Received: 7 September 2018 | Revised: 8 January 2019 | Accepted: 9 January 2019

DOI: 10.1002/mrm.27677

FULL PAPER

Magnetic Resonance in Medicine

Development and clinical implementation of SeedNet: A sliding-window convolutional neural network for radioactive seed identification in MRI-assisted radiosurgery (MARS)

Jeremiah W. Sanders[1,2] | **Steven J. Frank**[3] | **Rajat J. Kudchadker**[2,4] | **Teresa L. Bruno**[3] | **Jingfei Ma**[1,2]

[1]Department of Imaging Physics, University of Texas MD Anderson Cancer Center, Houston, Texas

[2]Medical Physics Graduate Program, University of Texas MD Anderson Cancer Center UTHealth Graduate School of Biomedical Sciences, Houston, Texas

[3]Department of Radiation Oncology, University of Texas MD Anderson Cancer Center, Houston, Texas

[4]Department of Radiation Physics, University of Texas MD Anderson Cancer Center, Houston, Texas

Correspondence
Jeremiah Sanders, Department of Imaging Physics, Unit 1472, The University of Texas MD Anderson Cancer Center, 1515 Holcombe Blvd., Houston, TX 77030.
Email: jsanders1@mdanderson.org.

Purpose: To develop and evaluate a sliding-window convolutional neural network (CNN) for radioactive seed identification in MRI of the prostate after permanent implant brachytherapy.

Methods: Sixty-eight patients underwent prostate cancer low-dose-rate (LDR) brachytherapy using radioactive seeds stranded with positive contrast MR-signal seed markers and were scanned using a balanced steady-state free precession pulse sequence with and without an endorectal coil (ERC). A sliding-window CNN algorithm (SeedNet) was developed to scan the prostate images using 3D sub-windows and to identify the implanted radioactive seeds. The algorithm was trained on sub-windows extracted from 18 patient images. Seed detection performance was evaluated by computing precision, recall, F_1-score, false discovery rate, and false–negative rate. Seed localization performance was evaluated by computing the RMS error (RMSE) between the manually identified and algorithm-inferred seed locations. SeedNet was implemented into a clinical software package and evaluated on sub-windows extracted from 40 test patients.

Results: SeedNet achieved 97.6 ± 2.2% recall and 97.2 ± 1.9% precision for radioactive seed detection and 0.19 ± 0.04 mm RMSE for seed localization in the images acquired with an ERC. Without the ERC, the recall remained high, but the false–positive rate increased; the RMSE of the seed locations increased marginally. The clinical integration of SeedNet slightly increased the run-time, but the overall run-time was still low.

Conclusion: SeedNet can be used to perform automated radioactive seed identification in prostate MRI after LDR brachytherapy. Image quality improvement through pulse sequence optimization is expected to improve SeedNet's performance when imaging without an ERC.

wileyonlinelibrary.com/journal/mrm

KEYWORDS
convolutional neural network (CNN), magnetic resonance imaging (MRI), prostate brachytherapy, radioactive seed

1 | INTRODUCTION

Radioactive seed localization after low-dose-rate (LDR) brachytherapy is essential to determine the radiation dose to the prostate and surrounding critical normal structures. X-ray CT is most commonly used for post-implant imaging because of the high contrast of the radio-opaque seeds,[1,2] and several algorithms for automatically localizing the seeds in CT images are currently available.[3-5] However, 1 major limitation of CT is its low soft-tissue contrast, which makes it difficult to delineate and contour the prostate and surrounding organs at risk. Additionally, the presence of the metallic seeds produces severe artifacts in the CT images, which increases the complexity of delineating the prostate boundary. In contrast, MRI provides high soft-tissue contrast[6] and may serve as an alternative to CT for post-implant imaging.

Two general approaches for using MRI to localize radioactive seeds after LDR brachytherapy have been reported. The first approach relies on co-registration of separate CT and MR images for seed identification and anatomy delineation, respectively.[7-11] The second approach is based on MRI only, using either a single pulse sequence or multiple pulse sequences for different image contrasts.[12-15] However, because the radioactive seeds do not produce MR signals, conclusively identifying their locations in MR images can be challenging, especially in the presence of other similar image features, such as needle tracks. A positive MRI contrast seed marker technology was recently developed to overcome this limitation[16,17]; with this approach, an MRI signal-producing solution (C4) is encapsulated and placed as functional spacers between adjacent radioactive seeds. This approach provides an accurate means of indirectly localizing the radioactive seeds through localization of the seed markers.

Imaging both the seed markers and anatomy previously required co-registration of multiple MR images with different contrasts.[18] Researchers recently demonstrated that the seed markers and anatomical structures could be imaged in a single acquisition using a balanced SSFP pulse sequence and an endorectal coil (ERC).[19] Although balanced SSFP imaging provides an accurate means for seed localization, certified medical dosimetrists (CMDs) must manually identify each seed in the MR images, which is time consuming and may require co-registration or cross-referencing with a post-implant CT acquisition to identify inconspicuous seeds near anatomic boundaries. In addition, although it provides intrinsically superior SNR, the use of an ERC increases study cost, reduces patient comfort, and decreases clinical throughput. Because of these limitations, the use of ERCs is largely limited to major academic medical centers. Methods to identify the radioactive seeds in MRI without an ERC should help expand the adoption of MRI for post-implant assessment.

A couple of automated algorithms for radioactive seed localization in MRI have been proposed. The first method uses morphological operations and Laplacian of a Gaussian feature detection to detect the signal voids produced by the implanted radioactive seeds.[20] This algorithm has only been applied to simplistic phantoms where the seed locations are easily identified. The second method also uses morphological operations along with heuristically chosen image filters to identify the radioactive seeds.[21] The average run-time of the algorithm is 48 min and only 89% of the seeds could be identified on average. The specific mechanics of both algorithms are very scarce and neither appear to be applicable in a routine clinical setting.

In theory, radioactive seed identification in post-implant imaging can be posed as a computer vision task. Convolutional neural networks (CNNs), recently re-popularized by the success of AlexNet,[22] have been applied successfully in several computer vision tasks, including object detection, classification, and localization. A common approach for object recognition is to train a CNN to predict the bounding box that surrounds a particular class of object in an image.[23-27] This is an appropriate choice when the size and pose of the object varies across images, such as in natural images. However, when the object of interest is of a fixed size and remains consistent across different images, it may not be appropriate to predict a bounding box for the object of interest, especially when there are multiple objects of the same type in close proximity to each other. In the targeted application of prostate brachytherapy, multiple seed strands may be implanted in proximity, resulting in closely spaced or overlapping signal voids from multiple seeds on the image. Single-shot bounding box detectors may identify these instances as a single seed, reducing the recall. In contrast, a sliding-window with a fixed sub-window size could be used to predict the class of the object that is centered within a given sub-window of the image. The location of the object could also be predicted within the sub-window and mapped back to the original image to predict the global location of the seed.

In this work, we present the development and clinical integration of a sliding-window CNN with a fixed sub-window size for identifying radioactive seeds in MR images of the prostate after LDR brachytherapy. With our approach, we effectively repose the task of radioactive seed identification as an image classification task at the sub-window level. The algorithm, SeedNet, demonstrates high precision and recall of the radioactive seeds, fast inference, and ease of integration via general purpose graphics processing units (GPUs), which may be sufficient for use in a clinical setting. SeedNet performs the seed identification in a fraction of the time required by CMDs, with or without the use of an ERC for imaging.

2 | METHODS

2.1 | Images

This study was performed under a retrospective Institutional Review Board-approved protocol. Sixty-eight patients were implanted with stranded radioactive seeds placed next to positive contrast MR-signal seed markers (Sirius, C4 Imaging, Houston, TX). All MRI scans were performed on the day of the seed implantation. The patients were scanned on a 1.5T Siemens Magnetom Aera scanner (Siemens Healthineers, Erlangen, Germany). Each scan was performed with 2 18-channel external pelvic arrays in combination with a 2-channel rigid ERC (Invivo, Gainesville, FL). Localizer images were used to guide the placement of the ERC and ensure its proper position relative to the prostate in both the superior–inferior and left–right directions. After the ERC was accurately positioned, patients were injected with glucagon to reduce peristaltic motion during image acquisition. The patients were scanned with a 3D constructive interference in steady-state sequence that provided a mix of T_1 and T_2 contrast, which permits visualization of both the seed markers and anatomy in a single sequence. Depending on the prostate size, 80–96 slices were acquired to ensure adequate coverage in the superior–inferior direction. Typical scan parameters were: TR/TE = 5.29/2.31 ms; readout bandwidth = 560 Hz/pixel; field of view = 15 × 15 cm; voxel dimensions = 0.59 × 0.59 × 1.20 mm (interpolated to 0.29 × 0.29 × 1.20 mm); flip angle = 52°; total scan time = 4–5 min.

After imaging with the ERC was completed, the ERC was removed in 48 of the 68 patients, and the patients were re-positioned and scanned again with the same 2 18-channel external pelvic array coils. Slight adjustments to the protocol were made to compensate for some of the SNR loss resulting from the removal of the ERC. The adjustments included increasing the voxel size and doubling the number of signal averages from 1 to 2. Additional information about the imaging protocol and scan parameters can be found in our previous work.[28]

2.2 | Image label creation

A CMD (T.L.B.) manually identified the radioactive seeds in the MR images using a clinical software package (MIM Software, Cleveland, OH). Positive identification of the radioactive seeds was facilitated through localization of the positive contrast MR-signal seed markers. The centroid coordinates of the radioactive seeds in each patient were exported to a text file to create the labels for the sub-windows. The seed coordinates were mapped to a voxel grid of the same size as the image to create a mask of the seed locations within the image grid (512 × 512 × number of slices). In the mask, background voxels were set to 0, seed voxels were set to 1, and seed marker voxels were set to 2. This mask was used to create the labels for the individual sub-windows, described below.

Because imaging without an ERC has been shown to significantly reduce CMDs' ability to manually identify seeds,[29] seed identification in the images acquired without the ERC was facilitated by localization of the seed markers that were visible and by fusing the MRI with a post-implant CT.

2.3 | Sliding window algorithm

For limiting the image context to the location of the implanted seeds, a rectangular cuboid region of interest (ROI) surrounding the prostate was cropped from the 3D image. The ROI can either be drawn manually, which was the method performed in this study, or automatically defined with one of a number of automated prostate segmentation algorithms for MRI.[30-36] A sliding-window algorithm was written to scan the entire ROI in 3D sub-windows with a user-defined stride. For a given ROI of size $i_x \times i_y \times i_z$, sub-window size of $w_x \times w_y \times w_z$, and stride of $s_x \times s_y \times s_z$, the number of sub-window translations in each orthogonal direction was computed as

$$\begin{aligned} n_x &= \left\lceil \frac{i_x - w_x + s_x}{s_x} \right\rceil, \\ n_y &= \left\lceil \frac{i_y - w_y + s_y}{s_y} \right\rceil, \\ \text{and } n_z &= \left\lceil \frac{i_z - w_z + s_z}{s_z} \right\rceil, \end{aligned} \tag{1}$$

where $\lceil \cdot \rceil$ indicates the ceiling operation. The total number of sub-windows analyzed for a given patient is the product of n_x, n_y, and n_z, $n_{total} = n_x n_y n_z$. As such, increasing both the size of the sub-window and stride of the sliding-window reduces the number of sub-windows to be processed per patient, which reduces the total computation time. However, increasing the stride of the sliding-window increases the probability of missing a seed if it is selected too large. Moreover, increasing the sub-window size may increase the number of other seeds or seed markers captured in the window, potentially reducing the ability of the model to discriminate among seeds, seed markers, and the background anatomy. The optimal sub-window dimensions

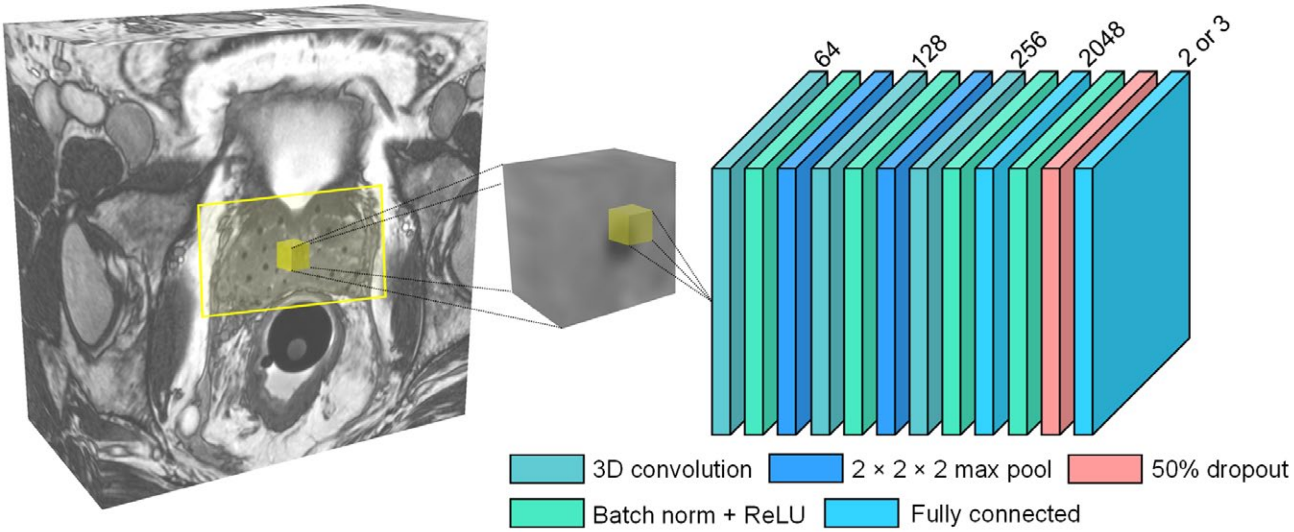

FIGURE 1 The SeedNet architecture. The windows were processed with 3D CNNs. Three separate CNNs, each with the same configuration of layers, were trained to perform seed detection, classification, and localization tasks

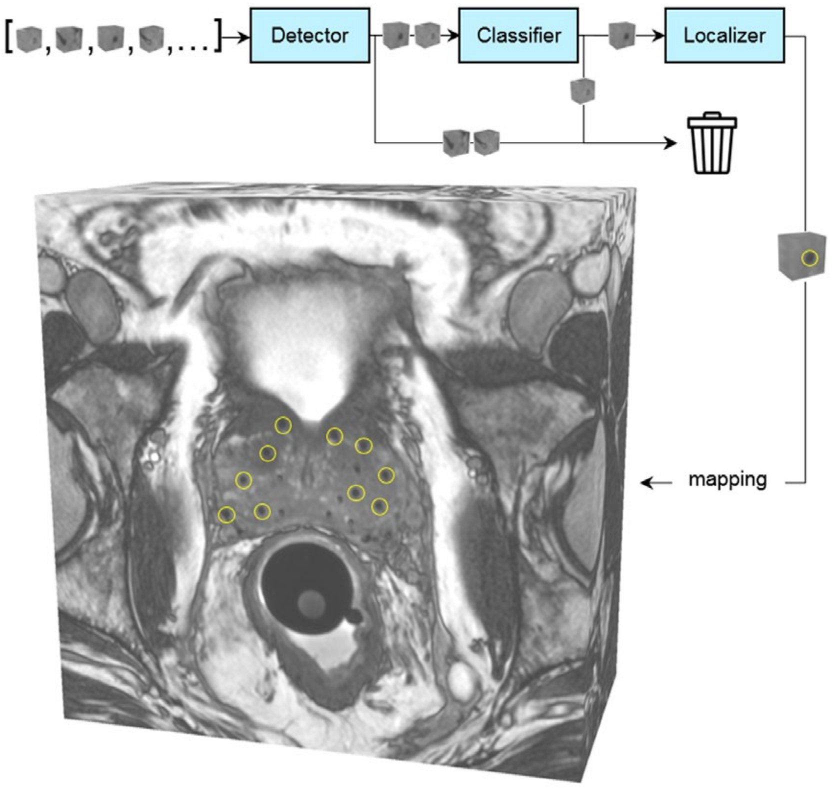

FIGURE 2 Sub-window processing at inference time. First, each sub-window was passed to the detector to identify windows containing seeds. The seed sub-windows were then passed to a classifier to reject seed marker sub-windows that, owing to their similar shape and proximity to the seeds, were incorrectly classified as seed sub-windows. The seed sub-windows from the classifier were passed to the localizer to pinpoint the precise location of the seed within the sub-window. Finally, the seed sub-windows were mapped back to their locations within the original image stack

were experimentally determined by evaluating the precision and recall for multiple values of w_x, w_y, and w_z, described below.

The sliding-window algorithm uses a *for* loop to loop over all combinations of the sub-window coordinates based on the user-defined stride and sub-window dimensions. The voxel location of a designated corner of each sub-window is stored to map the identified seed locations in the sub-windows back to the original image stack. Each iteration of the *for* loop is an independent process and may be parallelized over the individual cores of multi-core CPU architectures for reduced computation time.

2.4 | Network architecture

Model compactness and computational efficiency were major considerations in the design of the model to enable its deployment on computer architectures with relatively low computing power (such as those used by CMDs). As such, the network was constructed by starting with a single convolutional layer, a fully connected layer, and an output layer, and then the number of convolutional layers was increased until the accuracy on the cross-validation sub-windows plateaued at a maximum value. The resulting architecture is shown in Figure 1. The model consists of 3 3D convolutional layers with 3 × 3 × 3 convolution kernels. Maximum pooling (2 × 2 × 2, non-overlapping) followed by batch normalization and rectified linear unit activations were used between the convolutional layers. The third convolutional layer was connected to a fully connected layer with 2048 neurons. The output layer was a fully connected layer with 2 or 3 neurons, depending on the inference task. We note that this architecture is similar to the one developed by LeCun et al.[37] for classifying handwritten digits. The main differences, aside from the fact that SeedNet operates on 3D images, are that we have an additional convolutional layer before the fully connected layer and a different output size.

The sub-window processing is summarized in Figure 2. Each sub-window was first analyzed by a detector to identify the windows containing radioactive seeds. Because of their similar shape and proximity, some sub-windows containing seed markers could be misclassified as seed sub-windows by the detector. This is because, in a few instances, the seed markers may not produce their maximum MR-signal and may be misinterpreted as seeds. Therefore, the seed sub-windows from the detector were passed to a classifier to reject potential false–positives. Afterward, the seed sub-windows from the classifier were passed to a localizer to identify the precise locations of the seeds within the sub-windows. The sub-windows with the identified seed locations were then mapped back to the original image stack. The detector, classifier, and localizer were all constructed with the same configuration of layers as shown in Figure 1.

Two neurons (background, seed) were used in the output layer for the detection model. The classification model had 3 neurons (background, seed, seed marker) in the output layer. The localization model had 3 neurons (1 for each of the x, y, and z coordinates) in the output layer to estimate the seed location within the sub-window.

Two fully convolutional variants of the architecture shown in Figure 1 were tested. The first replaced the dense fully connected layer with a 3D convolutional layer using a 3 × 3 × 1 kernel and no zero padding. The second replaced the dense fully connected layer with a 3 × 3 × 1 average pooling layer followed by a 3D convolutional layer using a 1 × 1 × 1 kernel. The dense fully connected layer resulted in the lowest loss on the validation data set, being 2 orders of magnitude and 1 order of magnitude lower than the fully convolutional approaches, respectively.

2.5 | Class imbalance

For a given patient, the number of sub-windows to analyze is on the order of 10^4–10^5, depending on the size of the ROI encompassing the prostate. For example, consider a patient for whom $i_x = 180$, $i_y = 145$, $i_z = 36$, $w_x = 13$, $w_y = 13$, $w_z = 7$, $s_x = 2$, $s_y = 2$, and $s_z = 1$; from Equation (1), $n_{total} = 170{,}850$ sub-windows. A majority of these sub-windows contain the background anatomy, whereas only a fraction of them contain seeds. The exact number of sub-windows containing seeds depends on the number of seeds implanted and the stride of the sliding-window. Nevertheless, a large class imbalance exists between the sub-windows containing seeds and sub-windows containing background anatomy.

To simulate the class imbalance encountered during inference, the detection and classification models were trained with imbalanced data by randomly undersampling the majority class (background) and compensated by weighting the cross-entropy loss function.[38] We tested majority-to-minority class ratios of 10, 25, 50, and 100 and found that they all yielded the same performance. Therefore, we chose a majority-to-minority class ratio of 10 because it yielded identical performance with fewer training sub-windows. The number of background sub-windows, n_{bg}, in the detection training data was included at 10 times the number of seed sub-windows, n_{seed}. For the classification model, the number of seed sub-windows, n_{seed}, and number of marker sub-windows, n_{marker}, were compared, and the minimum, $n_{min} = \min\left(n_{seed}, n_{marker}\right)$, was computed. The number of seed and seed marker sub-windows included in the classification training data were each included at n_{min}, and the background windows were included at 10 times n_{min}. The class labels for the seeds and seed markers were weighted with a value of 10 in the cross-entropy loss functions.

2.6 | Data augmentation

Each sub-window in the training data set was processed with rotation affine transformations at 90°, 180°, and 270° in the axial (x–y) plane. Rotating the windows 180° about the x–z plane did not cause a reduction in the validation loss.

2.7 | Training

Eighteen of the 68 patients were randomly selected for training. Sub-windows from these patients were extracted and used to train the model. Training was conducted with an Adam optimizer.[39] The learning rate was set to 3×10^{-5}. The sub-windows were shuffled before training, and 20% of the sub-windows were reserved for cross-validation. Decaying the learning rate did not demonstrate an advantage in improving accuracy on the cross-validation sub-windows. A dropout rate of 50% was used.[40] Weights were initialized using Xavier's method.[41] The model was trained for 100 epochs using a mini batch size of 128, and checkpoints were saved after each epoch. The model achieving the lowest loss on the cross-validation sub-windows was selected for testing. Cross-entropy was used as the loss function for the detection and classification CNNs. Mean squared error was used as the loss function for the localization CNN. The models were written in Python using TensorFlow (Google, Mountain View, CA) and Keras.[42] The model was trained on a Dell 7920 computing server that operated under the Linux RedHat v7.2 operating system and contained 4 NVIDIA K2200 GPUs connected with SLI technology.

2.8 | Determining optimal sub-window dimensions

The context from which the model will make inferences is defined by the sub-window dimensions, w_x, w_y, and w_z. Increasing the sub-window dimensions gives the model more information with which to make inferences. However, because the seeds and seed markers are in proximity to one another, making the sub-window dimensions too large would capture multiple seeds and seed markers in the same window, potentially reducing the model's ability to discriminate between the classes. Similarly, making the sub-window dimensions too small would not capture the entire seed or marker in the sub-window. For evaluation of the impact of the sub-window size, models were trained with different values of $w_x \times w_y \times w_z$ and evaluated based on precision and recall. Ten of the 68 patients (different from those used for training) were selected at random and used for the evaluation. The models that achieved the highest precision and recall on these 10 patients were selected for testing on the remaining patients.

An estimate of the minimum sub-window size that fully encapsulates a seed can be computed with knowledge of the voxel dimensions of the MRI and physical dimensions of the radioactive seeds. Let the voxel dimensions be defined as Δ_x, Δ_y, and Δ_z and the diameter and length of the cylindrical seeds be defined as d_{seed} and l_{seed}, respectively. Assuming isotropic resolution in the transverse (x–y) plane, $\Delta_x = \Delta_y = \Delta_{xy}$, and the seeds would be approximately $p_{xy}^{seed} = \lceil \frac{d_{seed}}{\Delta_{xy}} \rceil$ pixels in diameter. In the z-dimension, the seeds would be approximately $p_z^{seed} = \lceil \frac{l_{seed}}{\Delta_z} \rceil$ pixels in length. The susceptibility artifacts around the seeds cause the signal voids to bloom slightly larger,[43] which increases p_{xy}^{seed} and p_z^{seed} by ~1–3 pixels depending on the seed orientation.

The patients in this study were implanted with common brachytherapy seeds, including palladium-103 seeds (Theragenics, Theraseed, Buford, GA), cesium-131 seeds (IsoRay, Proxelan, Richland, WA), and iodine-125 seeds (IsoAid, Advantage I-125, Port Richey, FL). The overall length and diameter of these seeds were approximately 4.5 mm and 0.8 mm, respectively. The seed markers were slightly smaller than the seeds. Additional pixels were added to the minimum sub-window size to avoid cases in which a portion of the seed or marker is cut off along the edges, as discussed below. With the above considerations, we investigated sub-window sizes of $11 \times 11 \times 5$, $11 \times 11 \times 7$, $11 \times 11 \times 9$, $13 \times 13 \times 5$, $13 \times 13 \times 7$, $13 \times 13 \times 9$, $15 \times 15 \times 5$, $15 \times 15 \times 7$, and $15 \times 15 \times 9$. A sliding-window stride of $2 \times 2 \times 1$ was used for all sub-window sizes. We did not use a stride >1 in the slice direction because it would cause the sliding-window to skip slices and potentially miss seeds.

2.9 | Sub-window labels

The seed and seed marker locations in the mask were cropped at the grid location where the prostate ROI was extracted from the image. The sliding-window algorithm was processed on the masks of the seed and marker centroid locations to create the labels for the sub-windows. Cases in which the seed or marker was partially cut off on the edge of the sub-window were avoided by requiring the centroids of the seeds and seed markers to be in the central $3 \times 3 \times 1$ portion of the sub-window to be classified as a seed or marker; otherwise, they were classified as background (Figure 3). This ensures that only seeds and seed markers that are fully encapsulated within the sub-windows are appropriately labeled.

Three sets of labels for each sub-window were created. The first set of labels was created to train the detection model. Sub-windows containing background were given a class label of 0, and sub-windows containing seeds were given a class label of 1. The second set of labels was created to train the classification model. Background and seed sub-windows were given class labels of 0 and 1, respectively, and seed marker sub-windows were given a class label of 2. Finally, a third set of labels was created to train the localization model. The voxel coordinates (x, y, and z) of the seed and seed marker centroids within the sub-window were used as the labels.

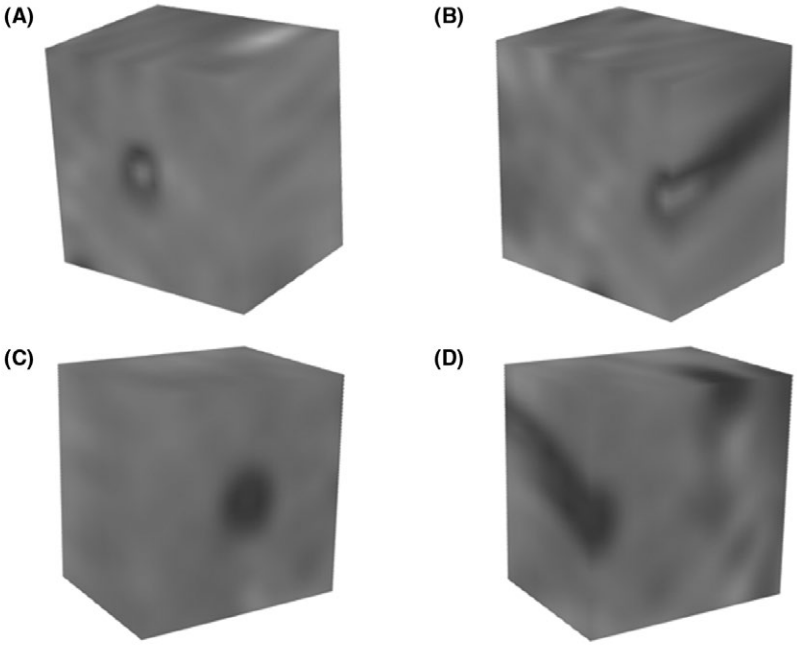

FIGURE 3 Example sub-windows of seeds and seed markers. (A) Sub-window of a hyperintense seed marker. (B) The seed marker lies on the edge of the window, and a portion of the seed marker is excluded. This window would be labeled as a background sub-window during training. (C) Sub-window of a hypointense seed. (D) The seed lies on the edge of the window. This window would also be labeled as a background sub-window during training

2.10 | Clinical implementation

SeedNet was integrated into a clinical software package used at our institution (MIM Software, Cleveland, OH) with the assistance of a deep learning application engine (DLAE) that we recently developed[44] (Figure 4). We used MIM's software development kit (SDK) to interface with the inference mechanics of DLAE. To demonstrate the integration, a prostate contour performed by a board-certified radiation oncologist (S.J.F.) and a patient's image volume were sent from MIM to a MATLAB-based SDK. In the SDK, a bounding box was extracted from the prostate contour and used to crop the prostate into a rectangular cuboid ROI. We note that variability in the prostate segmentation may produce slight variations in the ROI being scanned by the sliding-window. Although we did not investigate these slight variations in the present study, they are not expected to significantly impact the seed identification performance. The sliding-window algorithm was used to extract sub-windows from the ROI using a sub-window size of 13 × 13 × 7 and stride of 2 × 2 × 1. The sub-windows were sent into DLAE, which loaded the SeedNet model and made inferences on the sub-windows. The seed location inferences were mapped back to the original image volume and were sent to MIM using the SDK.

2.11 | Performance evaluation

SeedNet's seed detection performance was characterized by computing the precision, recall, F_1-score, false discovery rate, and false–negative rate of the detections. The localization performance was characterized by computing the RMS error (RMSE) between the inferred seed locations and the locations identified by the CMD.

3 | RESULTS

All model inferences were performed on a Dell Precision Tower 7910 that operated under the Windows 7 operating system and equipped with a single NVIDIA Quadro M2000 GPU. The performance metrics below are reported as the average ± 1 SD of the inferences for the cohort of test patients.

A sub-window 13 × 13 × 7 in size achieved the highest combined precision and recall on the 10 validation patients, which is consistent with our preliminary studies.[45-47] The precision and recall for this sub-window size were 97.2 ± 1.7% and 98.5 ± 1.2%, respectively. All test results reported below are based on a 13 × 13 × 7 sub-window size.

Representative examples of SeedNet inferences for 4 patients imaged with an ERC are shown in Figure 5. The precision and recall for the 40 test patients are plotted in Figure 6. The F1-scores, positive prediction summary, false discovery rates, and false–negative rates are shown in Figure 7 (left column). For the 40 test patients, the precision was 97.2 ± 1.9%, the recall was 97.6 ± 2.2%, and the inference time was 56.6 ± 18.2 s. The F_1-score was 97.4 ± 1.5%, the false discovery rate was 2.8 ± 1.9%, and the false–negative rate was 2.4 ± 2.2%. The RMSE of the seed locations was 0.19 ± 0.04 mm. The false–positives, which reduced the precision, primarily occurred in cylindrical dark crevices between anatomic boundaries, with the majority being near the apex of the prostate. No false–positives were observed in hemorrhaging areas or in regions where prior needle biopsies were performed. The false–negatives, which reduced the recall, occurred for seeds whose orientation deviated significantly from the rest of the seed population.

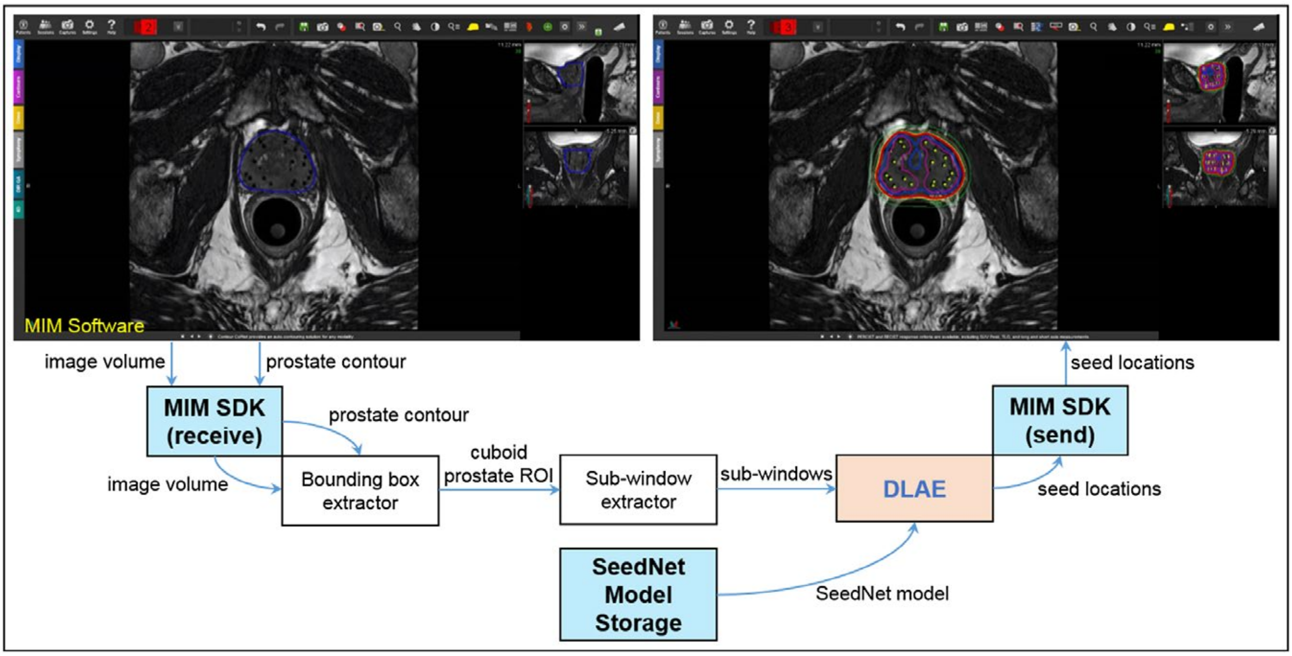

FIGURE 4 Data-flow diagram demonstrating the integration of SeedNet into a clinical software package using DLAE

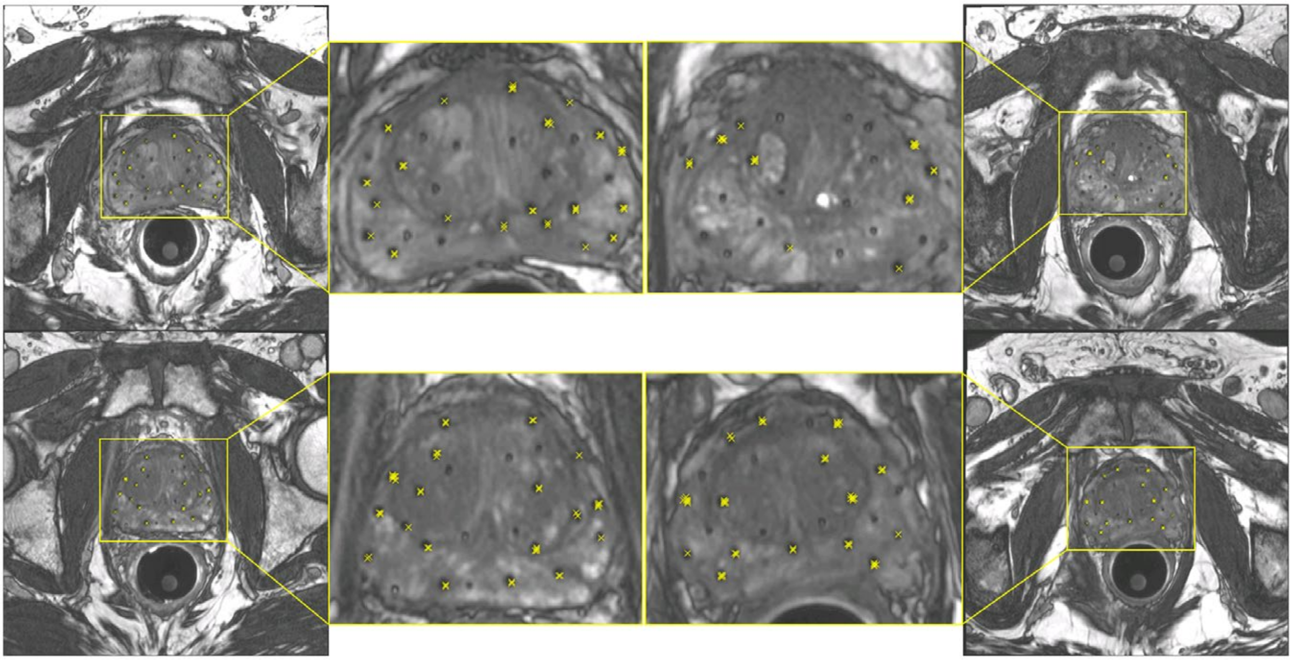

FIGURE 5 Examples of raw SeedNet inferences from SeedNet for 4 representative patients. Each yellow "×" indicates a seed location inference. As expected, multiple inferences were observed for each seed because the sliding window visited each seed multiple times. The seed markers are the circular hyperintense signals. Each circular signal void that appears to be a seed but is not indicated by a yellow "×" is actually an interface between a radioactive seed and seed marker. Connected components analysis was used to combine the clusters into individual seed locations

A comparison of SeedNet inferences for a patient imaged with and without an ERC is shown in Figure 8. A spherical signal void was present in the prostate (red arrow). Because the shape was spherical rather than cylindrical (i.e., like a seed), SeedNet correctly identified it as background (true negative) in images acquired both with and without an ERC. Two seed markers were clearly present near the patient's rectal wall in the image acquired with an ERC (yellow arrows). However, in the image acquired without the ERC, these 2 seed markers exhibited a reduced MR signal and made them

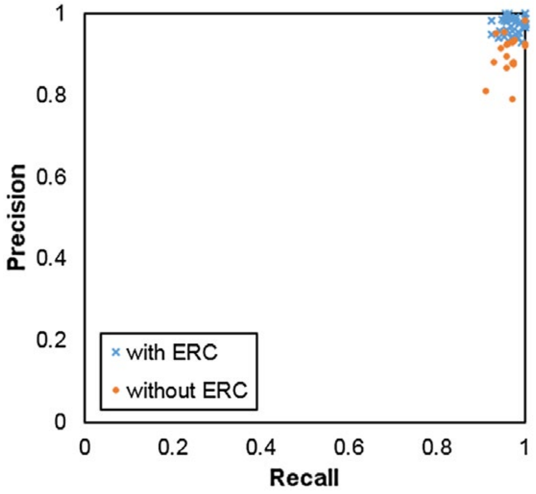

FIGURE 6 SeedNet's precision and recall for the 40 patient images acquired with an ERC and the 20 patient images acquired without an ERC

appear more as signal voids than as seed markers, causing them to be misclassified as seeds (false–positives). Finally, a region of negative contrast between anatomical boundaries was well visualized in the image acquired with the ERC (blue arrow). However, in the image acquired without the ERC, the same region appeared as a cylindrical signal void. The reduced SNR possibly caused SeedNet to incorrectly classify it as a radioactive seed (false–positive).

For the 20 test patient images acquired without the ERC, the precision was 90.5 ± 4.6%, the recall was 96.5 ± 2.3% (Figure 6), and the inference time was 50.5 ± 16.9 s. The F_1-score was 93.3 ± 3.0%, the false discovery rate was 9.5 ± 4.6%, and the false–negative rate was 3.5 ± 2.3% (Figure 7, right column). The RMSE of the seed locations was 0.24 mm ± 0.03 mm. Imaging without the ERC caused a reduction in SNR and overall image quality.[28,29] Additionally, some of the seed markers exhibited substantially less signal and visually appeared more like seeds than as seed markers, possibly contributing to SeedNet's reduced precision.

The clinical integration of SeedNet shown in Figure 4 was performed on a Dell XPS laptop equipped with an NVIDIA GTX 1050 GPU. A comparison of radioactive seed location inferences from SeedNet and a CMD is shown in Figure 9. The resultant radiation dose distributions computed by MIM Software are also shown. For this example, the time required to make seed location inferences using the MIM integration was 90 s. Approximately 25 s of this time was required to initialize the SDK, send the images and prostate contour to the SDK from MIM, and send the seed locations from the SDK to MIM. The remainder of the time (65 s) was required to extract the sub-windows, load the SeedNet models, and make inferences on the sub-windows. The patient shown was implanted with 56 radioactive seeds. SeedNet identified 54/56 of the seeds and 1 false–positive.

4 | DISCUSSION

Existing methods for seed identification in prostate brachytherapy primarily rely on post-implant CT images, and a number of automated seed localization algorithms have been developed and implemented into commercial software packages. Although CT provides excellent seed contrast, it lacks the necessary soft-tissue contrast to delineate the prostate and surrounding organs. To incorporate organ boundary information in the dosimetry, a dosimetrist typically must fuse the CT image with an intraoperative transrectal ultrasound image or an MRI acquired after the implant procedure. Moreover, the dosimetrist must manually co-register the images and identify the seeds. Here, we present an automated approach to MRI-based post-implant seed detection and localization. The high precision and recall, along with the low computation time, demonstrate the potential of fully automated, MRI-only seed identification for performing post-implant dosimetry.

Although it provides intrinsically superior image quality, the use of an ERC has several disadvantages. In our experience, some patients may find the insertion of the ERC painful, especially since most patients are scanned on the day of their implant procedure. In addition, preparation of the ERC may take up to 20–30 min to set up the patient table, which reduces clinical throughput, particularly in high-volume clinics. The rigid ERC is also expensive and requires a specialized sanitation station to sterilize the coil after each patient. For these reasons, the ERC has been limited to a few large academic centers and is seen as an obstacle to the adoption of MR-based post-implant dosimetry by the majority of the brachytherapy community. The impressive precision and recall of SeedNet demonstrate the potential for automated radioactive seed identification in images acquired with and without an ERC. However, degradation in seed identification performance was observed on images acquired without an ERC when compared to the images acquired with an ERC. Methods to improve the image quality when imaging without an ERC are being actively investigated[28] and are expected to improve SeedNet's performance.

A previous study investigated the performance of CMDs in manually identifying radioactive seeds with and without the use of an ERC for imaging.[27] The study revealed that CMDs could identify 90.8 ± 13.3% of the radioactive seeds in 27.2 ± 8.6 min when a rigid ERC was used for imaging. Without an ERC, the CMDs could identify 53.8 ± 33.4% of the radioactive seeds in 39.2 ± 10.4 min. SeedNet's seed identification performance and inference time were better than that of the CMDs in images acquired both with and without an ERC, making it an attractive alternative to manual seed identification. Our results showed that SeedNet can identify almost all of the radioactive seeds in images acquired without an ERC, which may enable fully automated and MRI-only

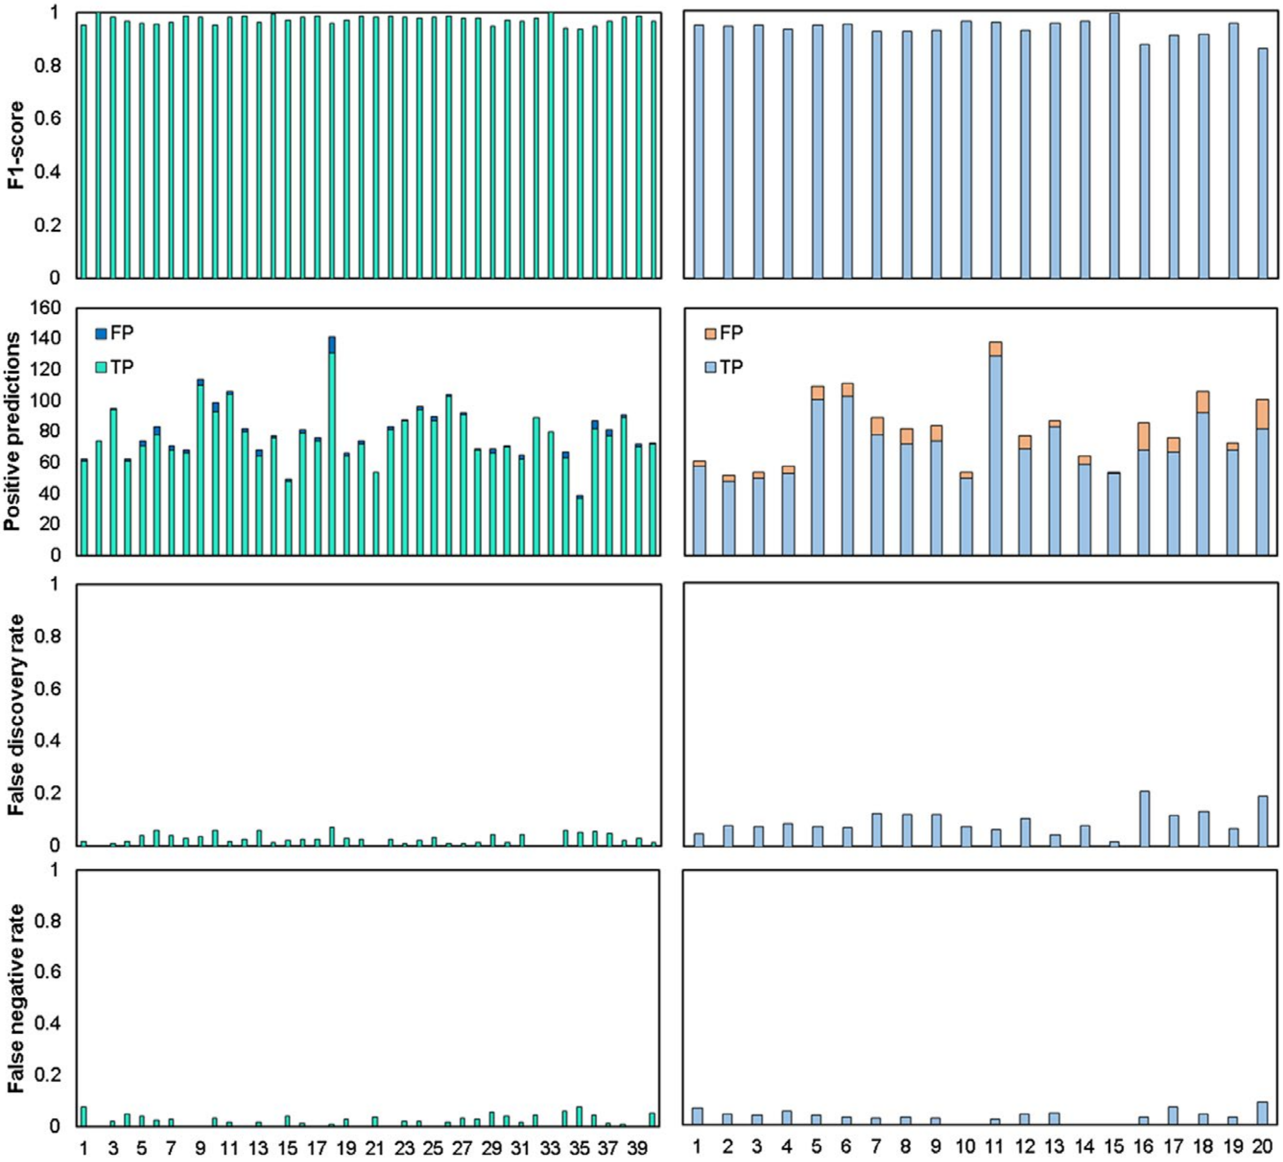

FIGURE 7 F_1-score, positive prediction summary, false discovery rate, and false–negative rate for the randomly selected test patients acquired with (left) and without (right) an ERC

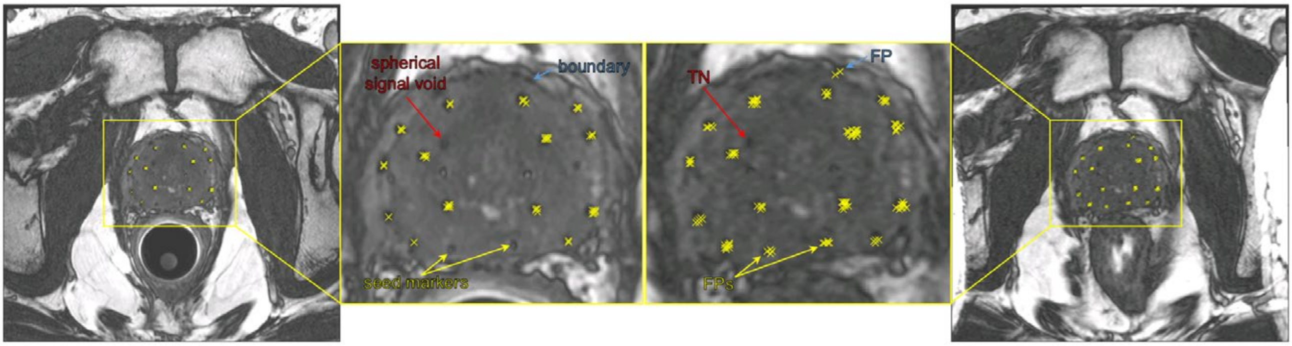

FIGURE 8 Comparison of raw SeedNet inferences for the same patient imaged with (left) and without (right) an ERC. TN, true–negative; FP, false–positive

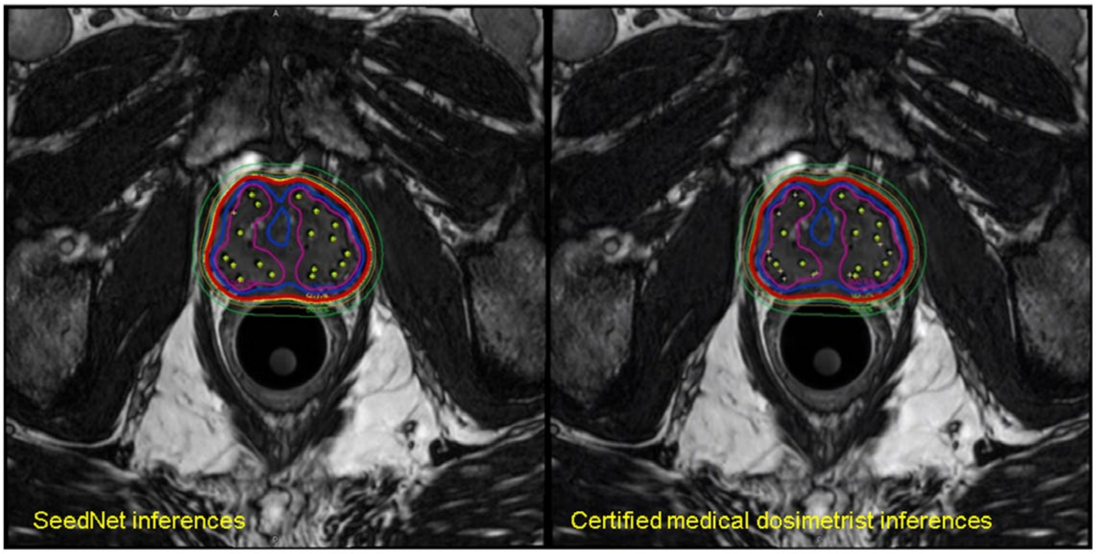

FIGURE 9 Example of SeedNet's (left) and a CMD's (right) seed location inferences and corresponding radiation dose distributions in MIM Software. The prostate contour was deactivated to allow for unobscured visualization of the isodose lines. The radioactive seeds are depicted as green circles

post-implant assessment without an ERC, thereby expanding the access of precision prostate brachytherapy to community hospitals and reducing the overall cost of brachytherapy. In addition, acquiring the post-implant images and performing the dosimetry are ideally done on the same day of the implant procedure. In the rare event that the patient does not receive the prescription dose to the prostate, the patient's treatment may need to be supplemented with additional radioactive seeds. On high-volume days, such patients may be required to stay an additional day before going home because of the lengthy time to perform the dosimetry. By significantly reducing the seed identification time to 1–2 min or less, SeedNet may make it possible to increase the daily capacity for prostate brachytherapy treatments and send all patients home on the day of their implant procedure.

The clinical implementation of SeedNet carries some practical considerations. The optimal sub-window size may vary depending on the scan parameters selected for imaging. For example, a sub-window size other than 13 × 13 × 7 may yield higher precision and recall for institutions using a spatial resolution different from that used in this study. In these cases, the sub-window size used by SeedNet may need to be reevaluated. In addition, although SeedNet has excellent precision and recall, some false–positives or false–negatives may occur. In such instances, a CMD may need to review the seed location inferences to augment or supplement the inferences. However, the CMD will have the entire image context to infer from (rather than a small sub-window), which should make it relatively easy to reject any remaining false–positives or correct false–negatives.

Although MRI may be the best approach to post-implant assessment, many clinics do not have the resources necessary to perform MRI-only dosimetry. The approach presented in this article could potentially be used to automatically localize radioactive seeds in CT images as well. Like post-implant MR images, post-implant CT images are acquired with consistent acquisition parameters across patients. On CT, radioactive seeds appear as clustered high-contrast regions, whereas the seed markers have Hounsfield units similar to those of soft-tissue. Therefore, the training data would contain sub-windows of only 2 classes: seed and background. As such, a separate classification CNN may not be required, which could simplify the model and reduce the computational complexity.

One notable limitation of our study is that all of the images were acquired on a single MRI scanner from a single vendor. Therefore, our training data did not include potential variability in scanner hardware, receiver coil sensitivities, pulse sequence implementations, and image reconstruction across different vendors. A future study could incorporate images acquired on multiple scanners into the training data to evaluate the generalizability of SeedNet. Another limitation of our study is that our training data only contained images acquired from our institution. As a result, variability in patient setup across institutions was not included in the training data. However, our model weights were trained on a large number of sub-windows, and the features learned by the model are expected to be useful for other institutions interested in implementing SeedNet. To facilitate investigation by interested readers, we have made our trained models available for download at https://github.com/jeremiahws/SeedNet. A transfer learning approach[48] could be used to tailor SeedNet to the images acquired at other institutions.

5 | CONCLUSION

We developed and evaluated a sliding-window CNN that uses a fixed sub-window size to detect and localize radioactive seeds in post-implant MRI of the prostate after LDR brachytherapy. SeedNet demonstrated high precision and recall for radioactive seed detection in a fraction of the time typically required of a CMD. SeedNet can localize implanted radioactive seeds in MR images acquired with or without ERC, albeit with reduced precision when an ERC is not used. Image quality improvements achieved through pulse sequence optimization may improve SeedNet's performance when imaging without a rigid ERC, potentially enabling more widespread access and use of MRI for precision LDR prostate brachytherapy. SeedNet can be incorporated into clinical software packages and clinical workflows for improved efficiency in LDR prostate brachytherapy.

ACKNOWLEDGMENTS

The authors would like to thank The University of Texas MD Anderson Cancer Center's Department of Scientific Publications for their help editing this paper.

CONFLICTS OF INTEREST

Steven J. Frank is a co-founder of C4 Imaging and has ownership interests in the company. Steven J. Frank is also an inventor of US and international patents on the C4 seed marker technology. Jingfei Ma is an inventor of US patents that are currently licensed to Siemens Healthineers and is a consultant for C4 Imaging.

ORCID

Jeremiah W. Sanders https://orcid.org/0000-0002-5342-4128

REFERENCES

1. Carey B, Swift S. The current role of imaging for prostate brachytherapy. *Cancer Imaging*. 2007;7:27–33.
2. Christodouleas J, Fox J, Song D, DeWeese T. Basic terms and concepts of radiation. In: Su L, editor. *Early diagnosis and treatment of cancer: prostate cancer*. Philadelphia: Saunders; 2010. p 159–176.
3. Brinkmann DH, Kline RW. Automated seed localization from CT datasets of the prostate. *Med Phys*. 1998;25:1667–1672.
4. Liu H, Cheng G, Yu Y, et al. Automatic localization of implanted seeds from post-implant CT images. *Phys Med Biol*. 2003;48:1191–1203.
5. Nguyen HG, Fouard C, Meneu F, et al. Automatic 3D seed location and orientation detection in CT image for prostate brachytherapy. In Proceedings of the 11th IEEE International Symposium on Biomedical Imaging, Beijing, China, 2014. p. 1320–1323.
6. Nitz WR, Reimer P. Contrast mechanisms in MR imaging. *Euro Radiol*. 1999;9:1032–1046.
7. Amdur RJ, Gladstone D, Leopold KA, et al. Prostate seed implant quality assessment using MR and CT image fusion. *Int J Radiat Oncol Biol Phys*. 1999;43:67–72.
8. De Brabandere M, Al-Qaisieh B, De Wever L, et al. CT- and MRI-based seed localization in postimplant evaluation after prostate brachytherapy. *Brachytherapy*. 2013;12:580–588.
9. Crook J, McLean M, Yeung I, et al. MRI-CT fusion to assess postbrachytherapy prostate volume and the effects of prolonged edema on dosimetry following transperineal interstitial permanent prostate brachytherapy. *Brachytherapy*. 2004;3:55–60.
10. Polo A, Cattani F, Vavassori A, et al. MR and CT image fusion for postimplant analysis in permanent prostate seed implants. *Int J Radiat Oncol Biol Phys*. 2004;60:1572–1579.
11. De Brabandere M, Hoskin P, Haustermans K, et al. Prostate postimplant dosimetry: Interobserver variability in seed localisation, contouring and fusion. *Radiother Oncol*. 2012;104:192–198.
12. Katayama N, Takemoto M, Yoshio K, et al. T2*-weighted image/T2-weighted image fusion in postimplant dosimetry of prostate brachytherapy. *J Radiat Res*. 2011;52:680–684.
13. Dubois DF, Prestidge BR, Hotchkiss LA, et al. Source localization following permanent transperineal prostate interstitial brachytherapy using magnetic resonance imaging. *Int J Radiat Oncol Biol Phys*. 1997;39:1037–1041.
14. Ohashi T, Momma T, Yamashita S, et al. Impact of MRI-based postimplant dosimetric assessment in prostate brachytherapy using contrast-enhanced T1-weighted images. *Brachytherapy*. 2012;11:468–475.
15. Moerland MA, Wijrdeman HK, Beersma R, et al. Evaluation of permanent I-125 prostate implants using radiography and magnetic resonance imaging. *Int J Radiat Oncol Biol Phys*. 1997;37:927–933.
16. Frank SJ, Stafford RJ, Bankson JA, et al. A novel MRI marker for prostate brachytherapy. *Int J Radiat Oncol Biol Phys*. 2008;71:5–8.
17. Frank SJ, Johansen MJ, Martirosyan KS, et al. A biodistribution and toxicity study of cobalt dichloride-n-acetyl cysteine in an implantable MRI marker for prostate cancer treatment. *Int J Radiat Oncol Biol Phys*. 2013;85:1024–1030.
18. Lim TY, Kudchadker RJ, Wang J, et al. Development of a magnetic resonance imaging protocol to visualize encapsulated contrast agent markers in prostate brachytherapy recipients: initial patient experience. *J Contemp Brachytherapy*. 2016;8:235–242.
19. Ma J, Moerland MA, Venkatesan AM, et al. Pulse sequence considerations for simulation and postimplant dosimetry of prostate brachytherapy. *Brachytherapy*. 2017;16:743–753.
20. Kuo N, Lee J, Tempany C, et al. MRI-based prostate brachytherapy seed localization. In Proceedings of the 7th IEEE International Symposium on Biomedical Imaging: From Nano to Macro, Rotterdam, Netherlands, 2010. p. 1397–1400.
21. Martin GV, Ma J, Kudchadker RJ, et al. Automated prostate brachytherapy seed detection on post-implant MRI using novel martin algorithm. *Brachytherapy*. 2017;16:S57.
22. Krizhevsky A, Sutskever I, Hinton GE. ImageNet classification with deep convolutional neural networks. *Adv Neural Inf Process Syst*. 2012;25.
23. Sermanet P, Eigen D, Zhang X, et al. OverFeat: integrated recognition, localization and detection using convolutional networks. arXiv preprint arXiv 2013:1312.6229.
24. Li X, Wang S. Object detection using convolutional neural networks in a coarse-to-fine manner. *IEEE Geosci Remote Sens Lett*. 2017;14:2037–2041.

25. Ren S, Kaiming H, Girshick R, et al. R-CNN: towards real-time object detection with region proposal networks. arXiv preprint. *arXiv*. 2015;1506.01497.
26. Caicedo JC, Lazebnik S. Active object localization with deep reinforcement learning. arXiv preprint. *arXiv*. 2015;1511.06015.
27. Najibi M, Rastegari M, Davis LS. G-CNN: an iterative grid based object detector. arXiv preprint. *arXiv*. 2015;1512.07729.
28. Sanders JW, Song H, Frank SJ, et al. Parallel imaging compressed sensing for accelerated imaging and improved signal-to-noise ratio in MRI-based post-implant dosimetry of prostate brachytherapy. *Brachytherapy*. 2018;17:816–824.
29. Sanders J, Frank S, Bathala T, et al. MRI-based prostate brachytherapy - imaging with and without an endorectal coil for post-implant quality assurance. *Brachytherapy*. 2017;16:S56.
30. Sanders JW, Lewis G, Frank S, et al. A fully convolutional network utilizing depth-wise separable convolutions for semantic segmentation of anatomy in MRI of the prostate after permanent implant brachytherapy. In Proceedings of the ISMRM Workshop on Machine Learning, Pacific Grove, CA, 2018.
31. Milletari F, Navab N, Ahmadi SA. V-Net: fully convolutional neural networks for volumetric medical image segmentation. arXiv preprint. *arXiv*. 2016;1606.04797.
32. Turkbey B, Fotin SV, Huang RJ, et al. Fully automated prostate segmentation on MRI: comparison with manual segmentation methods and specimen volumes. *AJR Am J Roentgenol*. 2013;201:W720–W729.
33. Ajani B, Krishnan K. Automatic and interactive prostate segmentation in MRI using learned contexts on a sparse graph template. In Proceedings of the 14th IEEE International Symposium on Biomedical Imaging, Melbourne, Australia, 2017. p. 315–318.
34. Zhu Q, Du B, Turkbey B, et al. Deeply-supervised CNN for prostate segmentation. arXiv preprint. *arXiv*. 2017;1703.07523.
35. Yu L, Yang X, Chen H, et al. Volumetric convnets with mixed residual connections for automated prostate segmentation from 3D MR images. In Proceesings of the AAAI Conference on Artificial Intelligence, 2017:66–72.
36. Litjens G, Toth R, van de Ven W, et al. Evaluation of prostate segmentation algorithms for MRI: the PROMISE12 challenge. *Med Image Anal*. 2014;18:359–373.
37. Lecun Y, Bottou L, Bengio Y, et al. Gradient-based learning applied to document recognition. *Proc IEEE*. 1998;86:2278–2324.
38. More A. Survey of resampling techniques for improving classification performance in unbalanced datasets. arXiv preprint. *arXiv*. 2016;1608.06048.
39. Kingma DP, Ba J. Adam: a method for stochastic optimization. arXiv preprint. *arXiv*. 2014;1412.6980.
40. Srivastava N, Hinton G, Krizhevsky, et al. Dropout: a simple way to prevent neural networks from overfitting. *J Mach Learn Res*. 2014;15:1929–1958.
41. Glorot X, Bengio Y. Understanding the difficulty of training deep feedforward neural networks. *J Mach Learn Res*. 2010;9:249–256.
42. Chollet F. Keras: the python deep learning library. https://keras.io/. Published March 27, 2015. Updated January 25, 2019. Accessed January 29, 2019.
43. Wachowicz K, Thomas SD, Fallone BG. Characterization of the susceptibility artifact around a prostate brachytherapy seed in MRI. *Med Phys*. 2006;33:4459–4467.
44. Sanders J, Fletcher J, Frank S, et al. Deep learning application engine (DLAE): end-to-end development and deployment of medical deep learning algorithms. In Proceedings of the ISMRM Workshop on Machine Learning Part II, Washington, DC, 2018.
45. Sanders J, Frank S, Ma J. Automated seed detection and localization in MRI of the prostate after permanent implant brachytherapy using a sliding-window convolutional neural network. In Proceedings of the ISMRM Workshop on Machine Learning, Pacific Grove, CA, 2018.
46. Sanders J, Frank S, Ma J. SeedNet: a sliding-window convolutional neural network for radioactive seed detection and localization in MRI. In Proceedings of the 26th Annual Meeting of the ISMRM, Paris, France, 2018. Abstract 3421.
47. Sanders J, Bruno T, Frank S, et al. SeedNet for automated detection and localization of radioactive seeds in post-implant MRI: a comparison with and without the use of an endorectal coil for imaging. *Proc Brachytherapy*. 2018;17:S80–S81.
48. Lisa T, Jude S. Transfer learning. In: Soria OE, Guerrero JDM, Sober MM, Benedito JRM, Lopez AJS, eds. *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques*. Hershey: IGI Global; 2010:242–264.

How to cite this article: Sanders JW, Frank SJ, Kudchaker RJ, Bruno TL, Ma J. Development and clinical implementation of SeedNet: A sliding-window convolutional neural network for radioactive seed identification in MRI-assisted radiosurgery (MARS). *Magn Reson Med*. 2019;00:1–13. https://doi.org/10.1002/mrm.27677