Noise2Contrast: Multi-contrast Fusion Enables Self-supervised Tomographic Image Denoising

Fabian Wagner[1(✉)], Mareike Thies[1], Laura Pfaff[1,2], Noah Maul[1,2],
Sabrina Pechmann[3], Mingxuan Gu[1], Jonas Utz[4], Oliver Aust[5],
Daniela Weidner[5], Georgiana Neag[5], Stefan Uderhardt[5], Jang-Hwan Choi[6],
and Andreas Maier[1]

[1] Pattern Recognition Lab, FAU Erlangen-Nürnberg, Erlangen, Germany
`fabian.wagner@fau.de`
[2] Siemens Healthcare GmbH, Erlangen, Germany
[3] Fraunhofer Institute for Ceramic Technologies and Systems IKTS,
Hermsdorf, Germany
[4] Department AIBE, FAU Erlangen-Nürnberg, Erlangen, Germany
[5] Department of Rheumatology and Immunology, FAU Erlangen-Nürnberg,
Erlangen, Germany
[6] Division of Mechanical and Biomedical Engineering,
Ewha Womans University, Seoul, Korea

Abstract. Self-supervised image denoising techniques emerged as convenient methods that allow training denoising models without requiring ground-truth noise-free data. Existing methods usually optimize loss metrics that are calculated from multiple noisy realizations of similar images, e.g., from neighboring tomographic slices. However, those approaches fail to utilize the multiple contrasts that are routinely acquired in medical imaging modalities like MRI or dual-energy CT. In this work, we propose the new self-supervised training scheme Noise2Contrast that combines information from multiple measured image contrasts to train a denoising model. We stack denoising with domain-transfer operators to utilize the independent noise realizations of different image contrasts to derive a self-supervised loss. The trained denoising operator achieves convincing quantitative and qualitative results, outperforming state-of-the-art self-supervised methods by 4.7–11.0%/4.8–7.3% (PSNR/SSIM) on brain MRI data and by 43.6–50.5%/57.1–77.1% (PSNR/SSIM) on dual-energy CT X-ray microscopy data with respect to the noisy baseline. Our experiments on different real measured data sets indicate that Noise2Contrast training generalizes to other multi-contrast imaging modalities.

Keywords: Self-Supervised Denoising · Known Operator Learning · Contrast Fusion

1 Introduction

Measured data is inherently affected by uncertainty determined by the measurement process and its related physics. In image data, that uncertainty appears
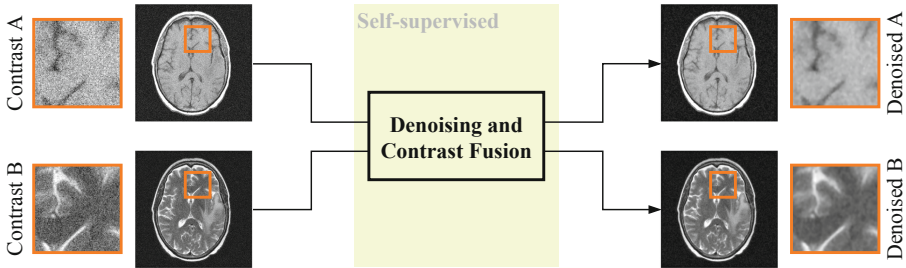
Fig. 1. Noise2Contrast: Fusion of image contrasts A and B enables self-supervised denoising, e.g., using T1 and T2 weighted MRI scans.

as image noise, disturbing an underlying ground truth image signal. Whereas imaging parameters like acquisition time, detector sensitivity, or illumination can be chosen to keep noise levels low, realistic imaging settings usually require a trade-off between acquisition parameters and image quality. In fact, some measurements, e.g., in clinical workflows, can only be carried out by accepting severe amounts of noise due to radiation exposure, acquisition time, or patient motion. Therefore, image processing algorithms were developed to reduce noise levels and extract the underlying noise-free signal. Conventional algorithms robustly denoise image data but require expert knowledge to adapt the algorithm to domain-specific conditions [14]. Unlike conventional filters, learning-based models can learn task-specific features purely from a training data distribution without domain-specific knowledge. However, deep neural networks inherently lack interpretability and were shown to be prone to prediction artifacts on out-of-domain samples [15]. Different hybrid approaches tried to combine data-driven optimization with conventional image filters to create reliable denoising operators with close to state-of-the-art performance [16].

Recently, multiple self-supervised denoising methods were proposed, circumventing the need for ground truth noise-free data during training [1,6–8]. Noise2Noise [8] and Noise2Void [7] allow self-supervised image denoising using two noisy representations of the same image or pixel-wise masking to calculate loss metrics that do not require a ground truth. Different other works applied these concepts to medical imaging modalities, e.g., by using neighboring volumetric slices [5,19] or time frames [18] as training targets following the Noise2Noise scheme.

Although self-supervised training on individual medical scans showed promising results, most existing approaches are not capable of using all available data. Many used medical imaging modalities like Magnetic Resonance Imaging (MRI) or dual-energy Computed Tomography (DECT) routinely acquire multiple image contrasts of the same scanned object that remain so far unused in self-supervised denoising approaches. In this work, we present the novel denoising method Noise2Contrast which is capable of using multiple image contrasts to train a denoising model in a fully self-supervised manner. An overview of Noise2Contrast

is illustrated in Fig. 1. Our method is able to employ the independent noise realizations in different image contrasts of medical imaging modalities to train a robust denoising operator. We confirm our theoretical considerations with extensive experiments on real medical data. Our contributions are three-fold.

– We present the self-supervised denoising method Noise2Contrast combining image information from different acquired image contrasts.
– We demonstrate how to train a robust denoising operator using our proposed scheme by simultaneously learning denoising and domain transformation.
– Extensive experiments quantitatively and qualitatively confirm the applicability of our method on different real medical data sets.

2 Methods

2.1 Self-supervised Image Denoising

Each image acquisition j introduces noise n through the image formation and detection processes to the ground truth object y

$$x_i^{(j)} = y_i + n_j. \tag{1}$$

Image denoising then aims to find an operator f_w that maps noise-affected images $x_i^{(j)}$ to a denoised prediction $\hat{y}_i$ close to the noise-free ground truth y_i by minimizing

$$\underset{w}{\operatorname{argmin}} \sum_i \mathcal{L}\left(f_w\left(x_i^{(1)}\right), y_i\right) \tag{2}$$

based on a loss metric $\mathcal{L}$ and parameters w. Supervised learning methods typically use a training set of N paired samples $\left(x_i^{(1)}, y_i\right)$ with $i \in \{1, \ldots, N\}$ to train a neural network data driven to predict denoised images from the learned training data distribution. As paired ground truth images are often difficult to obtain in real applications, self-supervised training methods aim to find an optimal denoising operator while having solely access to noisy images. Lehtinen et al. [8] demonstrated that learning the mapping of the noisy measurement to a second image with the same content but a different noise realization $x_i^{(2)}$, e.g., a second photo taken, is similar to solving the supervised problem in Eq. 2

$$\underset{w}{\operatorname{argmin}} \sum_i \mathcal{L}\left(f_w\left(x_i^{(1)}\right), x_i^{(2)}\right). \tag{3}$$

Although many works adopt this so-called Noise2Noise training scheme, the method requires at least two images with equivalent content and contrast per sample during training which might not be available in reality. Other works, e.g., Noise2Void [7], propose masking individual pixels of noisy images to create pseudo-paired training samples $x_i^{(1\star)}$. Subsequently, a denoising model can be trained by learning to predict the correct intensity values at the masked positions. However, Noise2Void demands pixel-wise statistically independent noise which is often not satisfied in particular on real detector data and for medical imaging modalities [16].

2.2   Denoising Using Known Operators

Including prior knowledge in neural network architectures has been shown bene-
ficial in terms of model performance, generalizability, and prediction robustness
[9,15]. We adapt the known operator learning concept in our proposed method
by separating denoising and domain-transfer tasks through the network archi-
tecture as described in Sect. 2.3. As the denoising operator, we use trainable
bilateral filter layers [16] that can be trained via gradient-based optimization
like any other neural network layer. The filter forward operation smooths image
content in homogeneous regions (spatial kernel) while preserving edges through
a range kernel

$$\hat{Y}_k = \frac{1}{\alpha_k} \sum_{n \in \mathcal{N}} G_{\sigma_s}(\|k - n\|) G_{\sigma_r}(X_k - X_n) X_n \tag{4}$$

with

$$\alpha_k = \sum_{n \in \mathcal{N}} G_{\sigma_s}(\|k - n\|) G_{\sigma_r}(X_k - X_n) \tag{5}$$

and G_{σ_s} and G_{σ_r} denoting Gaussian spatial and range kernel of width σ_s and σ_r
respectively. $\|\dots\|$ indicates the spatial distance between pixels of index k and
n and $\mathcal{N}$ is the filter window. The differentiable implementation of Wagner et
al. [16] allows optimizing all filter parameters σ_s and σ_r data driven using deep
learning frameworks. The algorithmic filter design from Eq. 4 proves that the
bilateral filter can solely act as a denoising operator as it is not able to extract
complex features or modify the images besides local pixel intensity averaging.

In addition to data-driven optimization of a known denoising algorithm, we
demonstrate how to employ a neural network as an independent denoising oper-
ator. By training denoising and domain-transfer networks subsequently, different
image processing tasks can be entirely separated into independent network parts
to enable self-supervised denoising of multi-contrast data. The proposed training
schemes are presented in the following section.

2.3   Multi-contrast Fusion Through Domain Transfer

Measuring two images with the same content to perform Noise2Noise denoising
is often infeasible in medical imaging due to radiation and time constraints. How-
ever, modalities like MRI or DECT routinely acquire multiple image contrasts
that show the same anatomical structures but highlight different biological fea-
tures. A second noisy image contrast indicated by † (e.g., T1 and T2 weighting
in MRI imaging)

$$x_i^{\dagger(j)} = y_i^{\dagger(j)} + n \tag{6}$$

could be used as a noise-affected target image in the setting of Eq. 3. In such
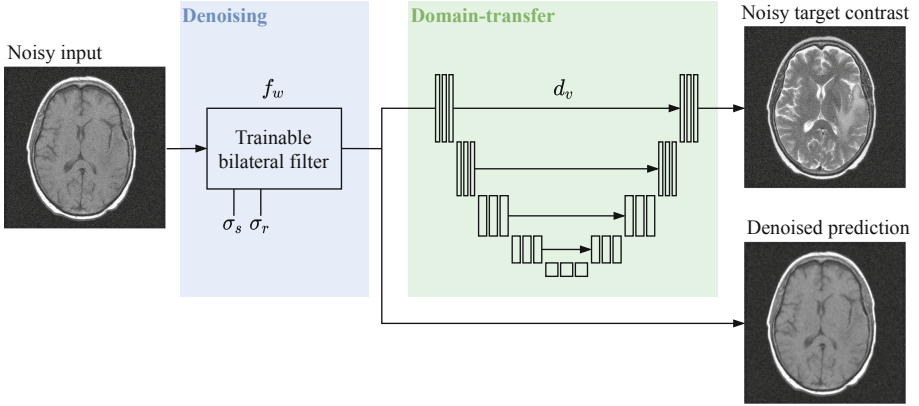a setting, a network would learn to predict a denoised image with the target

Fig. 2. Illustration of the proposed domain-transfer-based self-supervised denoising approach Noise2Contrast on the example of MRI T1 and T2-weighted contrasts. A noisy input of contrast one is processed by subsequent denoising (blue) and domain-transfer (green) operators. This allows deriving a self-supervised loss metric $\mathcal{L}$ using the noisy target with contrast two. The denoised input image is obtained by removing the domain-transfer operator. (Color figure online)

contrast. However, it is not possible to extract a solely denoised image $\hat{y}_i$ from the network prediction with preserved contrast. To avoid mixing both tasks, we propose separating the trained model into known operators to allow using the network parts individually during inference as we are only interested in the denoised prediction but want to preserve the original image contrast. An illustration of the presented training scheme is illustrated in Fig. 2. We present two solutions how to separate the denoising and domain translation tasks to enable self-supervised denoising.

Known Operator-Based. First, a known denoising operator is used in combination with a domain translation neural network d_v and trained self supervised. We use a trainable bilateral filter layer (Sect. 2.2) as the filter operation can not perform complex domain translations or intensity shifts by design and thus can be considered as a known denoising operator. Therefore, denoising and domain translation are inherently separated through the pipeline's architecture when training the chained operators d_v and f_w. The following training task results

$$\operatorname*{argmin}_{w,v} \sum_i \mathcal{L}\left(d_v\left(f_w\left(x_i^{(1)}\right)\right), x_i^{\dagger(1)}\right) \tag{7}$$

containing the domain translation operator d_v represented by a U-Net [11] with trainable parameters v. A self-supervised mean squared error loss is calculated between the denoised and domain-translated input image and the target contrast image $x_i^{\dagger(1)}$ with independent noise.

Network Operator-Based. Second, a neural network is trained as a denoising operator in the same setting as Eq. 7. To enforce a strict separation of denoising and domain translation, operators d_v and f_w are trained in a subsequent fashion. First, the domain translation network is trained in the known operator-based setting to predict images of target contrast $y_i^\dagger$ from denoised input contrast images $\hat{y}_i$. Subsequently, that trained network is frozen and employed as a domain translation operator to transfer the predictions of a denoising neural network to the target contrast domain. The sequential training of denoising and domain translation operator enforces the networks to learn tasks independently and use them as separate image processing operators.

3   Experiments

3.1   Data

We perform multiple experiments to investigate how noise can be effectively reduced in multi-contrast medical data without requiring noise-free ground truth data. First, we evaluate our method on three different MRI contrasts that are routinely used to identify tissue-specific properties and abnormalities: T1, T2, and Fluid Attenuated Inversion Recovery (FLAIR)-weighting. We use the public Brain-Tumor-Progression data set [12] consisting of clinical MRI head scans of 20 brain tumor patients and split it into twelve training, two validation, and six test patients. Each scan contains T1, T2, and FLAIR-weighted reconstructions that are used as input and target data to evaluate the proposed self-supervised denoising method. We simulate Gaussian noise as present in the real and imaginary part of complex-valued reconstructed MR images or in the phase-corrected magnitude images [10] and choose the noise standard deviation as 5 % of the maximum scan intensity.

In a second experiment, we compare denoising methods on a mouse tibia bone sample scanned in a dual-energy Zeiss Xradia 620 Versa X-ray Microscope (XRM). Tomographic XRM imaging is instructive for investigating bone-remodeling and bone-related diseases on the micrometer scale due to its high bone-to-soft tissue contrast. Here, dual-energy acquisitions allow quantitative measurements of bone density and sample composition [4]. However, dual-energy XRM measurements contain severe noise levels due to finite scan times and dose concerns in potential in vivo measurements [17]. We denoise a 1.5 h dual-energy scan (50 kV and 70 kV) and compare the predictions with a 14 h high-SNR acquisition that is regarded as ground truth. XRM scans are reconstructed using the pipeline of Thies et al. [13]. The two settings LE (low-energy) → HE (high-energy) and HE → LE are investigated.

3.2   Networks

Three stacked trainable bilateral filter layers [16] are employed as known operator-based denoising model f_w. The domain translation network d_v is represented by a standard U-Net [11] with 16 input features and around 1.1 Mio

trainable parameters v. We use the Adam optimizer with learning rate $5 \cdot 10^{-5}$ in all our experiments. Models are trained until convergence of the self-supervised training loss computed on the validation scans in each epoch (MRI data) or on the training scan (XRM data).

3.3 Denoising Experiments

Different contrast combinations are investigated for the MRI data set to evaluate the generalizability of our proposed self-supervised denoising approach. We chose the settings T1 → T2, T2 → T1, and T2 → FLAIR for our experiments with the respective input x_i and target $x_i^{\dagger}$ contrast domains input → target and trained a model for each setting individually.

We compare our methods to the state-of-the-art blind-spot training scheme Noise2Void (N2V) [7]. In addition, we compare to a different reference method where a target image $x_i^{(2)}$ is chosen as the neighboring slice of the input image $x_i^{(1)}$. Multiple related works apply this or similar principles to create pseudo-pairs of noisy images [2,5,19]. We denote the reference approach as Noise2Neighbor (N2N) in the following as it comes close to the initial Noise2Noise idea where two noisy images of the same contrast are available. Our known operator and network operator-based methods are denoted as Noise2Contrast (BFs) and Noise2Contrast (U-Net) respectively.

4 Results

Table 1. Quantitative denoising results on the Brain-Tumor-Progression [12] MRI test data set. (mean ± std) is calculated over the patients. The best-performing method is highlighted in bold.

Setting	Method	PSNR (mean ± std)	SSIM (mean ± std)
T1 → T2	Noisy baseline	26.02 ± 0.01	0.384 ± 0.059
	Noise2Contrast (BFs)	**36.76 ± 1.40**	**0.869 ± 0.021**
	Noise2Contrast (U-Net)	30.43 ± 0.20	0.385 ± 0.119
	Noise2Void (BFs) [7]	35.81 ± 1.56	0.847 ± 0.023
	Noise2Neighbor (BFs) [2,19]	32.49 ± 4.36	0.867 ± 0.047
T2 → T1	Noisy baseline	26.02 ± 0.02	0.444 ± 0.071
	Noise2Contrast (BFs)	**34.69 ± 1.91**	**0.865 ± 0.023**
	Noise2Contrast (U-Net)	33.19 ± 1.03	0.698 ± 0.070
	Noise2Void (BFs) [7]	34.30 ± 1.69	0.841 ± 0.024
	Noise2Neighbor (BFs) [2,19]	29.91 ± 3.80	0.828 ± 0.067
T2 → FLAIR	Noisy baseline	26.02 ± 0.02	0.444 ± 0.071
	Noise2Contrast (BFs)	**35.21 ± 1.70**	**0.871 ± 0.022**
	Noise2Contrast (U-Net)	21.74 ± 0.05	0.221 ± 0.135
	Noise2Void (BFs) [7]	34.30 ± 1.69	0.842 ± 0.023
	Noise2Neighbor (BFs) [2,19]	29.90 ± 3.80	0.828 ± 0.067

Table 2. Quantitative denoising results on the dual-energy XRM bone scan. (mean ± std) is calculated over the z-slices. The best-performing method is highlighted in bold.

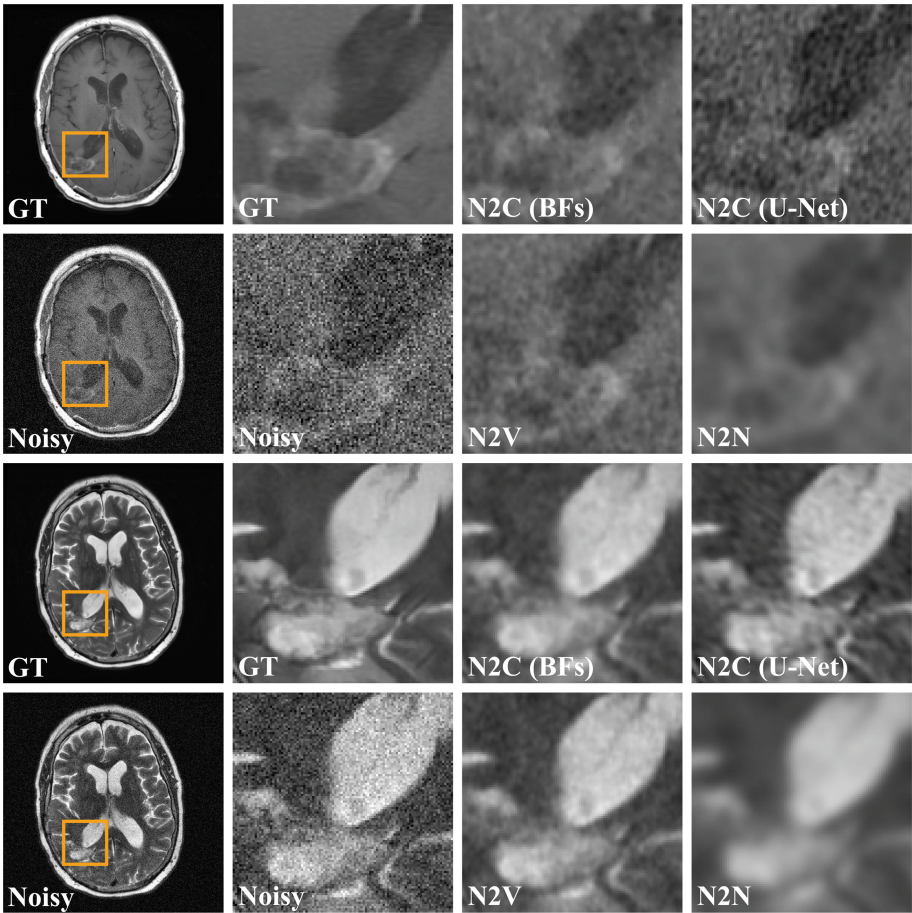| Setting	Method	PSNR (mean ± std)	SSIM (mean ± std)
LE → HE	Noisy baseline	22.17 ± 0.29	0.158 ± 0.008
	Noise2Contrast (BFs)	**29.86 ± 0.19**	**0.622 ± 0.015**
	Noise2Void (BFs) [7]	27.28 ± 0.22	0.420 ± 0.015
HE → LE	Noisy baseline	23.15 ± 0.29	0.178 ± 0.010
	Noise2Contrast (BFs)	**30.63 ± 0.22**	**0.610 ± 0.015**
	Noise2Void (BFs) [7]	28.36 ± 0.24	0.453 ± 0.015

Fig. 3. Qualitative denoising results on the Brain-Tumor-Progression [12] MRI test data set of T1 → T2 (top) and T2 → T1 predictions. The images are displayed in equal windows.
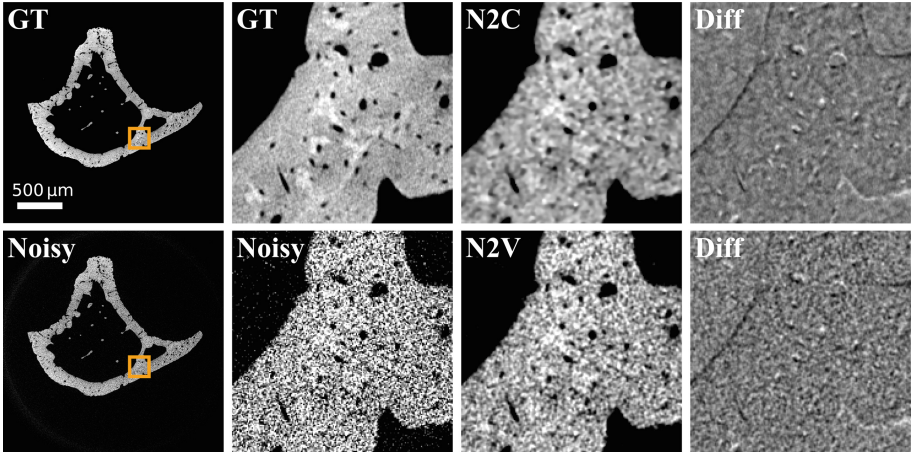
Fig. 4. Qualitative denoising results on the dual-energy XRM bone scan in the LE →
HE setting. The images are displayed in equal windows. Diff denotes the difference
images between the respective method and the high-dose ground truth.

We compute the quantitative image quality metrics peak signal-to-noise
ratio (PSNR) and structural similarity index measure (SSIM) for all model
predictions. Results on the Brain-Tumor-Progression data are presented in
Table 1. Our proposed multi-contrast training scheme using known opera-
tors Noise2Contrast (BFs) quantitatively outperforms all comparison methods.
Noise2Contrast (BFs) improves the results of Noise2Void by 4.7–11.0% PSNR
and by 4.8–7.3% SSIM with respect to the noisy baseline. Exemplary predic-
tions visualized in Fig. 3 confirm the quantitative findings and show that our
training scheme converges to a solution that preserves features while remov-
ing the image noise. Predictions of our additional experiment using a U-Net
for denoising (Noise2Contrast (U-Net)) exhibit lower noise removal compared to
the known operator-based method. State-of-the-art Noise2Void training achieves
similar visual results compared to our method, however, predictions contain a
slightly higher noise level. Noise2Neighbor fails to predict reasonable images and
blurs high-frequency features. The lower half of the magnified regions in Fig. 3
contains a brain lesion that allows comparing perceptual noise levels on a clinical
pathology.

Results on the dual-energy XRM data are presented in Table 2 and Fig. 4.
Here, Noise2Contrast (BFs) improves the results of Noise2Void by 43.6–50.5%
PSNR and by 57.1–77.1% SSIM with respect to the noisy baseline. On par
with the quantitative metrics, the visual predictions of Noise2Void contain con-
siderably more noise than our presented Noise2Contrast (BFs) training. This is
particularly visible in the provided difference images that are calculated between
the model predictions and the 14 h high-SNR XRM acquisitions. Note that the
low-dose network input and the high-dose ground-truth scans are independently

acquired scans. Despite the high mechanical precision of the used XRM, subsequent scanning results in small micrometer-scale shifts that are visible as thin edges in the difference images.

5   Discussion

The training configuration using a U-Net-based denoising model in combination with a domain-transfer model Noise2Contrast (U-Net) achieved promising visual results. However, the quantitative performance left room for improvement compared to the best-performing methods. We recognized that the trained denoising U-Net predicted visually appealing results but did not always fully preserve all input contrast intensities which led to poor quantitative metrics. We believe that a better-designed and more thoroughly trained domain-transfer model would help to provide more reasonable image gradients to the denoising network and improve the overall denoising performance. Pre-trained domain transfer models that are trained on ground truth data [3] can be employed here to improve the domain transfer operation. Alternatively, a regularizing loss term calculated between denoised input contrast and noisy input image can be investigated to enforce preserved intensities.

In the case of abnormalities being highlighted by the former image contrast but not being visible in the latter one or vice versa, the question arises if the Noise2Contrast training scheme can handle such samples and provide meaningful gradients to the denoising operator. We believe that as long as the domain transfer operator can learn a reasonable contrast transformation in particular a known and stable denoising operator like trainable bilateral filters that focus on local image properties can learn a reasonable denoising. In fact, the presented experiments on the Brain-Tumor-Progression MRI data contain such samples as tumors greatly vary in their visibility for different MRI contrasts. However, the generalizability of this finding must be proven on more clinical data and evaluated individually on the given imaging modality and set of pathologies.

We performed additional experiments directly mapping the input contrast to the target contrast image with a single model following the standard Noise2Noise approach. Although such models learned to simultaneously denoise and map to the target domain, their clinical application remains very limited as the model predictions inherently alter the image contrast which is usually not desired. In this Noise2Noise setting, image quality metrics calculated between model prediction and input contrast ground truth yielded poor scores as expected due to the modified prediction contrast. Additionally, we investigated a setting with a known denoising operator like the trainable bilateral filter used to predict the denoised input contrast by mapping on the target contrast without using a domain translation network. This yielded poor results likewise as the known denoising operator is not capable of learning the contrast mapping such that it only predicted blurred images to minimize the training loss.

Only a few fully self-supervised denoising techniques exist that can remove noise while preserving high-frequency image features. Blind-spot methods like

Noise2Void can achieve impressive results on certain data sets but are limited to pixel-wise independent noise statistics by design. Therefore, compelling results can be achieved on imaging modalities with simple noise characteristics and simulated data like the Brain-Tumor-Progression MRI scans in the first part of our study. Real measured data and computed tomography scans generally contain correlated noise caused by the detection process and the image reconstruction algorithm. Our experiments on real measured dual-energy XRM data confirm this limitation of Noise2Void. In contrast, our proposed known operator-based training scheme Noise2Contrast achieves considerably better quantitative and qualitative results as it relaxes prerequisites for specific noise properties like pixel-wise independent signals in the measured and reconstructed image data as demonstrated by the XRM experiments. Therefore, we conclude that Noise2Contrast is better suited to train models on modalities with correlated noise patterns like dual-energy CT compared to state-of-the-art Noise2Void training.

6   Conclusion

In this work, we presented the Noise2Contrast training scheme that allows self-supervised image denoising using multi-contrast data. Noise2Contrast combines information from independently measured image contrasts through an operator-based pipeline to train a denoising model. Our experiments on routine clinical MRI contrasts and on a pre-clinical dual-energy tomographic X-ray Microscope bone scan demonstrate superior performance of Noise2Contrast compared to the few other existing self-supervised denoising techniques. We believe that the universal Noise2Contrast training scheme can be applied on data from many more multi-contrast imaging modalities like photon-counting-CT, confocal microscopy, or hyperspectral imaging.

References

1. Batson, J., Royer, L.: Noise2Self: blind denoising by self-supervision. In: Proceedings of the ICML, pp. 524–533. PMLR (2019)
2. Choi, K., Lim, J.S., Kim, S.: Self-supervised inter-and intra-slice correlation learning for low-dose CT image restoration without ground truth. Expert Syst. Appl. **209**, 118072 (2022)

3. Denck, J., Guehring, J., Maier, A., Rothgang, E.: Enhanced magnetic resonance image synthesis with contrast-aware generative adversarial networks. J. Imaging **7**(8), 133 (2021)
4. Genant, H.K., Boyd, D.: Quantitative bone mineral analysis using dual energy computed tomography. Investig. Radiol. **12**(6), 545–551 (1977)
5. Jeon, S.Y., Kim, W., Choi, J.H.: MM-net: multi-frame and multi-mask-based unsupervised deep denoising for low-dose computed tomography. IEEE TRPMS 1–12 (2022)
6. Kim, K., Kwon, T., Ye, J.C.: Noise distribution adaptive self-supervised image denoising using tweedie distribution and score matching. In: Proceedings of the CVPR, pp. 2008–2016 (2022)
7. Krull, A., Buchholz, T.O., Jug, F.: Noise2Void-learning denoising from single noisy images. In: Proceedings of the CVPR, pp. 2129–2137 (2019)
8. Lehtinen, J., et al.: Noise2Noise: learning image restoration without clean data. In: Proceedings of the PMLR, vol. 80, pp. 2965–2974. PMLR (2018)
9. Maier, A.K., et al.: Learning with known operators reduces maximum error bounds. Nat. Mach. Intell. **1**(8), 373–380 (2019)
10. Prah, D.E., Paulson, E.S., Nencka, A.S., Schmainda, K.M.: A simple method for rectified noise floor suppression: phase-corrected real data reconstruction with application to diffusion-weighted imaging. Magn. Reson. Med. **64**(2), 418–429 (2010)
11. Ronneberger, O., Fischer, P., Brox, T.: U-net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
12. Schmainda, K.M., Prah, M.A.: Data from brain-tumor-progression. Technical report Version 1, The Cancer Imaging Archive (2018). https://doi.org/10.7937/K9/TCIA.2018.15quzvnb
13. Thies, M., et al.: Calibration by differentiation-self-supervised calibration for X-ray microscopy using a differentiable cone-beam reconstruction operator. J. Microsc. **287**(2), 81–92 (2022)
14. Tomasi, C., Manduchi, R.: Bilateral filtering for gray and color images. In: Proceedings of the ICCV, pp. 839–846. IEEE (1998)
15. Wagner, F., et al.: Trainable joint bilateral filters for enhanced prediction stability in low-dose CT. Sci. Rep. **12**(1), 1–9 (2022)
16. Wagner, F., et al.: Ultralow-parameter denoising: trainable bilateral filter layers in computed tomography. Med. Phys. **49**(8), 5107–5120 (2022)
17. Wagner, F., et al.: Monte Carlo dose simulation for in-vivo X-ray nanoscopy. In: Maier-Hein, K., Deserno, T.M., Handels, H., Maier, A., Palm, C., Tolxdorff, T. (eds.) Bildverarbeitung für die Medizin 2022. Informatik aktuell, pp. 107–112. Springer, Wiesbaden (2022). https://doi.org/10.1007/978-3-658-36932-3_22
18. Wu, D., Ren, H., Li, Q.: Self-supervised dynamic CT perfusion image denoising with deep neural networks. IEEE Trans. Radiat. Plasma Med. Sci. **5**(3), 350–361 (2020)
19. Zhang, Z., Liang, X., Zhao, W., Xing, L.: Noise2Context: context-assisted learning 3D thin-layer for low-dose CT. Med. Phys. **48**(10), 5794–5803 (2021)