, 6, 7, 8, 19
- 27. Liang, J., Cao, J., Sun, G., Zhang, K., Van Gool, L., Timofte, R.: SwinIR: Image restoration using swin transformer (Aug 2021) 2
- Lowe, D.G.: Distinctive image features from Scale-Invariant keypoints. Int. J. Comput. Vis. 60(2), 91–110 (Nov 2004) 20
- Lui, H., Zhao, J., McLean, D., Zeng, H.: Real-time raman spectroscopy for in vivo skin cancer diagnosis. Cancer Res. 72(10), 2491–2500 (May 2012) 1
- Manifold, B., Thomas, E., Francis, A.T., Hill, A.H., Fu, D.: Denoising of stimulated raman scattering microscopy images via deep learning. Biomed. Opt. Express 10(8), 3860–3874 (Aug 2019) 1, 2, 6, 7, 8, 16, 19
- Mannas, M.P., Jones, D., Deng, F.M., Hoskoppal, D., Melamed, J., Orringer, D.A., Taneja, S.S.: Stimulated raman histology, a novel method to allow for rapid pathologic examination of unprocessed, fresh prostate biopsies. Prostate 83(11), 1060–1067 (Aug 2023) 1
- Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.Y., Ermon, S.: Sdedit: Guided image synthesis and editing with stochastic differential equations. arXiv preprint arXiv:2108.01073 (2021) 11
- Ongie, G., Jalal, A., Metzler, C.A., Baraniuk, R.G., Dimakis, A.G., Willett, R.: Deep learning techniques for inverse problems in imaging. IEEE Journal on Selected Areas in Information Theory 1(1), 39–56 (May 2020) 16
- 34. 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.: Rapid intraoperative histology of unprocessed surgical specimens via fibre-laser-based stimulated raman scattering microscopy. Nat Biomed Eng 1 (Feb 2017) 1, 4, 5, 16
- Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation (May 2015) 16
- 36. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S.K.S., Gontijo-Lopes, R., Ayan, B.K., Salimans, T., Ho, J., Fleet, D.J., Norouzi, M.: Photorealistic Text-to-Image diffusion models with deep language understanding (Oct 2022) 6, 7, 8, 16, 19

- 37. Sasaki, H., Willcocks, C.G., Breckon, T.P.: UNIT-DDPM: UNpaired image translation with denoising diffusion probabilistic models (Apr 2021) 16
- Ulyanov, D., Vedaldi, A., Lempitsky, V.: Instance normalization: The missing ingredient for fast stylization (Jul 2016) 2, 6, 7, 8, 16, 20
- Wang, Y., Yu, J., Zhang, J.: Zero-shot image restoration using denoising diffusion null-space model. arXiv preprint arXiv:2212.00490 (2022) 16
- 40. Wang, Z., Bovik, A.C.: Mean squared error: Love it or leave it? a new look at signal fidelity measures. IEEE Signal Processing Magazine **26**(1), 98-117 (2009). https://doi.org/10.1109/MSP.2008.930649 9
- Waterhouse, D.J., Fitzpatrick, C.R.M., Pogue, B.W., O'Connor, J.P.B., Bohndiek, S.E.: A roadmap for the clinical implementation of optical-imaging biomarkers. Nat Biomed Eng 3(5), 339–353 (May 2019) 1, 2
- Weigert, M., Schmidt, U., Boothe, T., Müller, A., Dibrov, A., Jain, A., Wilhelm, B., Schmidt, D., Broaddus, C., Culley, S., Rocha-Martins, M., Segovia-Miranda, F., Norden, C., Henriques, R., Zerial, M., Solimena, M., Rink, J., Tomancak, P., Royer, L., Jug, F., Myers, E.W.: Content-aware image restoration: pushing the limits of fluorescence microscopy. Nat. Methods 15(12), 1090–1097 (Dec 2018) 2
- Zhang, Y., Shi, X., Li, D., Wang, X., Wang, J., Li, H.: A unified conditional framework for diffusion-based image restoration. Adv. Neural Inf. Process. Syst. 36 (2024) 12
- 44. Zheng, D., Zhang, X., Ma, K., Bao, C.: Learn from unpaired data for image restoration: A variational bayes approach. IEEE Trans. Pattern Anal. Mach. Intell. **45**(5), 5889–5903 (May 2023) **16**
- 45. Zheng, H., He, P., Chen, W., Zhou, M.: Truncated diffusion probabilistic models and diffusion-based adversarial auto-encoders. arXiv preprint arXiv:2202.09671 (2022) 11
- Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired Image-to-Image translation using Cycle-Consistent adversarial networks (Mar 2017) 6, 7, 8, 16, 19
- Zhu, X., Yao, J., Zhu, F., Huang, J.: WSISA: Making survival prediction from whole slide histopathological images. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6855–6863 (Jul 2017) 2

A Additional Related Works

A.1 Image Denoising and Restoration

Image denoising and restoration belong to a set of problems called image inverse problems [33]. There are several different levels of difficulty for image denoising and restoration problems. The easiest level is when the corruption process is known as a linear process (such as blurring, inpainting, uniform Gaussian noise, etc) and can be programmatically replicated onto clean images, and the goal is to create an inverse process to the corruption. There have been extensive studies of restoring known corruptions, including supervised approaches [33] and unsupervised approaches [24, 25, 39]. The second level of difficulty is when the corruption process is not reproducible, but there are existing pairs of clean and corrupted images allowing for supervised training, and the typical solution is training models to generate the clean image given the corrupted images in a supervised way, such as UNet [30, 35] or Conditional Diffusion [4, 25, 36]. The most difficult setting is when we don't have paired data and also cannot reproduce the corruption programmatically; instead, we only have access to unpaired clean images and corrupted images. Existing approaches for this challenging setting include cycle-consistency-based approaches such as CycleGAN [46] and UnitDDPM [37], VAE-based approaches [44], and flow-based approaches [9]. In addition to the above settings, there are also denoising methods that do not require a training dataset, such as MedianBlur or Deep Image Prior [38]. General Diffusion Prior [11] is a recent unsupervised diffusion-based method that can perform unpaired non-linear blind image restoration, but it is extremely computationally costly as it requires full generative diffusion with multiple gradient steps within each reverse diffusion step.

Restoring noisy SRH images generally falls within the third difficulty level (i.e. no paired data and not a linear inverse problem), as the distribution of noise is unpredictable. Although it is possible to apply supervised methods from the second level by training to remove Gaussian and Laser noises (such as our UNet and Conditional Diffusion baselines), they are generally outperformed by unpaired methods as shown in our experiments. Several recent approaches (including [6], [5], and [13]) also used the last part of the reverse diffusion process for unpaired image restoration. We discussed these closely related works in more detail in Section 6.

A.2 Deep Learning applications on SRH

Deep learning techniques have been playing an important role ever since SRH was used for intra-operative neurosurgery, as the very first paper that introduces this optical diagnostic technique [34] trained a multilayer perceptron to automatically predicts brain tumor subtypes from quantified SRH image attributes. Later, Hollon et al. [20] proposed a CNN-based model that can directly interpret raw SRH images (instead of quantified attributes) and make diagnosis at over 94% accuracy. The first public dataset of clinical SRH images, OpenSRH [22], was released in 2022, and it contains SRH images from 300+ brain tumor patients and 1300+ unique whole slide optical images. DeepGlioma [17] has allowed deep-learning-based models to perform molecular classification of diffuse glioma types via a multimodal model that takes in both SRH images and diffuse glioma genomic data. HiDisc [23] and S3L [21] are self-supervised visual representation learning methods that were developed to classify whole-slide SRH images.

All the above works focused on the classification of SRH images, and there have been few previous works that attempted to tackle the unavoidable noise problem of SRH images through image restoration. To the best of our knowledge, the only previous work on this topic was [30], in which the authors collected a paired dataset using low and high laser energy, and trained a U-Net to perform the denoising. In comparison, our method works for all types of noisy SRH images (not just those resulting from low laser power), and we have shown that our method outperforms their U-Net approach.

B Details about SRH images

We include examples of image degradation for different cytologic/histologic features and brain tumor types in Figure 10. The examples show there exist many vastly different types of degradation that can randomly occur in low-quality SRH images, therefore making SRH image restoration a challenging task, but RSCD was able to restore them quite well.

C Methodology Additional Details

C.1 Step Calibrator Training and Calibration Distribution

We show the full algorithm of training the step calibrator with manually added Gaussian noise and augmentations in Algorithm 2, as well as one example of the augmentation in Figure 11.

Algorithm 2 Training step calibrator with augmentations

```
1: Requires: Set of high quality images X, total steps T, hyperparameters based on noise schedule \bar{\alpha}_1,...,\bar{\alpha}_T
 2: S \leftarrow \text{ResNet-50}
 3: repeat
           Sample x \in X
 4:
 5:
           t \sim \text{Uniform}(\{1, \dots, T\})
 6:
           \epsilon \sim N(0,1)
 7:
           x_t \leftarrow \sqrt{\bar{\alpha}_t}x + \sqrt{1 - \bar{\alpha}_t}\epsilon
           t' \sim \text{Uniform}(\{1,...,t\})
 8:
           \epsilon' \sim N(0,1)
 9:
10:
           x_{t'} \leftarrow \sqrt{\bar{\alpha}_{t'}}x + \sqrt{1 - \bar{\alpha}_{t'}}\epsilon'
11:
           m \leftarrow \text{Binary mask of random region}
12:
           x_t = mx_t + (1-m)x_{t'}
           L \leftarrow ||t - S(x_t)||_2
13:
           Train S with loss L
14:
15: until converged
```

To determine a proper T', we ran our trained step calibrator on 63K real low-quality SRH images, and the distribution of t_{pred} is shown in Figure 12. We can see that the number of images decreases exponentially with its t_{pred} , and t_{pred} never exceeded 200, which justifies our choice of T' = 200 in our diffusion model training.

C.2 Diffusion Model Training

We show the full training algorithm in Algorithm 3 below. It is almost identical to the training algorithm from [15], except that we only sample t up to T' = 200.

D Experiment details and additional results

D.1 Baseline and Ablation details

Below are the modeling and training details of the baselines:



Fig. 10: Examples of tissue-dependent sources of image degradation. Each row is an example of cytologic or histologic tissue features that can result in excessive laser scattering, absorption, and image degradation. First row: hypercellularity and dense chromatin (yellow arrows) increase the overall density of tissues. Second row: melanin and other pigments (yellow arrows) can have autofluorescence and reduce signal-to-noise ratios. Third row, blood vessels (yellow borders) contain collagen and elastin which are dense, proteinaceous structures that scatter laser. Last row: the same is true for lipid- or protein-dense intracellular inclusions and red blood cells (yellow arrows).

Algorithm 3 Restorative Step-Calibrated Diffusion: Training with shortcut

- 1: Requires: Set of high-quality images X, diffusion model ϵ_{θ} , total steps T, max training steps T', noise schedule hyperparameters $\bar{\alpha}_1, ..., \bar{\alpha}_T$
- 2: repeat 3: Sample $x_0 \in X$ 4: $t \sim \text{Uniform}(\{1,...,T'\})$ 5: $\epsilon \sim N(0,1)$ 6: Take gradient descent step on 7: $\nabla_{\theta} ||\epsilon - \epsilon_{\theta}(\sqrt{\overline{\alpha}_t}x_0 + \sqrt{1 - \overline{\alpha}_t}\epsilon, t)||^2$ 8: until converged

CycleGAN [46]: We trained a CycleGAN model for image-to-image translation between unpaired low-quality and high-quality images. We used the 4.5K hand-picked high-quality and 3.8K hand-picked low-quality images for training the CyleGAN. Both the high-quality and low-quality image sets contained at least 10 different tumor types. The CycleGAN model performed better when trained on the manually selected images compared to training on the full 840K filtered high-quality images.

Synthetic noise [26,30]: A U-Net model is trained by adding Gaussian noise to high-quality images to generate synthetic low-quality images. We train only with the 4.5K manually selected high-quality images because we found that the model performs better compared to training on the full 840K filtered high-quality images.

Conditional Diffusion [25, 36]: We train a conditional generative diffusion model, where we sample from a random Gaussian prior, and generate the restored image conditioned on the low-quality image. The conditioning method is concatenation (i.e. the input image x_t and the conditioning image is concatenated as input to the UNet). We train the model with the 840K HQ training images and fine-tune it with the 4.5K hand-picked images. During training, we add a random amount of Gaussian noise to the high-quality images as conditioning images. We show the full training algorithm in Algorithm.

Algorithm 4 Conditional Diffusion: Training

- 1: Requires: Set of high-quality images X, diffusion model ϵ_{θ} , total steps T, max training steps T', noise schedule hyperparameters $\bar{\alpha}_1, ..., \bar{\alpha}_T$
- 2: repeat 3: Sample $x_0 \in X$ 4: $t' \sim \text{Uniform}(\{1, \dots, T'\})$ $\epsilon' \sim N(0,1)$ 5: $x_{cond} \leftarrow \sqrt{\bar{\alpha}_{t'}} x_0 + \sqrt{1 - \bar{\alpha}_{t'}} \epsilon$ 6: 7: $t \sim \text{Uniform}(\{1, \dots, T\})$ 8: $\epsilon \sim N(0,1)$ 9: Take gradient descent step on 10: $\nabla_{\theta} || \epsilon - \epsilon_{\theta} (\sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, t, x_{cond}) ||^2$ 11: **until** converged

Regularized Reverse Diffusion (RRD) [5] is a recent method developed for restoring noisy medical images. It also uses part of the reverse diffusion process as restoration, but it employs a non-parametric noise estimator [3] to determine the step numbers, and uses a Fourier-feature-based regularization to reduce hallucination. Since we were unable to find any official code for the RRD paper and we deal with images in a vastly different domain from theirs (RRD was originally designed for denoising single-channel MRI

images), we trained our own diffusion model following the same process as [15] and using the same cosine noise schedule as RSCD.

Deep Image Prior [38] is an unsupervised image restoration method that works by training a ResNet to reconstruct the noisy image but takes the generation of the model before its convergence as the denoised image. We take the output of step 1800 as our baseline, as images generated from this training step tend to have the best quality.

Median Blur is a traditional, but effective, method to smooth noisy images. It replaces each pixel value with the median of the $k \times k$ pixel kernel (k = 5 for experiments).

Below are details of the ablations:

No Dynamic Recalibration: RSCD without Dynamic Recalibration, i.e. Algorithm 1 without the blue parts. We take the initial step calibration and directly run that many steps of denoising to the end.

Nonparametric Noise Estimator: We replace our ResNet-based step calibrator with the nonparametric noise estimator [3] used in RRD [5]. We determine the step number t as the closest cumulative noise $\sqrt{1-\bar{\alpha}_t}$ to the predicted noise level from the nonparametric noise estimator.

No Step Calibrator, 10 steps always: Instead of using a step calibrator and Dynamic Recalibration, we always only perform 10 steps of diffusion-based denoising on the low-quality images.

No Step Calibrator, 50 steps always: Instead of using a step calibrator and Dynamic Recalibration, we always only perform 50 steps of diffusion-based denoising on the low-quality images.

Linear noise schedule: We use a linear noise schedule instead of the cosine linear schedule.

No Augmentation Step Calibrator: We use a step calibrator that is trained without augmentations (i.e. Algorithm 2 without the red parts).

D.2 Unpaired Images Evaluation additional qualitative results

We show additional qualitative comparisons of restorations of unpaired images in Figure 13.

D.3 Near-registered dataset collection details

Imaging with low laser power is one of the known ways to produce noisy images. We generate a paired set of near-registered SRH images by imaging a surgical specimen immediately after the SRH imaging system is started (without laser warm-up and optimization), and then image the same specimen after 20 minutes, once the laser has been optimized and sufficiently warmed up. Next, image registration is required to spatially align the low-quality and high-quality images because the tissue specimen may be shifted due to gravity between the two image acquisitions. Paired images were registered by fitting an affine transformation using scale-invariant feature transform [28] and RANSAC [7].

While the near-registered images are matched enough for meaningful evaluation of fidelity, they are not fitting for training supervised image restoration because (1) the images are not perfectly paired (free-floating structures like red blood cells may have rotated/shifted relative to nearby cells), so training with them is just promoting hallucinations; (2) the noise in this dataset results from cold laser alone, so it does not well represent all types of random noise that could occur in SRH images from real intraoperative settings.

D.4 Hyperparameters

We list all architecture/training/inference hyperparameters for all components of RSCD in Table 3.

	Base Model	ResNet-50		
Step		2-Layer MLP		
Calibrator	Prediction Head	(2048, 100, 1)		
Training		ReLU activation		
maining	Learning Rate	0.00001		
	Optimizer	Adam		
	Epochs	500		
	Batch Size	150		
	Model	UNet		
	In channels	2		
Diffusion	Init dim	128		
Model	Out channels	2		
Architecture	Dim mult	[1,2,2,4]		
	Resnet block groups	8		
	Convnext Mult	2		
	Noise Schedule	Cosine $(s=0.008)$		
	Total Steps (T)	1000		
Diffusion	Max Shortcut Training	200		
Model	Steps (T')	200		
Training	Batch Size	8		
ITailling	Learning Rate	0.00001		
	Loss type	L1		
	Freeha	1 on 840K (full data)		
	Epocits	20 on 4.5 K (finetuning)		
	Optimizer	Adam		
Dynamic	Recalibration interval (d)	10		
Recalibration	necampiation interval (d)	10		

Table 3: Modeling, training and inference hyperparameters for all components of RSCD

D.5 Computing Resource and Inference Efficiency

Training the step calibrator takes about 5 hours on 1 A40 GPU. Training the diffusion model on 840K data took about 4 days on 1 A40 GPU, and fine-tuning on 4.5K hand-picked data took about 1 day on 1 A40 GPU.

We analyze the inference efficiency of RSCD on a single RTX 2080 GPU, in both runtime and NFE of the forward function of the UNet. On average over 1000 randomly selected low-quality SRH images, RSCD takes about 14.5 seconds to denoise a 256x256 patch (34 NFE), while methods that runs full generative reverse diffusion (e.g. Conditional Diffusion) take 431 seconds per patch (1000 NFE). Thus, RSCD is about 30x faster than running a full generative reverse diffusion process. Also, note that this number is obtained by denoising each patch individually without batching on an old and slow GPU for fair comparison. The amortized computation time can be improved by a lot if we perform the denoising steps in large batches and with a better GPU.

D.6 Human Expert Preference Details

We show the human expert preference collection interface in Figure 14. The expert is given the original LQ image as well as 2 restorations (one from RSCD and one from a baseline, in random order; the expert is not told which restoration is baseline and which one is RSCD), and then the expert is asked to (1) select one restoration that the expert believes has better overall quality, and (2) determine which restoration has clinically significant hallucinations (can choose "neither" if neither restoration hallucinates). We perform a total of 1200 such pair-wise preference collections (100 LQ images x 6 baseline comparisons x 2 preferences from different experts per comparison), distributed among all participating experts.

E Downstream task details

E.1 Deep learning-based tumor diagnosis

Data: A portion of our SRH images have ground truth labels on whether the tissue is normal or tumor. Therefore, we took the intersection of that labeled portion and the 12K unpaired low-quality images used for the experiment in section 4.2 as the dataset for deep learning-based tumor diagnosis. The dataset contains 8,372 low-quality images, each with its ground truth label of tumor/normal.

Model: We used the segmentation model from [20], the most influential and widely accepted study on deep learning-based brain tumor diagnosis through SRH images, as our classification model. The model classifies each SRH image into one of three classes: normal, tumor, and non-diagnostic. The model was trained to classify an SRH image as non-diagnostic if the model is unable to extract sufficient information from the image to make a diagnosis, and there are several possibilities for an image to be non-diagnostic, including: (1) the image scanning an irrelevant part of the slide, or (2) the image quality is too low.

Evaluation: For each image in our dataset, we run the same classification model on both the original and the restored version of the image. We categorize the predictions into 3 categories as follows:

- Correct: If the ground truth label matches the predicted label (i.e. predicting an image from a tumor tissue as tumor, or predicting an image from a normal tissue as normal)
- Wrong: If the ground truth label is the opposite of the predicted label (i.e. predicting an image from a tumor tissue as normal, or predicting an image from a normal tissue as tumor)
- Non-Diagnostic: If the predicted label is non-diagnostic

As shown in Table 7, the restored images have a significantly higher "correct" rate while having a much lower percentage of misclassified or non-diagnostic images. This means that RSCD can restore a significant portion of the images that were originally classified as non-diagnostic or wrongly classified due to poor quality, and make them high-quality enough to be correctly classified.

We include examples of original/restored images with their automated diagnosis results in Figure 15.

E.2 Z-stack Image Restoration

Image acquisition. Non-destructive z-stack SRH imaging often results in noisy images in deeper slices due to tissue scattering from the structures in the slices above, especially in specimens that contain blood or necrotic tissue. Z-stack SRH imaging follows the protocol for normal SRH imaging. Once the scanning of a whole slide is complete, the SRH imaging system automatically adjusts the laser focus, and an image at a higher depth is subsequently acquired. The z-stacked images are acquired starting at the default depth of 22 μm , with a z-resolution of 2 μm for 20 z-sections (reaching a final depth of 60 μm). Each z-stacked whole-slide is then patched into 256 × 256 × 20 volumes for denoising.

Additional results. We include additional examples of restored z-stack images spanning a variety of tumor types in Figure 16.



Fig. 11: An example of the noise augmentation applied for training better step calibrator models. In this example, we can see that the augmented image is more noisy around the edges (t = 100) while is less noisy around the center (t' = 32).



Fig. 12: Distribution of the initial step calibrations for low-quality SRH image restoration. The number of images decreases exponentially as the step number increases.



Fig. 13: Additional qualitative results for the unpaired experiment.



Fig. 14: The interface used for human expert preference. The expert is presented with the original LQ image on the left and 2 different restorations on the right. The expert is then asked to indicate their preference on which restoration has better overall quality and which ones have clinically significant hallucinations.



Fig. 15: Examples of RSCD improving automated diagnosis. Left: Panel of images before and after restoration. The classification model [20] incorrectly classified low-quality images of tumor tissue as normal brain tissue prior to image restoration. Following image restoration with RSCD, images were correctly classified as tumor tissue. Incorrect normal classification errors most commonly arise when tumor cells infiltrate a normal brain, which is the most common growth pattern for diffuse gliomas (shown here). Right: Panel of images classified as non-diagnostic prior to image restoration. Interestingly, the classification model can be sensitive to *even small amounts of image degradation*. In this clinical context, image degradation functions as adversarial noise with respect to brain tumor classification. Following image restoration with RSCD, the model correctly classifies the image as tumor tissue.



Fig. 16: Additional depth z-stack restoration results. The raw 3D z-stack volumes are shown on the left and the restored volumes using RSCD are shown on the right. The center images are the cross-section views of the first and the last slices, at 22 and 60 microns, respectively. Image noise increases as the tissue depth increases due to tissue scattering. With RSCD, the restored images show little to no additional noise in deep images.