



OPEN

Trainable joint bilateral filters for enhanced prediction stability in low-dose CT

Fabian Wagner[✉], Mareike Thies, Felix Denzinger, Mingxuan Gu, Mayank Patwari, Stefan Ploner, Noah Maul, Laura Pfaff, Yixing Huang & Andreas Maier

Low-dose computed tomography (CT) denoising algorithms aim to enable reduced patient dose in routine CT acquisitions while maintaining high image quality. Recently, deep learning (DL)-based methods were introduced, outperforming conventional denoising algorithms on this task due to their high model capacity. However, for the transition of DL-based denoising to clinical practice, these data-driven approaches must generalize robustly beyond the seen training data. We, therefore, propose a hybrid denoising approach consisting of a set of trainable joint bilateral filters (JBFs) combined with a convolutional DL-based denoising network to predict the guidance image. Our proposed denoising pipeline combines the high model capacity enabled by DL-based feature extraction with the reliability of the conventional JBF. The pipeline's ability to generalize is demonstrated by training on abdomen CT scans without metal implants and testing on abdomen scans with metal implants as well as on head CT data. When embedding RED-CNN/QAE, two well-established DL-based denoisers in our pipeline, the denoising performance is improved by 10%/82% (RMSE) and 3%/81% (PSNR) in regions containing metal and by 6%/78% (RMSE) and 2%/4% (PSNR) on head CT data, compared to the respective vanilla model. Concluding, the proposed trainable JBFs limit the error bound of deep neural networks to facilitate the applicability of DL-based denoisers in low-dose CT pipelines.

Minimizing patient dose in computed tomography (CT) is necessary to avoid radiation-related diseases¹, especially with the number of conducted diagnostic CT scans increasing every year². Low-dose CT acquisitions reduce patient dose^{3,4} but contain higher noise levels in the measured data^{5,6}. To enhance the image quality of low-dose CT acquisitions, image-based denoising approaches have been proposed, which aim to preserve clinically relevant features compromised with noise. Classical approaches are based on physically motivated conventional filters, considering the inherent properties of the image features^{7–11}. Although such filters produce reliable results through a clear algorithmic formulation, their performance is restricted by a limited capability to extract complex features. In addition, conventional filters often require hyperparameters that have to be tuned by hand. Therefore, deep learning (DL)-based denoising methods gained interest due to their flexibility, strong performance, and data-driven optimization^{12–17}. However, deep neural networks usually do not robustly generalize beyond their finite training data distribution, which so far limits clinical applications of DL-based denoising for low-dose CT^{18,19}.

Previously, Maier et al. proved that including physical knowledge in terms of known operators in neural networks reduces the absolute error bound of the model^{20–22}. Consequently, different image processing pipelines were proposed, employing physical assumptions about noise characteristics to leverage prediction reliability of DL-based methods in the context of image denoising^{23,24}. The joint bilateral filter (JBF) is a conventional denoising filter that allows edge-preserving denoising while considering additional information in terms of a guidance image during its filter operation. Imitating the JBF with a shallow convolutional network led to a reduction of trainable parameters in the JBFnet²³ and the MJBFB architecture²⁵. Although both network architectures are inspired by the JBF operation, they both learn filter operations through fully convolutional neural networks with a relatively large number of free parameters compared to the JBF. Therefore, both architectures can learn any possible filter kernels and are not enforced to perform the well-known JBF operation, which raises questions on data integrity and interpretability likewise to other DL methods²⁴. A different approach employs a custom bilateral filter approximation built from neural network building blocks that can be optimized²⁶, but it does not allow integration of additional learned information into the filter process. Other works presented methods to find optimal filter²⁷ or training²⁸ hyperparameters by predicting them through external neural networks. However,

Pattern Recognition Lab, Friedrich-Alexander-Universität Erlangen-Nürnberg, 91058 Erlangen, Germany. ✉email: fabian.wagner@fau.de

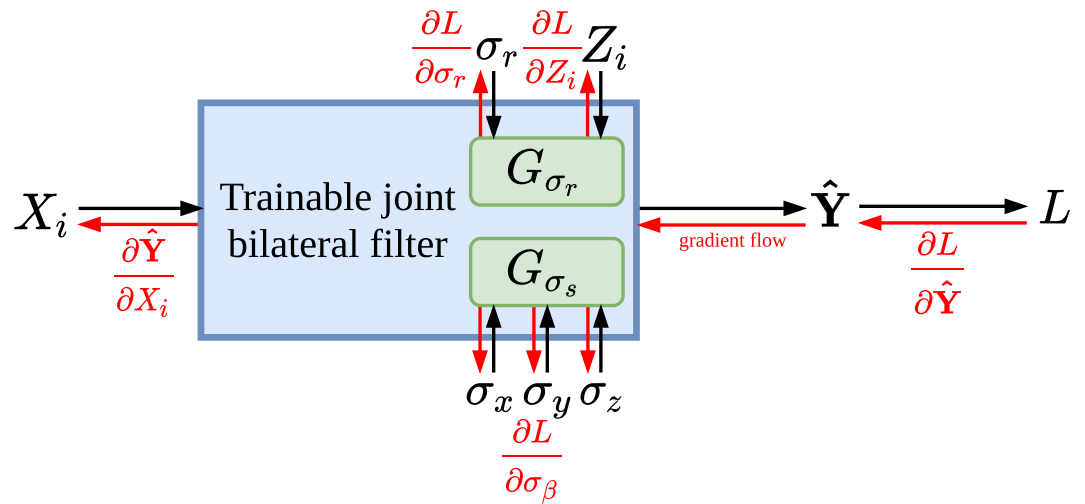


Figure 1. Illustration of the proposed trainable joint bilateral filter layer. In the forward pass (black arrows), the input X_i is filtered using parameters σ_γ ($\gamma \in \{x, y, z, r\}$) and the guidance image Z_i to predict the denoised image \hat{Y} . The model's loss is indicated as L . Analytical derivatives are calculated in the backward pass (red arrows) toward filter input, guide, and parameters.

such approaches do not allow for direct integration into DL models as they can not compute gradients toward those hyperparameters.

In our previous work, we presented a trainable bilateral filter with competitive denoising performance that can be included in a differentiable pipeline and optimized in a data-driven fashion²⁹. However, the prediction of bilateral filter layers is solely dependent on three learned spatial parameters and one intensity parameter⁹. Therefore, the bilateral filter operation is conceptually different from the joint bilateral filter algorithm, as JBFs allow considering additional information in terms of a guidance image in their denoising algorithm³⁰. In this work, we extend our research on bilateral filtering by proposing a fully differentiable, trainable joint bilateral filter that allows denoising using a learned guidance image which broadens its applicability. Our filter layer derives analytical gradients toward the filter input, the image guide, and all filter parameters to achieve differentiability and enable data-driven optimization. Guidance images are estimated using two well-established denoising algorithms: RED-CNN¹², an encoder–decoder architecture achieving competitive performance in recent works^{31,32}, and Quadratic Autoencoder (QAE)¹³, employing quadratic neurons. Our proposed hybrid filter model bridges the gap between deep neural networks' high model capacities and the robustness of conventional denoising filters due to the well-defined, restricted influence of the learned guide.

Contributions. Our contributions are threefold. First, we propose a GPU-based, trainable JBF based on an analytical gradient that can be included in any differentiable pipeline. To the best of our knowledge a directly trainable JBF was never presented before. Second, we introduce a hybrid denoising pipeline combining the flexibility of deep neural networks with the robustness of the trainable JBF. Third, we demonstrate the robustness of our model on abdomen CT scans containing metal, with metal not being present in the training data distribution and on out-of-domain head CT scans. Our hybrid JBF-based denoising setting improves the prediction reliability of existing DL-based models with limited computational overhead.

Methods

Artificial neural networks are generally trained via gradient descent optimization by minimizing a loss metric L calculated from network predictions to fulfill a desired task³³. This requires calculating the derivative of the loss L with respect to each trainable model parameter to iteratively update the network during training.

In this section, the analytical gradient of the proposed trainable JBF layer with respect to filter input, guidance image, and filter parameters is derived as the algorithmic contribution of our work. Figure 1 illustrates the general working principle of the denoising layer. In the forward filter operation an input image is convolved with two Gaussian kernels, namely one spatial and one range kernel. The spatial kernel averages pixels within the distance of the filter kernel like a conventional Gaussian filter that smooths the image. An additional, so-called range kernel weighs the influence of pixels from the neighborhood dependent on their intensity difference to the filtered pixel to prevent blurring of edges. The JBF derives its range kernel on an external guidance image which allows employing additional information during the filter operation.

In the following, bold letters are used to indicate vectors. According to Petschnigg et al.³⁰ the JBF operation is defined as

$$\hat{Y}_k = \frac{1}{w_k} \underbrace{\sum_{n \in \mathcal{N}} G_{\sigma_s}(\mathbf{p}_k - \mathbf{p}_n) G_{\sigma_r}(Z_k - Z_n) X_n}_{=: \alpha_k} \quad (1)$$

and the normalizing factor w_k as

$$w_k := \sum_{n \in \mathcal{N}} G_{\sigma_s}(\mathbf{p}_k - \mathbf{p}_n) G_{\sigma_r}(Z_k - Z_n), \quad (2)$$

with the denoised prediction \hat{Y} indexed by $k \in \mathbb{N}$, the noisy input image \mathbf{X} in the voxel neighborhood $n \in \mathcal{N}$ around k , and a guidance image \mathbf{Z} . Guidance images should provide additional information to the filter operation and can be, e.g., additional images paired with the filter input or learned predictions from a neural network as later introduced in this work. The Gaussian intensity range kernel

$$G_{\sigma_r}(c) := \exp\left(-\frac{c^2}{2\sigma_r^2}\right) \quad (3)$$

is derived from intensity differences on the guidance image \mathbf{Z} and enforces edge sensitivity of the filtering operation. A second, spatial filter kernel G_{σ_s} weights voxels according to their spatial distance derived from the positions $\mathbf{p}_k \in \mathbb{N}^d$ and $\mathbf{p}_n \in \mathbb{N}^d$ with $d = 3$ for three-dimensional filtering

$$G_{\sigma_s}(\mathbf{c}) = \prod_{s \in \{x, y, z\}} \exp\left(-\frac{c_s^2}{2\sigma_s^2}\right). \quad (4)$$

DL pipelines require gradient calculation of the loss function L with respect to each trainable parameter to enable data-driven optimization. We can calculate the gradient for our joint bilateral filter layer by using the chain rule

$$\frac{\partial L}{\partial \sigma_\gamma} = \frac{\partial L}{\partial \hat{Y}} \frac{\partial \hat{Y}}{\partial \sigma_\gamma} = \sum_k \frac{\partial L}{\partial \hat{Y}_k} \frac{\partial \hat{Y}_k}{\partial \sigma_\gamma} \quad (5)$$

with the four kernel widths σ_γ representing the only trainable weights of the proposed layer when filtering in three dimensions ($\gamma \in \{x, y, z, r\}$). The derivative of the loss function with respect to the filter prediction $\frac{\partial L}{\partial \hat{Y}}$ is provided by the backpropagation of the loss through differentiable operations applied on the JBF layer output, e.g., subsequent convolutional layers or the loss function itself. The term $\frac{\partial \hat{Y}_k}{\partial \sigma_\gamma}$ can be written using the definition of the joint bilateral filter algorithm from Eq. (1) together with the product and chain rule of differentiation

$$\frac{\partial \hat{Y}_k}{\partial \sigma_\gamma} = -w_k^{-2} \alpha_k \frac{\partial w_k}{\partial \sigma_\gamma} + w_k^{-1} \frac{\partial \alpha_k}{\partial \sigma_\gamma}, \quad (6)$$

the partial derivatives

$$\frac{\partial w_k}{\partial \sigma_\gamma} = \sum_{n \in \mathcal{N}} \frac{\partial}{\partial \sigma_\gamma} G_{\sigma_s}(\mathbf{p}_k - \mathbf{p}_n) G_{\sigma_r}(Z_k - Z_n), \quad (7)$$

$$\frac{\partial \alpha_k}{\partial \sigma_\gamma} = \sum_{n \in \mathcal{N}} X_n \frac{\partial}{\partial \sigma_\gamma} G_{\sigma_s}(\mathbf{p}_k - \mathbf{p}_n) G_{\sigma_r}(Z_k - Z_n), \quad (8)$$

and the Gaussian terms

$$\frac{\partial}{\partial \sigma_\gamma} G_\sigma(c) = G_\sigma(c) \frac{c^2}{\sigma_\gamma^3}. \quad (9)$$

In addition, the derivative of the loss with respect to each input voxel X_i of the joint bilateral filter yields

$$\begin{aligned} \frac{\partial L}{\partial X_i} &= \frac{\partial L}{\partial \hat{Y}} \frac{\partial \hat{Y}}{\partial X_i} = \sum_k \frac{\partial L}{\partial \hat{Y}_k} \frac{\partial \hat{Y}_k}{\partial X_i} \\ &= \sum_k \frac{\partial L}{\partial \hat{Y}_k} w_k^{-1} G_{\sigma_s}(\mathbf{p}_k - \mathbf{p}_i) G_{\sigma_r}(Z_k - Z_i) \end{aligned} \quad (10)$$

using the definition of the JBF from Eq. (1). This gradient calculation to the filter input is required to allow including the filter as a trainable layer into a differentiable pipeline. The derivative of L with respect to each voxel of the guidance image Z_i can be calculated as

$$\frac{\partial L}{\partial Z_i} = \sum_k \frac{\partial L}{\partial \hat{Y}_k} \frac{\partial \hat{Y}_k}{\partial Z_i} = \sum_k \frac{\partial L}{\partial \hat{Y}_k} \left(-w_k^{-2} \alpha_k \frac{\partial w_k}{\partial Z_i} + w_k^{-1} \frac{\partial \alpha_k}{\partial Z_i} \right) \quad (11)$$

where the following two cases must be distinguished: Case 1 derives gradients to arbitrary voxels located in the filter neighborhood ($k \neq i$) of the guidance image. In contrast, Case 2 defines the gradient to the center voxel ($k = i$) of the respective filter window.

Case 1: ($k \neq i$).

$$\begin{aligned}\left. \frac{\partial w_k}{\partial Z_i} \right|_{k \neq i} &= G_{\sigma_s}(\mathbf{p}_k - \mathbf{p}_i) G_{\sigma_r}(Z_k - Z_i) \frac{Z_k - Z_i}{\sigma_r^2} \\ \left. \frac{\partial \alpha_k}{\partial Z_i} \right|_{k \neq i} &= G_{\sigma_s}(\mathbf{p}_k - \mathbf{p}_i) G_{\sigma_r}(Z_k - Z_i) \frac{Z_k - Z_i}{\sigma_r^2} X_i\end{aligned}\quad (12)$$

Case 2: ($k = i$).

$$\begin{aligned}\left. \frac{\partial w_k}{\partial Z_i} \right|_{k=i} &= \sum_{n \in \mathcal{N}} G_{\sigma_s}(\mathbf{p}_i - \mathbf{p}_n) G_{\sigma_r}(Z_i - Z_n) \frac{Z_n - Z_i}{\sigma_r^2} \\ \left. \frac{\partial \alpha_k}{\partial Z_i} \right|_{k=i} &= \sum_{n \in \mathcal{N}} G_{\sigma_s}(\mathbf{p}_i - \mathbf{p}_n) G_{\sigma_r}(Z_i - Z_n) \frac{Z_n - Z_i}{\sigma_r^2} X_n.\end{aligned}\quad (13)$$

We calculate the analytical gradients in the backward pass of a fully trainable JBF using the CUDA binding of the *PyTorch* deep learning framework³⁴ to leverage computational performance. The processing time of one 512×512 image using $5 \times 5/11 \times 11$ pixel kernel windows is around 1.8 ms/8.0 ms on the GPU and 69 ms/350 ms on the CPU. In comparison, *torch.nn.Conv2d* layers (*PyTorch*) approximately require 0.1 ms/0.2 ms (GPU) and 8 ms/20 ms (CPU) for processing the single channel image. For both layers, gradient calculations have comparable run times as their forward passes. All run times were estimated by averaging 50 repeated forward/backward passes through the respective layers using a *NVIDIA Quadro RTX 4000* GPU. Please note that run times can strongly vary depending on the used hardware.

The filter window size of the JBF is chosen dynamically dependent on the spatial kernel sizes as $5 \cdot \sigma_s$. This ensures that $> 98\%$ of the Gaussian filter kernel mass is contained by the filter window which turned out to be a reasonable trade-off between accuracy and computational complexity of the algorithm.

Our filter layer is publicly available at <https://github.com/faebstn96/trainable-joint-bilateral-filter-source> and can be installed via the well-known Python Package Index (PyPI) as plug-and-play *PyTorch* layer. In addition, our code repository contains example scripts and a test script that compares the implementation of the analytical gradients with numerical gradient approximations using the *torch.autograd.gradcheck* function to make sure the filter derivative is correctly implemented.

Experimental setup

Denosing pipeline. Our denosing pipeline, illustrated in Fig. 2, is built on three consecutive trainable JBF layers. The iterative composition of filtering blocks is inspired by the design of the deep convolutional architecture JBFnet²³ and our previous experimental findings on using multiple stacked bilateral filters²⁹ which improved performance compared to employing only a single denosing step. The three trainable JBFs add in total twelve independently trainable parameters to the denosing model. The forward pass of each filter layer is calculated as written in Eq. (1). A guidance image is predicted from a deep convolutional network and used to derive the weighting of the intensity range kernel G_{σ_r} in each JBF. Multiple network configurations are presented in the following, investigating the influence of JBF layers on the denoised prediction.

Experiments. Our experiments are particularly designed to investigate the prediction robustness of hybrid JBF + DL-based denosing models compared to the respective vanilla DL model. We perform experiments with two different well-established low-dose CT denosing architectures predicting the guidance image: RED-CNN¹² and QAE¹³. In all our experiments, we train the two reference models independently as described in their works until full convergence of the validation loss, occurring after up to 300 epochs. Subsequently, we place the models in our denosing pipeline and optimize the JBFs for additional 200 epochs until convergence of the validation loss. Both trained vanilla deep neural networks are used as performance reference. We use the mean squared error loss and two separate Adam optimizers for σ_r ($l_r = 1 \cdot 10^{-2}$) and σ_s ($l_s = 5 \cdot 10^{-4}$) during training as both sets of parameters define filter kernels that act on independent scales. However, additional experiment where we used only a single Adam optimizer converged to very similar sets of filter parameters within comparable numbers of optimization steps. Therefore, we conclude that the network convergence is not overly sensitive to learning rate configurations when using an Adam optimizer.

Data. All used abdomen and head CT scans are from the public *TCIA Low Dose CT Image and Projection* data set (Version 4)³⁵, containing paired low-dose (25% dose) and high-dose CT volumes. The goal of our experiment is to quantitatively evaluate the robustness of the introduced denosing models and compare them with the vanilla DL-based denosing models RED-CNN and QAE. Therefore, we manually split the abdomen data into two domains. First, patients without metal pieces and second reconstructions containing pieces of metal like implants or catheters that appear as bright regions due to their strong x-ray absorption. Only data from the first domain not containing metal is used for training (21 scans) and validation (two scans). Subsequently, we

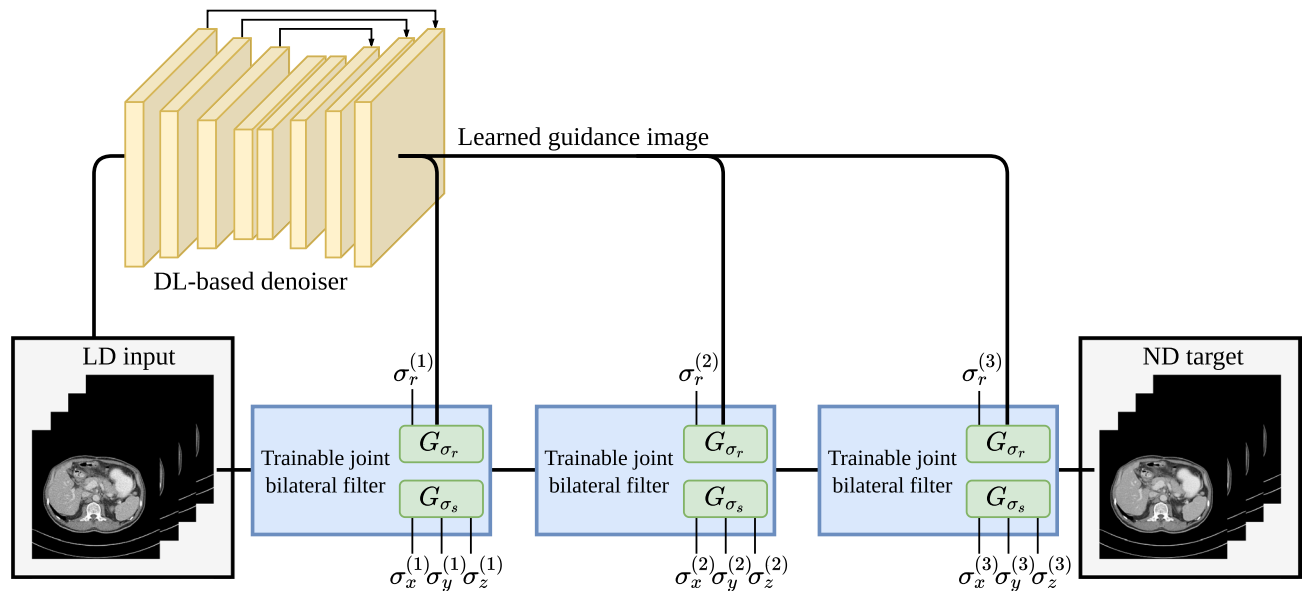


Figure 2. The investigated JBF-based denoising pipeline consists of three stacked trainable JBF layers to iteratively remove noise from the low-dose input reconstruction. Pre-trained RED-CNN and QAE models are employed as DL-based denoiser to predict guidance images. The model is trained supervised on the training domain and tested on CT data from other domains to investigate robustness properties of the pipeline. Indices (v) with $v \in \{1, 2, 3\}$ name the individual trainable JBF layers.

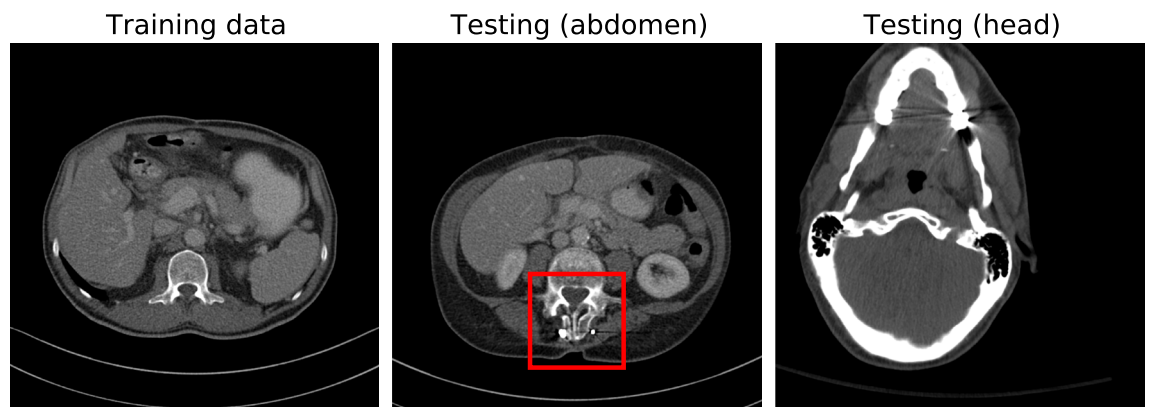


Figure 3. Exemplary slices from the training (left, abdomen without metal) and two testing (middle, right) data sets. The abdomen test patients contain regions with metal implants like the one highlighted in red. The reconstruction window is $[-150, 500]$ HU.

test our models on the previously unseen metal domain scans (24 scans) to evaluate how the different architectures can handle examples that are insufficiently represented by the training data domain. As the metal pieces are usually located in small sub-volumes of the reconstructions, we additionally define 17 three-dimensional regions of interest (ROIs) that are evaluated separately to get more expressive results on the sensitivity to the out-of-domain features. The coordinates of all 17 ROIs are provided in the supplementary material together with exemplary abdomen slices containing the respective ROIs to facilitate reproducibility. Additionally, we test our models on data from a separate domain, namely head CT scans (20 scans), to investigate prediction robustness on a different anatomy. Figure 3 shows example slices from the training and testing data sets with a highlighted abdomen ROI containing metal parts. Note that all scans are directly taken from the public data set without further modification such that they well represent clinical routine head and abdomen CT acquisitions of patients with and without metal implants³⁵.

Results

Quantitative results. We present quantitative denoising results on the entire abdomen test data set and only on the abdomen ROIs containing metal pieces in Table 1. Performance metrics for the investigated out-of-domain head CT data set are listed in Table 2. The three established image quality metrics root-mean-square error (RMSE), peak signal-to-noise ratio (PSNR), and structural similarity (SSIM)³⁶ are calculated to compare

	Full metal data			Metal ROIs		
	RMSE [HU] ↓	PSNR ↑	SSIM ↑	RMSE [HU] ↓	PSNR ↑	SSIM ↑
Low-dose CT	18.6 ± 6.1	41.51 ± 2.13	0.9512 ± 0.031	26.7 ± 15.9	38.98 ± 4.28	0.9594 ± 0.030
RED-CNN ¹²	12.0 ± 3.9	45.38 ± 1.93	0.9816 ± 0.012	24.6 ± 13.1	39.68 ± 3.75	0.9788 ± 0.015
RED+JBFs	12.7 ± 4.5	45.01 ± 2.02	0.9797 ± 0.016	22.1 ± 14.2	40.82 ± 4.27	0.9797 ± 0.013
QAE ¹³	44.7 ± 36.1	43.17 ± 4.49	0.9811 ± 0.014	837.8 ± 478.3	14.88 ± 10.92	0.6967 ± 0.232
QAE+JBFs	14.5 ± 5.4	44.61 ± 2.23	0.9827 ± 0.010	150.9 ± 87.5	26.99 ± 7.19	0.9075 ± 0.062

Table 1. Quantitative denoising results on the full abdomen scans containing metal implants as well as on 17 ROIs that contain metal parts. The metrics are averaged over all test patients or all test ROIs and provided as mean ± std. The respectively better performing network modification is highlighted in bold.

	Head CT data		
	RMSE [HU] ↓	PSNR ↑	SSIM ↑
Low-dose CT	5.3 ± 1.1	38.36 ± 1.58	0.9338 ± 0.017
RED-CNN ¹²	5.1 ± 0.6	38.42 ± 0.88	0.9646 ± 0.006
RED+JBFs	4.8 ± 0.6	39.01 ± 0.98	0.9662 ± 0.006
QAE ¹³	44.8 ± 40.9	37.30 ± 2.58	0.9679 ± 0.007
QAE+JBFs	9.8 ± 6.4	38.97 ± 1.77	0.9695 ± 0.007

Table 2. Quantitative denoising results on the full head CT data set. The metrics are averaged over all test patients and provided as mean ± std. The respectively better performing network modification is highlighted in bold.

model prediction with their respective high-dose target reconstruction. RMSE and PSNR particularly assess deviations from the target image intensities, whereas SSIM aims to imitate human perception to compare image content. We found that all performance differences between vanilla and respective JBF-based model in Table 1 are significant based on a Wilcoxon signed-rank test³⁷ on a p -value $p < 0.005$. The Wilcoxon signed-rank test is particularly suited to test the paired model predictions at hand without presuming an underlying statistical model.

Whereas the hybrid JBF layer-based pipelines perform comparably to the vanilla deep denoising models over the entire abdomen test data, an explicit performance improvement is recognized on the 17 abdomen ROIs as well as on the head CT scans on all three investigated image quality metrics. Both JBF-based pipelines decrease the RMSE by 10%/82% and improve the PSNR and SSIM by 3%/81% and 0.1%/30% around the out-of-training-domain metal features compared to the vanilla RED-CNN and QAE respectively. The denoising performance on the head CT data is improved by 6%/78% (RMSE), 2%/4% (PSNR), and 0.1%/0.1% (SSIM).

Qualitative results. Visual results on one ROI and a head CT slice are displayed in Fig. 4. Provided difference images between model prediction and high-dose target prediction particularly highlight disturbed features and erroneous predictions. Intensity distortions in close proximity to metal implants and in the skull region can be recognized for the RED-CNN, which get almost entirely removed using the RED-CNN prediction as guidance image in a JBF-based setting. Here, in particular the intensity shifts visible as shadows of the skull in the difference images of the vanilla model prediction are fully restored by the proposed hybrid JBF-based model.

The QAE predicts strong artifacts that are visible in the abdomen intensity images and difference images surrounding metal implants. Using such predictions as an image guide in a JBF-based pipeline produces results that visually look much closer to the high-dose target where features like the shape of metal pieces or the adjacent anatomy are visible. Further, intensity distortions in QAE predictions on the head CT data set are removed using the combined QAE+JBFs filtering approach. Only regions around the dental crowns with heavy metal reconstruction artifacts remain disturbed.

Discussion

Although one could simply add abdomen scans containing metal pieces or head CT data to the training data set to improve denoising performance, our experiment is particularly designed to evaluate and quantify robustness to real CT data that is underrepresented in the training data. Our experiment, therefore, mimics the present clinical scenario where a model is only trained on a limited number of scans but must also handle differing anatomies or scanning parameters. The denoising performance of a JBF depends on an optimal intensity range kernel G_{σ} , to avoid blurring edges. Here, the proposed pipeline can benefit from the guidance image that is predicted by a deep model that is capable of employing global image features to facilitate extracting sharp edges needed for the filter kernel computation. In case of prediction failures like in regions around metal implants or at the skull, the intensity range kernel contribution is either over- or underestimated. This results in over- or under-smoothing of the respective image region but is always based on the local content of the input image. Therefore, the intensity range kernel design of Gaussian shape prevents the output from large prediction errors by design.

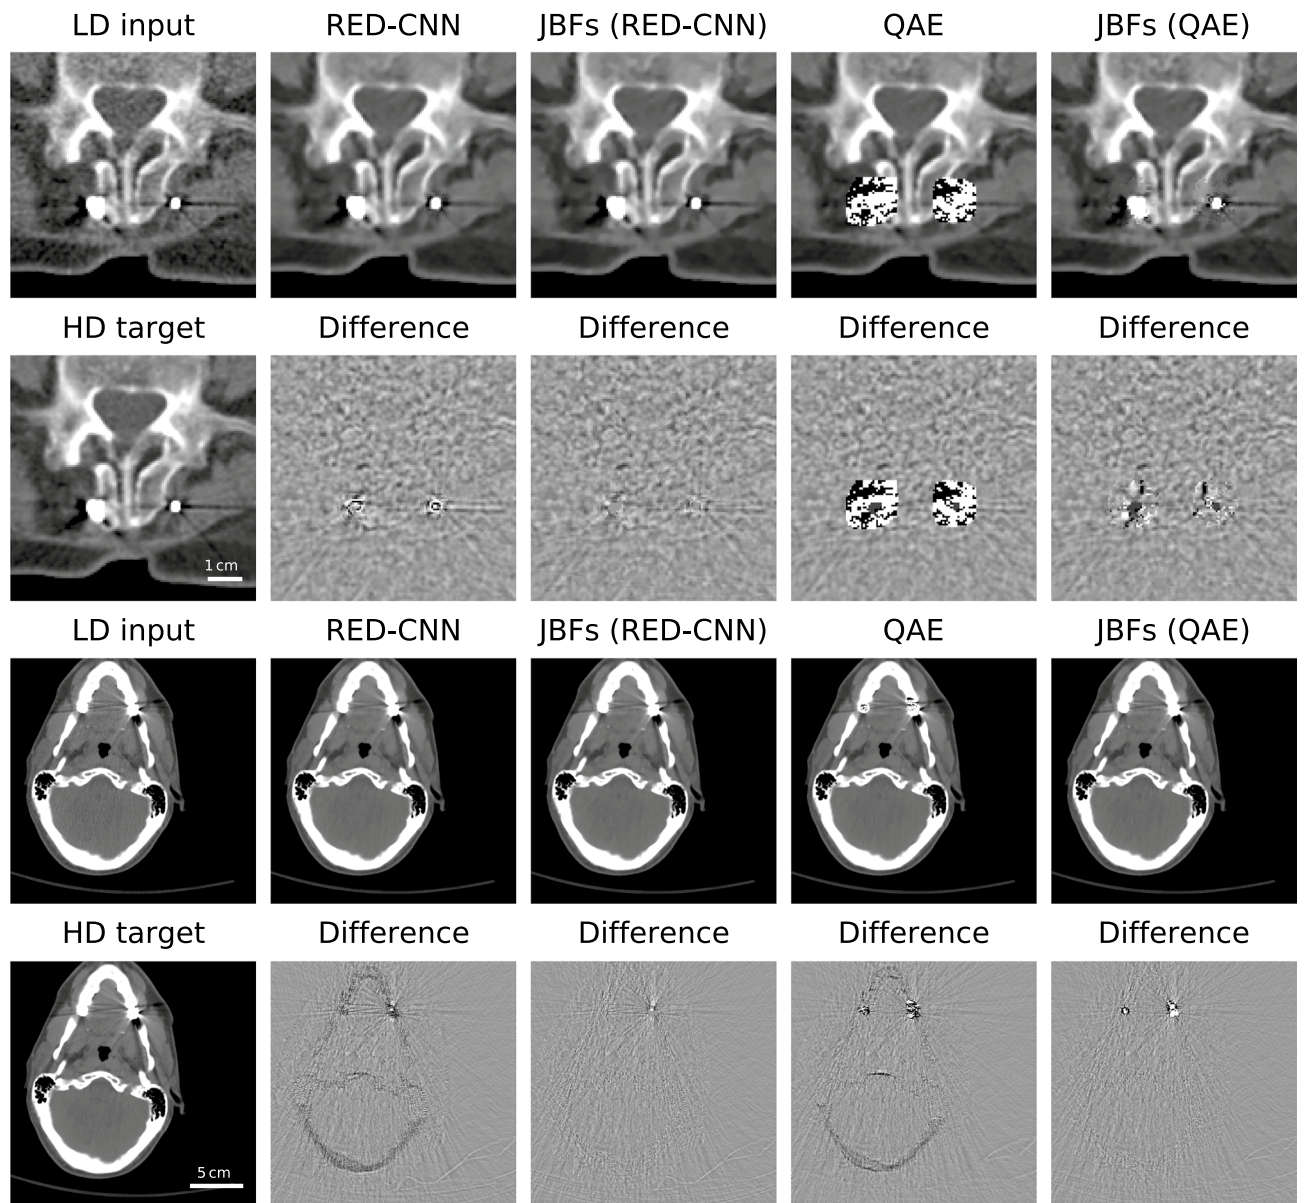


Figure 4. Qualitative denoising results on the ROI highlighted in Fig. 3 and on a head CT slice. Difference images are calculated between model prediction and high-dose (HD) target and are shown in the window $[-50, 50]$ HU for abdomen data and $[-100, 100]$ HU for head data. The reconstruction window is $[-150, 500]$ HU. Our hybrid models visually outperform the respective vanilla deep models.

In our conducted experiments, pre-trained denoising networks predict the guidance images that are input to the JBF layers. We performed additional experiments, training the JBFs together with the denoising networks in a combined end-to-end setting. Although this setting enhanced performance within the training data domain, we did not recognize explicit performance improvements in terms of robustness on the investigated out-of-training-domain data sets. Eventually, we did not design our experiments to answer the question how a guidance image that is optimized for JBFs is handled in the training data domain but we particularly want to investigate how JBFs handle the displayed artifacts predicted by the denoising networks as the primary goal of our study.

JBF-based pipelines almost entirely prevent the predictions from artifacts introduced by the DL-based models but the combined QAE+JBFs predictions still contain some slight distortions around the spine metal implant in Fig. 4. These results visualize that the JBF, although enforcing proximity to the noisy input, is still dependent on a reasonable guidance image. This dependence is desired as learned information from the guidance image should be employed during filtering. Our experiments show that mainly artifacts where image content is entirely removed in large areas of the guidance image are difficult to restore through JBFs. Please note that, the shown artifacts introduced by the QAE network can be regarded as worst-case in a clinical pipeline and are still satisfactorily handled by the JBFs considering the original, unfiltered QAE predictions.

DL frameworks like *PyTorch*³⁴ allow an automatic calculation of gradients in their operators. Therefore, one could think of implementing a JBF directly from *PyTorch* tensors instead of using analytical gradients to make its

parameters trainable. Although this is possible, training such a filter would require expensive Python loops over the training batches and kernel windows which would accumulate huge computational graphs for the gradient calculation. In practice, training such a model with reasonable image and batch sizes, therefore, is infeasible in terms of computational time and GPU memory. The analytical filter derivative presented in this work greatly simplifies the required computations to enable data-driven optimization and limit the computational overhead through adding JBF layers as shown by comparing run times with convolutional layers. Eventually, we believe that our open-source filter layer can be useful in further hybrid applications as a known denoising operator that can be optimized in a data-driven manner.

Conclusion

In this work, we presented a trainable JBF layer that can be incorporated into any deep model. We propose a hybrid denoising pipeline using these JBF layers and pre-trained deep denoising neural networks. The latter can produce faulty predictions when tested on data that is insufficiently represented in the training domain. In our experiments, we show that JBFs prevent DL-based models from severe prediction failures although the JBFs make use of distorted guidance images predicted from the neural networks. These results are explained by the clear algorithmic design of the JBF that limits the influence of the guidance image to the contribution of the intensity range filter kernel. We think that JBF layers can combine the flexibility of deep neural networks with the prediction reliability of conventional methods to leverage the power of deep models in clinical low-dose CT applications.

Data availability

The data sets analysed during the current study are publicly available in the *TCIA Low Dose CT Image and Projection Data* repository (Version 4)³⁵, <https://doi.org/10.7937/9NPB-2637>. Coordinates and exemplary slices of all analyzed abdomen ROIs are included in the supplementary material.

Code availability

The implementation of our open-source CUDA-accelerated trainable bilateral filter layer (*PyTorch*) together with example scripts and tests is publicly available at <https://github.com/faebstn96/trainable-joint-bilateral-filter-source>.

Received: 18 July 2022; Accepted: 17 October 2022

Published online: 20 October 2022

References

- Boone, J. M., Hendee, W. R., McNitt-Gray, M. F. & Seltzer, S. E. Radiation exposure from CT scans: How to close our knowledge gaps, monitor and safeguard exposure-proceedings and recommendations of the Radiation Dose Summit, sponsored by NIBIB, February 24–25, 2011. *Radiology* **265**, 544–554 (2012).
- Hess, E. P. *et al.* Trends in computed tomography utilization rates. *J. Patient Saf.* **10**, 52–58 (2014).
- Wagner, F. *et al.* Monte Carlo dose simulation for in-vivo X-ray nanoscopy. In *Bildverarbeitung für die Medizin* 107–112 (Springer, 2022).
- Huang, Y. *et al.* Semi-permeable filters for interior region of interest dose reduction in X-ray microscopy. In *Bildverarbeitung für die Medizin* 61–66 (Springer, 2021).
- Barrett, H. H., Gordon, S. & Hershel, R. Statistical limitations in transaxial tomography. *Comput. Biol. Med.* **6**, 307–323 (1976).
- Maier, A. & Fahrig, R. GPU denoising for computed tomography. *Graph.Process. Unit Based High Perform. Comput. Radiat. Ther.* **1**, 113–128 (2015).
- Dabov, K., Foi, A., Katkovnik, V. & Egiazarian, K. Image denoising with block-matching and 3D filtering. In *Image Processing: Algorithms and Systems, Neural Networks, and Machine Learning*, vol. 6064, 606414 (International Society for Optics and Photonics, 2006).
- Giraldo, J. C. R. *et al.* Comparative study of two image space noise reduction methods for computed tomography: Bilateral filter and nonlocal means. In *2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society* 3529–3532 (IEEE, 2009).
- Tomasi, C. & Manduchi, R. Bilateral filtering for gray and color images. In *Sixth International Conference on Computer Vision* 839–846 (IEEE, 1998).
- Zhao, T., Hoffman, J., McNitt-Gray, M. & Ruan, D. Ultra-low-dose CT image denoising using modified BM3D scheme tailored to data statistics. *Med. Phys.* **46**, 190–198 (2019).
- Maier, A. *et al.* Three-dimensional anisotropic adaptive filtering of projection data for noise reduction in cone beam CT. *Med. Phys.* **38**, 5896–5909 (2011).
- Chen, H. *et al.* Low-dose CT with a residual encoder–decoder convolutional neural network. *IEEE Trans. Med. Imaging* **36**, 2524–2535. <https://doi.org/10.1109/TMI.2017.2715284> (2017).
- Fan, F. *et al.* Quadratic autoencoder (Q-AE) for low-dose CT denoising. *IEEE Trans. Med. Imaging* **39**, 2035–2050. <https://doi.org/10.1109/TMI.2019.2963248> (2019).
- Wu, D., Kim, K. & Li, Q. Low-dose CT reconstruction with Noise2Noise network and testing-time fine-tuning. *Med. Phys.* **48**, 7657–7672 (2021).
- Gu, J. & Ye, J. C. AdaIN-based tunable CycleGAN for efficient unsupervised low-dose CT denoising. *IEEE Trans. Comput. Imaging* **7**, 73–85 (2021).
- Li, M., Hsu, W., Xie, X., Cong, J. & Gao, W. SACNN: Self-attention convolutional neural network for low-dose CT denoising with self-supervised perceptual loss network. *IEEE Trans. Med. Imaging* **39**, 2289–2301 (2020).
- Patwari, M., Gutjahr, R., Raupach, R. & Maier, A. Low dose CT denoising via joint bilateral filtering and intelligent parameter optimization. In *Sixth International Conference on Image Formation in X-Ray Computed Tomography* 174–177 (2020).
- Antun, V., Renna, F., Poon, C., Adcock, B. & Hansen, A. C. On instabilities of deep learning in image reconstruction and the potential costs of AI. *Proc. Natl. Acad. Sci.* **117**, 30088–30095 (2020).
- Hirano, H., Minagi, A. & Takemoto, K. Universal adversarial attacks on deep neural networks for medical image classification. *BMC Med. Imaging* **21**, 1–13 (2021).

20. Maier, A. *et al.* Precision learning: Towards use of known operators in neural networks. In *2018 24th International Conference on Pattern Recognition* 183–188 (IEEE, 2018).
21. Maier, A. *et al.* Learning with known operators reduces maximum error bounds. *Nat. Mach. Intell.* **1**, 373–380. <https://doi.org/10.1038/s42256-019-0077-5> (2019).
22. Thies, M. *et al.* Calibration by differentiation—Self-supervised calibration for X-ray microscopy using a differentiable cone-beam reconstruction operator. *J. Microsc.* **287**, 81–92 (2022).
23. Patwari, M., Gutjahr, R., Raupach, R. & Maier, A. JBFnet—Low dose CT denoising by trainable joint bilateral filtering. In *International Conference on Medical Image Computing and Computer-Assisted Intervention—MICCAI 2020* 506–515 (Springer, 2020).
24. Wu, H., Zheng, S., Zhang, J. & Huang, K. Fast end-to-end trainable guided filter. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 1838–1847 (2018).
25. Wu, Q., Tang, H., Liu, H. & Chen, Y. Masked joint bilateral filtering via deep image prior for digital X-ray image denoising. *IEEE J. Biomed. Health Inform.* **26**, 4008–4019 (2022).
26. Gadde, R., Jampani, V., Kiefel, M., Kappler, D. & Gehler, P. V. Superpixel convolutional networks using bilateral inceptions. In *European Conference on Computer Vision* 597–613 (Springer, 2016).
27. Patwari, M., Gutjahr, R., Raupach, R. & Maier, A. Limited parameter denoising for low-dose X-ray computed tomography using deep reinforcement learning. *Med. Phys.* **49**, 4540–4553 (2022).
28. Xu, J. & Noo, F. Efficient gradient computation for optimization of hyperparameters. *Phys. Med. Biol.* **67**, 03NT01 (2022).
29. Wagner, F. *et al.* Ultralow-parameter denoising: Trainable bilateral filter layers in computed tomography. *Med. Phys.* <https://doi.org/10.1002/mp.15718> (2022).
30. Petschnigg, G. *et al.* Digital photography with flash and no-flash image pairs. *ACM Trans. Graph. (TOG)* **23**, 664–672. <https://doi.org/10.1145/1015706.1015777> (2004).
31. Bera, S. & Biswas, P. K. Noise conscious training of non local neural network powered by self attentive spectral normalized Markovian patch GAN for low dose CT denoising. *IEEE Trans. Med. Imaging* **40**, 3663–3673 (2021).
32. Huang, Z. *et al.* DaNet: Dose-aware network embedded with dose-level estimation for low-dose CT imaging. *Phys. Med. Biol.* **66**, 015005 (2021).
33. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
34. Paszke, A. *et al.* PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems* Vol. 32 (Curran Associates, Inc., 2019).
35. Moen, T. R. *et al.* Low-dose CT image and projection dataset. *Med. Phys.* **48**, 902–911. <https://doi.org/10.1002/mp.14594> (2021).
36. Wang, Z., Bovik, A. C., Sheikh, H. R. & Simoncelli, E. P. Image quality assessment: From error visibility to structural similarity. *IEEE Trans. Image Process.* **13**, 600–612 (2004).
37. Wilcoxon, F. Individual comparisons by ranking methods. In *Breakthroughs in Statistics* 196–202 (Springer, 1992).

Acknowledgements

The research leading to these results has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation program (ERC Grant No. 810316). Further, we thank the NVIDIA Corporation for their GPU donation through the NVIDIA Hardware Grant Program. Open Access funding enabled and organized by Projekt DEAL.

Author contributions

F.W. conceived and conducted the experiments and the algorithm derivation. M.T., F.D., and M.G. contributed on the filter algorithm and experimental design. M.P. provided network implementations. S.P. assisted with the CUDA kernels. N.M., L.P., and Y.H. provided valuable technical feedback during development. A.M. supervised the project. All authors reviewed the manuscript.

Competing interests

F.D., N.M., and L.P. are employees of Siemens Healthcare GmbH. All other authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-022-22530-4>.

Correspondence and requests for materials should be addressed to F.W.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022