

**UCSF**

**UC San Francisco Electronic Theses and Dissertations**

**Title**

A Faster R-CNN Model for Detection of MR-Compatible Catheters

**Permalink**

<https://escholarship.org/uc/item/5629h4qz>

**Author**

Alcantar, Casey

**Publication Date**

2020

Peer reviewed|Thesis/dissertation

A Faster R-CNN Model for Detection of MR-Compatible Catheters

by  
Casey Alcantar

THESIS

Submitted in partial satisfaction of the requirements for degree of  
MASTER OF SCIENCE

in

Biomedical Imaging

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

Approved:

DocuSigned by:

*Steven Hetts*

Steven Hetts

DA41E0D34A4A463...

Chair

DocuSigned by:

*Alastair Martin*

Alastair Martin

DocuSigned by:

*Peder Larson*

Peder Larson

DocuSigned by:

*Valentina Padoia*

Valentina Padoia

CB0AB38D57D842C...

Committee Members

Copyright 2020  
by  
Casey Alcantar

## **Acknowledgements**

Thank you to everyone who gave me advice and guidance for this project, especially during these unprecedented times. Thank you to my project advisory committee members Dr. Steven Hetts, Dr. Alastair Martin, Dr. Peder Larson, and Dr. Valentina Padoia. Additional thanks to Teri Moore, Bridget Kilbride, Dr. Caroline Jordan, and the IR Lab for all your support.

# **A Faster R-CNN Model for Detection of MR-Compatible Catheters**

**By Casey Alcantar**

## **Abstract**

This project sought to create a deep learning model to detect and track MR-compatible catheter tips under Magnetic Resonance Imaging. Interventional MRI, or iMRI, has many advantages over traditional x-ray angiography methods, yet the path towards adoption is hindered by many obstacles, including the lack of easily visualizable catheter tips. The model, the Faster Region-based Convolutional Neural Network (Faster R-CNN), was chosen due to its well-balanced speed and accuracy over other model architectures. The dataset included MR images of passive and resonant catheter tips alone and as well as passive catheter tips in an abdominal aorta phantom. The Faster R-CNN was trained over many iterations and over the best run it was able to draw bounding boxes over the tip of the catheter with an overall mean average precision of 0.59 and overall average recall of 0.66. Further optimization of training parameters will be needed to create a model that can achieve a better mean average precision. This study opens the possibility of applying artificial intelligence models towards iMRI methods, which helps push towards the goal of proving the safety and efficacy of iMRI procedures. These foundational elements are critical to smoothing the adoption of iMRI for guiding endovascular procedures.

## Table of Contents

Introduction.....	1
Overview .....	1
Hypothesis.....	4
Methods.....	4
Overview .....	4
Dataset Preparation .....	5
System Architecture .....	10
Results.....	13
Classifier.....	13
Faster R-CNN.....	14
Discussion.....	18
Conclusion .....	20
References.....	21

## List of Figures

Figure 1 .....	3
Figure 2 .....	5
Figure 3 .....	6
Figure 4 .....	8
Figure 5 .....	9
Figure 6 .....	9
Figure 7 .....	10
Figure 8 .....	11
Figure 9 .....	12
Figure 10 .....	13
Figure 11 .....	15
Figure 12 .....	15
Figure 13 .....	16
Figure 14 .....	17
Figure 15 .....	18
Figure 16 .....	18

**List of Tables**

Table 1 .....16

Table 2 .....16

## **Introduction**

### *Overview*

The phrase, “time is brain” means that the longer the brain is deprived of blood flow in an ischemic stroke, the further the extent of brain cell death<sup>1</sup>. For an interventionalist, the concern is that the longer that brain tissue is infarcted, the more likely that there will be hemorrhage after blood flow is restored to the occluded artery. Every year, more than 780,000 people in the United States have an acute ischemic stroke (AIS)<sup>2</sup>. The current recommendations for endovascular treatment of acute ischemic stroke is within 3 to 6 hours of stroke onset, but fewer than 10% of AIS patients receive treatment due to this narrow treatment window<sup>3</sup>.

Advances in stroke diagnosis, including diffusion weighted MRI, have shifted the paradigm to “physiology is brain”, meaning that clinicians can now visualize brain tissue to determine if it is salvageable when flow is restored, regardless of how long it has been infarcted<sup>4</sup>. Stroke treatment outside of the 6 hour window requires initial MRI imaging to diagnose the location and extent of the stroke using diffusion-weighted MRI, and treatable patients are subsequently taken to the x-ray angiography suite for clot extraction<sup>5</sup>. With high temporal and spatial resolution, x-ray angiography is the current clinical standard for treatment of AIS<sup>6</sup>. However, x-ray angiography exposes patients and staff to ionizing radiation, is unable to provide quantitative physiologic information, and provides poor soft tissue characterization. Over the past two decades, research has been performed to make interventional MRI possible, citing that MRI has many advantages, including the ability to determine tissue viability intraprocedurally, and the lack of exposure to ionizing radiation<sup>7</sup>.

The ability to diagnose a treatable AIS and to immediately progress to an interventional procedure under MR guidance is appealing due to the time efficiencies and the ability to

immediately assess the impact of a therapeutic intervention. This approach requires MR methods to provide vascular road-mapping and adequate visualization and tracking of devices<sup>8</sup>. Recent advances have demonstrated the ability of MRI to provide these capabilities, albeit with inferior performance when compared to x-ray digital subtraction angiography (DSA).

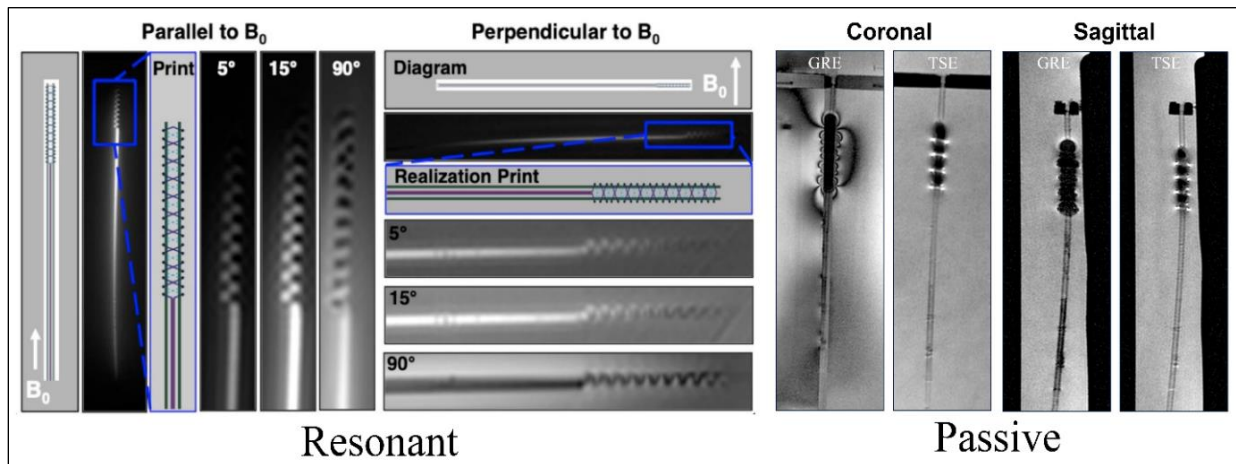
MR visualizes a 2-dimensional scan plane within a 3-dimensional volume, whereas x-ray DSA uses projections where overlying structures are superimposed on one another<sup>9</sup>.

Furthermore, spatial and temporal resolutions are typically traded off against each other. This creates a problem when attempting to visualize the catheter in the vasculature, since the catheter is traveling within the 3D volume and may move out of the scan plane as it is advanced.

Traditional catheters used in x-ray DSA are constructed out of metallic materials, which can cause radiofrequency-induced heating in MRI<sup>10</sup>. MR safe catheters can be constructed from plastic polymers but can be hard to localize in MR images and tend to have inferior mechanical properties. Various MR visible markers have been created to increase the visualization of catheters, and most importantly, the catheter tip.

Methods for MRI markers can be categorized as active, resonant, and passive<sup>10</sup>. Active markers generally use a micro-coil at the tip with of the catheter and have wires leading down the length of the catheter to receive information about the location of the catheter tip<sup>9</sup>. However due to the long length of the conductor, active catheters also run the risk of radiofrequency-induced heating of surrounding tissue<sup>6</sup>. Resonant catheters, also known as semi-active catheters, use circuits with capacitors connected to a coil on the catheter tip, which locally amplifies the  $B_1+$  field and enhances the signal or alters contrast<sup>9</sup>. Passive markers utilize ferromagnetic, ferrimagnetic, or strongly paramagnetic materials that cause local  $B_0$  field inhomogeneities and cause the catheter to be displayed as a signal void in the image. An interventionalist that is

relying on MR fluoroscopy to track a catheter must recognize these markers against a variable background and further be aware of the potential for the catheter to migrate out of the visualized imaging plane. Methods to support the clinician in this regard would have substantial benefit.



**Figure 1**

MR images of resonant catheter tips (left) and passive catheter tips (right)<sup>9</sup>. In the resonant catheter tip, the coil is printed in a double helical shape and the bright areas are where the coil amplifies the signal. The dark areas are where the coil is not present and there is a lack of signal. The signal amplification varies with the flip angle of the sequence and also the orientation of the catheter in relation to the main magnetic field. In the passive catheter tip, radiopaque ink mixed with iron oxide nanoparticles is painted onto the catheter in four bands that are 5mm apart from each other. The iron in the ink causes disruption of the main magnetic field so that no coherent signal is created. A signal void is created directly over the bands and also extends into the immediately surrounding area, causing the tip to be visualized as a large signal void.

Machine learning, and more recently, deep learning methods have seen wide adoption in the medical imaging space and MRI<sup>11</sup>. While deep learning methods have been used in interventional x-ray and ultrasound modalities, there have been few methods developed using interventional MRI<sup>12,13</sup>. Prior research on real-time catheter segmentation in x-ray fluoroscopy has demonstrated that it is feasible to perform segmentations in real-time scales, with segmentation times of 125 ms per frame<sup>12</sup>. With this in mind, researching a new neural network to segment an MRI-compatible catheter is a plausible next step.

The overall problem of detecting a catheter can be simplified into two problems, object classification and object localization, or determining the type of object and predicting the coordinates of the object if it is in the image. A Faster Region-based Convolutional Neural Network (Faster R-CNN) is able to perform both tasks at a computational speed of 10 ms per image<sup>14</sup>. This is faster than the fps of a real-time interventional MRI sequence, which is approximately 10 fps<sup>15</sup>. While there are other deep convolutional methods that are able to detect objects at faster speeds, such as YOLO or SSD, they are generally less accurate<sup>16</sup>.

### *Hypothesis*

In this work, we propose to utilize a Faster Region-based Convolutional Neural Network (Faster R-CNN) to track resonant and passive MR-compatible catheters in real-time MR imaging and alert when the catheter leaves the imaging volume. This Faster R-CNN will accomplish multiple steps, including tracking the catheter tip in real-time as it moves throughout the field of view and drawing a bounding box over the tip. Using a deep learning model to quickly visualize and track MR-compatible catheters can improve operator confidence and facilitate the adoption of interventional MRI.

## **Methods**

### *Overview*

Training and deploying a neural network required multiple steps, including creating both a dataset and system architecture. Creating the dataset involved collection, categorization, and preparation of data. Creating the system architecture involved preprocessing, training, and model selection. Each of these steps are highlighted in detail below.

## *Dataset preparation*

### *Classifier dataset*


MRI images of the passive and resonant catheter tips were obtained from prior pre-clinical research in the IR lab. The resonant catheter images were acquired at 3.0 T (Discovery MR 750 w, GE Healthcare, Chicago, IL) using a balanced steady state free precession sequence with varying flip angles of 5, 15, and 90 degrees, where the catheter was oriented either parallel or perpendicular to the  $B_0$  field. The passive catheter images were acquired at 1.5 T MRI (Achieva, Philips Medical Systems, Best, The Netherlands) using both gradient echo (GRE) and turbo spin echo (TSE) sequences. Both coronal and sagittal slices were acquired along the length of the catheter. Images were uploaded into Horos software and were exported to JPEG format in a 240 x 240 pixel image size (Nimble Co LLC d/b/a Purview, Annapolis, MD). The images were processed using Matlab 2019b for contrast stretching and conversion of images from red-green-blue (RGB) to grayscale (MathWorks Inc., Natick, MA). Contrast stretching boosts contrast based on the image's histogram to improve normalization and line detection in varying lighting conditions.



**Figure 2**

Examples of MR images of passive (left) and resonant (right) catheters used for classification layers.

The size of both datasets was relatively small, where the passive dataset consisted of 160 images and the resonant dataset consisted of 386 images. Both datasets were augmented offline using python skimage<sup>17</sup>. The augments included random 90° vertical flips to change the orientation of the catheter, random rotations between -35° to 35°, and random addition of noise not exceeding 5% of the pixels in the image. These augmentations were performed on both datasets until both reached a total of 2,000 images each. The datasets were then separately split into train, testing and validation sets by 80:10:10 ratios.




**Figure 3**

Examples of augmented images of passive (A, B) and resonant catheters (C, D). Image A depicts the addition of noise, Image B depicts image rotation, Image C depicts noise, and Image D depicts image rotation.

### *Tracking dataset*

Similarly, the dataset of the passive catheter moving through a simplified abdominal aortic vascular phantom model was acquired on a 1.5 T MRI (Achieva, Philips Medical Systems, Best, The Netherlands) using a GRE sequence and taken at approximately 3 frames per second (fps). Both coronal and sagittal slices were acquired. The abdominal aorta phantom model was made out of polyvinyl alcohol cryogel (Sevol Grade 165 PVA powder, Sekisui Specialty Chemicals America, Dallas, TX), and created using Delrin rods (McMaster-Carr, Elmhurst, Ill), to generate physiologically similar vessel trajectories and angles, including the celiac artery, superior mesenteric artery, bilateral renal arteries, and inferior mesenteric artery<sup>18</sup>. The phantom was submerged in water and an MRI-safe 15-F Check-Flo Performance Introducer (Cook, Bloomington, IN) was used to mimic vascular access to the aorta. Three different series of images were acquired using differing slab thicknesses: 5 mm, 10 mm, and 20 mm. Overall, 2749 usable images were obtained. Images were uploaded into Horos software and exported to JPEG format in a 240 by 240 pixel image size. The images were preprocessed using Matlab 2019b, including contrast stretching for normalization and conversion from RGB to grayscale. Images of the resonant catheter were unable to be obtained for this study.


The passive dataset was split into an approximate 80:10:10 ratio, where it was confirmed that the validation set contained images from the 5mm, 10mm, 20mm slab thickness series. The training dataset also underwent offline augmentation. The augmentations included random 90° vertical flips to change the orientation of the catheter, random rotations between -35° to 35°, and random addition of noise not exceeding 5% of the pixels in the image. These augmentations doubled the training set from 2043 images to 4086.



**Figure 4**

MR images of a passive catheter in a vascular phantom with varied slab thicknesses, including 5mm (left), 10mm (middle), and 20mm (right).

The ground truth bounding box annotations were created using the MD.ai website software. Bounding boxes were created by using the ruler tool to measure a length of 46 mm from the less distal signal void towards the distal end of the catheter. The actual physical measurements of the most proximal marker on the passive catheter was 40 mm from the distal tip, but under MRI, the signal void extended approximately 6 mm from this marker. The bounding box tool was then used to draw a box around this area with a 1 mm buffer on each side. This resulted in a box with a length of 48 mm when the catheter was most parallel to the aortic vessel. The width varied in size due to the change in angle of the catheter, and ranged from 10 mm when parallel, to 38 mm when it was angled while it was advancing through the smaller vessels. The bounding box annotations were reviewed by experienced radiographers for quality assurance. To prepare the files for use with Tensorflow, the annotation files were exported from the MD.ai website into a json file and were subsequently converted to a csv file. The image files and corresponding csv files were then used to create TFRecord files for the training, test, and validation sets.



**Figure 5**

Depiction of passive catheter measurements (left), and measurements taken of image using ruler tool on the MD.ai website (right). The four markers were painted over a 20 mm area and the marker closest to the distal tip (most distal) of the catheter was 20 mm away.



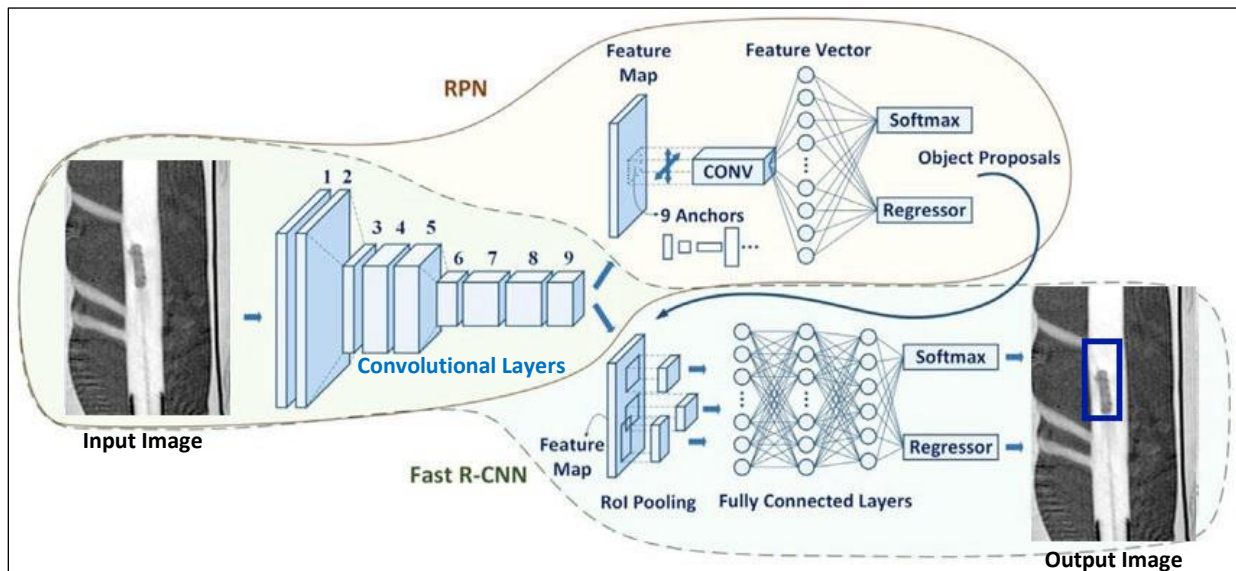
**Figure 6**

Screenshot of MD.ai website with user-generated bounding box to be used as the ground truth.

## System Architecture

### Faster Region-based Convolutional Neural Network Model

A Faster R-CNN model was implemented to quickly detect and track the catheter tip<sup>14</sup>. The Faster R-CNN model was composed of two modules, a Fast R-CNN detector and a region proposal network (RPN), that share convolutional layers within the object detection network<sup>19</sup>. The Fast R-CNN detector used a CNN to extract features from input images and classify the features to a specific class. The RPN was used to create bounding boxes over the area where there could be a particular object using a second convolutional network to decrease computational time from 2 s to 10 ms per image. The RPN first slid a small  $n \times n$  window over the feature map of the entire image made from the pretrained classification CNN and predicted a number of possible bounding box regions of various scales and ratios at the center of each sliding window.



**Figure 7**

Representation of a Faster R-CNN architecture, modified from Ren et al<sup>14</sup>. Two modules are shown, the Region Proposal Network (RPN) and the Fast R-CNN. Both modules share the initial Convolutional layers. Input and output images are of a passive catheter in an arterial phantom, with a blue bounding box superimposed over the catheter in the output image.

## Loss function

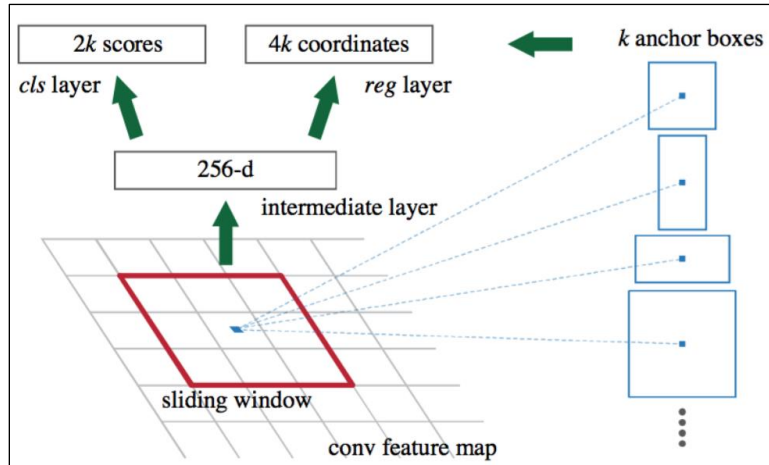
The loss function of the Faster R-CNN network used a combined loss function for both the classification and bounding box regression:

$$\mathcal{L} = \mathcal{L}_{cls} + \mathcal{L}_{box}$$

$$\mathcal{L}(\{p_i\}, \{t_i\}) = \frac{1}{N_{cls}} \sum_i \mathcal{L}_{cls}(p_i, p_i^*) + \frac{\lambda}{N_{box}} \sum_i p_i^* \cdot L_1^{smooth}(t_i - t_i^*)$$

$$\mathcal{L}_{cls}(p_i, p_i^*) = -p_i^* \log p_i - (1 - p_i^*) \log(1 - p_i)$$

where  $L_{cls}$  was the log loss function over two classes,  $L_{box}$  was the log loss function over the bounding box,  $L_1^{smooth}$  was the smooth L1 loss,  $p_i$  was the predicted probability of anchor  $I$  being an object ( $p_i = 1$ ),  $p_i^*$  was the ground truth label of whether anchor  $I$  is an object,  $t_i$  was the predicted four parameterized coordinates, and  $t_i^*$  were the ground truth coordinates.  $N_{cls}$  was the normalization term set to be a mini-batch size, which at default is set to 256,  $N_{box}$  was the normalization term set to the number of anchor locations, and  $\lambda$  was a balancing parameter for the weighting of  $L_{cls}$  and  $L_{box}$ .




**Figure 8**

Representation of the Region Proposal Network (RPN)<sup>14</sup>. The RPN generates proposals by using an  $n \times n$  sliding window over the output of the convolution feature map which then proposes a  $k$  number of anchor boxes over each anchor point. These outputs are mapped to an intermediate layer and subsequently fed into the box-classification (*cls*) and box regression (*reg*) layer.

## Classifier

The initial step in the Faster R-CNN architecture were the initial convolutional layers that created a feature map which were then fed into the RPN and Fast R-CNN detector. In the original Faster R-CNN paper by Ren et al, these layers were based on the VGG16 architecture<sup>14</sup>. The structure of VGG16 consisted of a total of thirteen convolutional layers with intermittent max pooling layers, three fully connected layers, and a 1000-way softmax classifier at the end<sup>20</sup>. The Faster R-CNN was intended to detect catheter tips between Passive and Resonant catheters, and as such, the classification layers were created using the VGG16 architecture and were fed images of both passive and resonant catheters in the classification dataset.



**Figure 9**

Representation of the VGG16 structure<sup>20</sup>. VGG16 is composed of a total of 16 weight layers, including the 13 convolution layers (convolution + ReLU) and the 3 fully connected layers (fully connected+ReLU).

## Transfer learning


Transfer learning is a method used in machine learning to apply an already preexisting model to another dataset. The Faster R-CNN model was additionally trained using transfer learning which was originally planned to be used as a baseline model from which the model trained from scratch was to be measured against. Many datasets have been created for the

application of object detection, including COCO or Common Objects in Context, which is a dataset that includes 120,000 images of 80 different types of common objects<sup>21</sup>. Unfortunately there is no current equivalent dataset for medical images, so choosing a model that was trained on a large dataset with many different types of objects the best available option. Instead of using the VGG16 model as the classification model, an Inception SSD model that was previously trained on the COCO dataset, which is available from the TensorFlow Object Detection API, was used as the classification layers in the Faster R-CNN model<sup>22</sup>.

## Results

### *Classifier Performance*

The VGG16 model was trained in a Google Colab notebook using a Tesla K90 GPU and using Keras with a Tensorflow backend. The training parameters included training for 25 epochs with 100 steps per epoch and 10 validation steps per epoch. The model reached an accuracy of 1.0, or 100%, validation accuracy of 1.0, loss of 0, and validation loss of 0.0003. The trained VGG16 model was then ready to be used in the overall Faster R-CNN model.

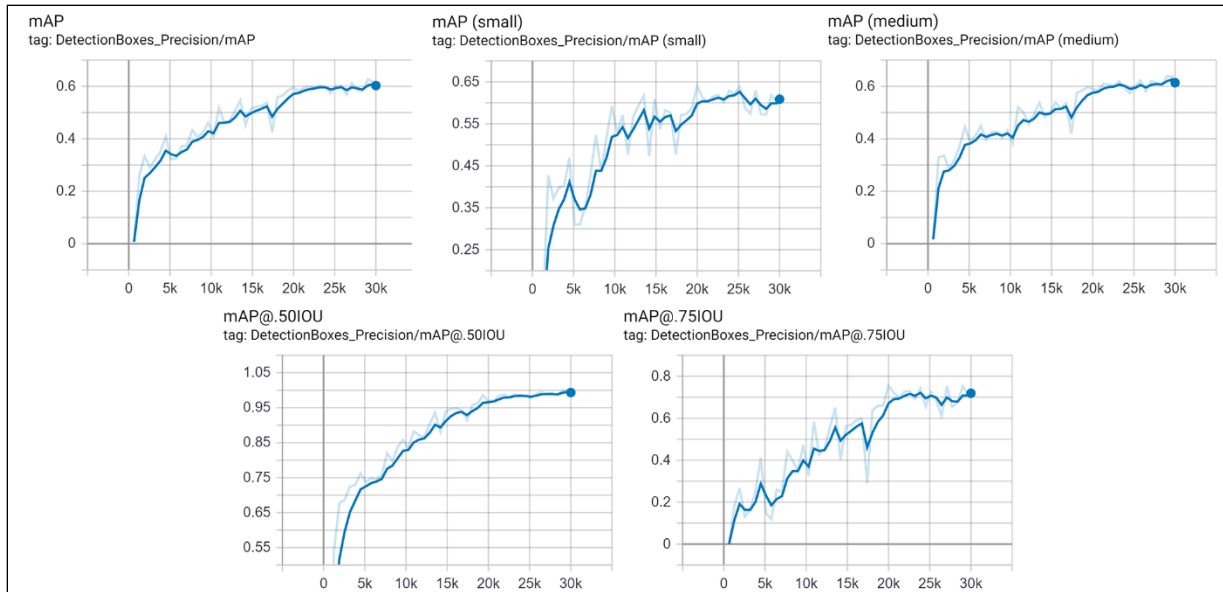


**Figure 10**

Results of Training the catheter tip classification VGG16 model.

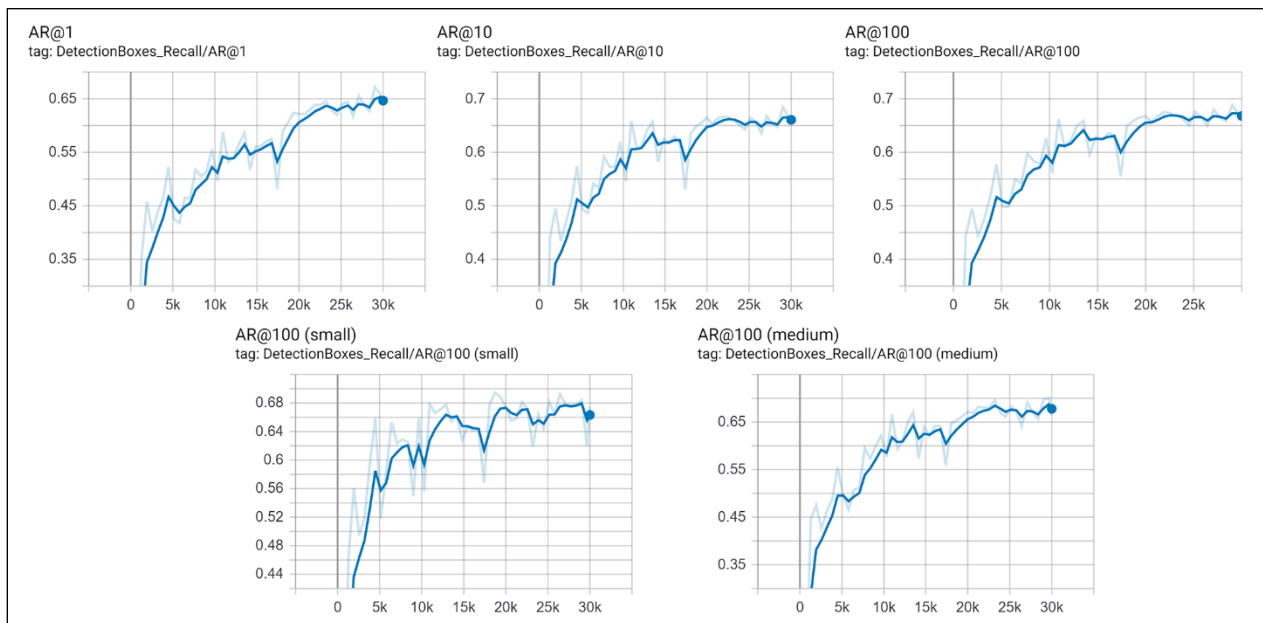
### *Faster R-CNN Model Performance*

The entire Faster R-CNN model was trained in a Google Colab notebook using a Tesla T4 GPU and using Tensorflow. Unfortunately, the VGG16 model classifier trained on the passive and resonant catheters was unable to provide results due to the Region Proposal Network (PRN) not creating bounding boxes over the catheter. The classifier was not creating an appropriate feature map where the RPN could not distinguish between an object and the background. The Faster R-CNN model was trained instead using transfer learning using Inception SSD classification model trained on the COCO dataset<sup>22</sup>. The training parameters included training for 30,000 epochs with 100 steps per epoch and 50 validation steps per epoch. Total training time was 7 hours and 35 minutes. The overall mean average prescription (mAP) was 0.594, mAP for small bounding boxes was 0.62, and mAP for medium bounding boxes was 0.59. The mAP at 0.5 intersection over union (IOU) was 0.99, and the mAP at 0.75 IOU was 0.73. The overall average recall (AR) at 1 maximum detections was 0.64, the AR at 10 maximum detections was 0.65, and the AR at 100 maximum detections was 0.66. The AR at 100 for small boxes was 0.68, and for medium boxes was 0.66. The classification loss for the box classifier was 0.029, the localization loss for the box classifier was 0.027, the localization loss of the RPN loss was 0.005589, the objectness loss of the RPN was 0.105, and the total loss was 0.165.



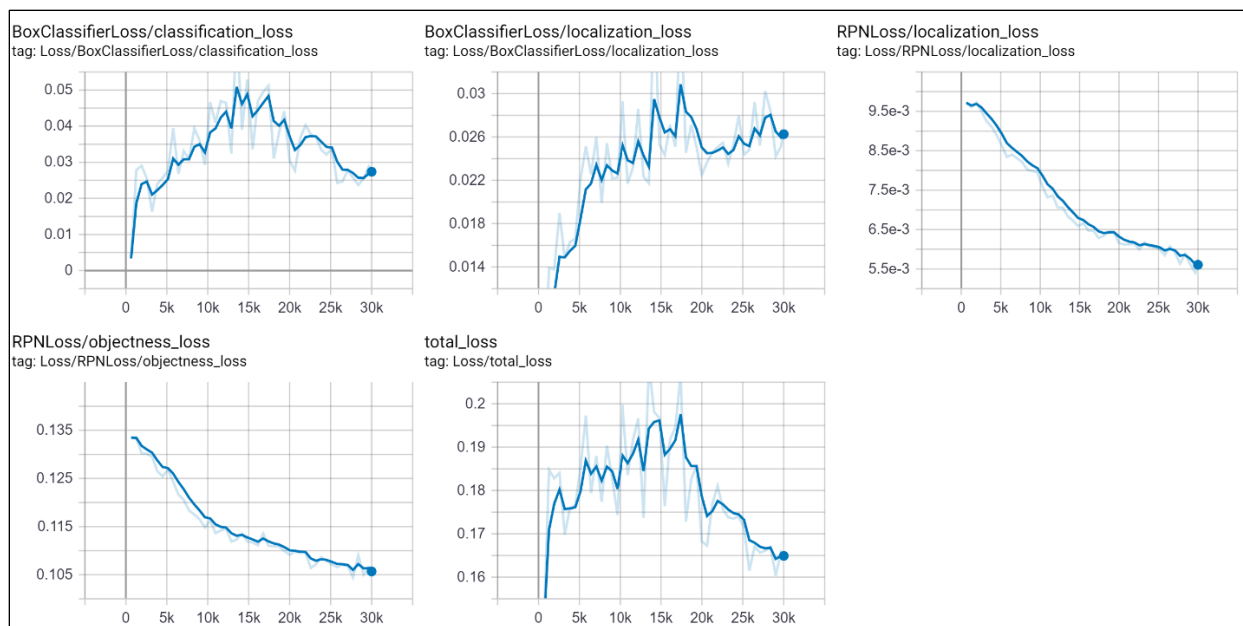
**Figure 11**

Results of training the Faster R-CNN model: The overall mAP was 0.59 after 30,000 epochs (top left). The mAP for small bounding boxes was 0.62 (top middle) and for medium bounding boxes was 0.59 (top right). The mAP at 0.5 IOU was 0.99 (bottom left), and the mAP at 0.75 IOU was 0.73 (bottom right).



**Figure 12**

Results of training the Faster R-CNN model: The overall AR at 1 maximum detections was 0.64 (top left), the AR at 10 maximum detections was 0.65 (top middle), and the AR at 100 maximum detections was 0.66 (top right). The AR at 100 for small boxes was 0.68 (bottom left), and for medium boxes was 0.66 (bottom right).



**Figure 13**

Results of training the Faster R-CNN model: The classification loss for the box classifier was 0.029 (top left), the localization loss for the box classifier was 0.027 (top middle), the localization loss of the RPN loss was 0.005589 (top right), the objectness loss of the RPN was 0.105 (bottom left), and the total loss was 0.165 (bottom middle).

**Table 1**

Evaluation of Mean Average Precision and Intersection Over Union of the Faster R-CNN model


	Overall	Small	Medium	0.5 IOU	0.75 IOU
Mean Average precision	0.594	0.623	0.594	0.66	0.68

**Table 2**

Evaluation of Average Recall of the Faster R-CNN model


	Overall at 1 maximum detections	Overall at 10 maximum detections	Overall at 100 maximum detections	Small boxes at 100 maximum detections	Medium boxes at 100 maximum detections
Average Recall	0.637	0.653	0.661	0.675	0.663

Examples of the validation images passed through the final Faster R-CNN model are shown in **Figures 14, 15, and 16**. The model will not draw a bounding box over a figure with a prediction probability lower than 50%. **Figure 14** shows examples where the model correctly drew bounding boxes with 99% to 100% prediction probability over the passive catheter tip. **Figure 15** shows examples where the model correctly drew bounding boxes over the image but the prediction probability of the bounding box was lower than 99%. **Figure 16** shows examples where the model incorrectly drew bounding boxes and missed drawing a box over a clearly visible catheter.




**Figure 14**

Examples of correctly inferred images on the validation set. In the images where the catheter are visible, bounding boxes were correctly drawn with a high prediction probability (99-100%). In the image where the catheter is out of the slab, a bounding box was not drawn.



**Figure 15**

Examples of correctly inferred images on the validation set, but with less confidence. The bounding boxes are drawn correctly, but was inferred with lower prediction probability (50% or greater).



**Figure 16**

Examples of incorrectly drawn bounding boxes (left, right) and missing bounding boxes (middle) on the validation set.

## Discussion

With a mean average precision (mAP) of 0.59, the model was successful in the fact that it was able to draw boxes correctly on the catheter tip more than half the time. However, in order for the model to be seriously considered for use in a clinical setting, the mAP would need to be significantly higher. The main problem with the model was that the VGG16 classifier, which was trained on the passive and resonant catheter images, did not work when it was added into the

Faster R-CNN model. Further troubleshooting will need to be performed in order to see why the RPN was not working with these layers. One thought is that there might not have been enough images of the resonant catheter for the classifier, even after data augmentation was performed. Another problem may be in the RPN code itself; there are many parameters associated with training the RPN that can be changed, including the stride of the sliding window, number of anchor proposals, and size of the proposed bounding boxes. The training process can be considerably altered even if one parameter is not set correctly.

Another goal that was not accomplished was to run inference on the validation images and measure how quickly the model could detect boxes on a CPU and GPU. This would show that the model is fast enough to be used in a real-time interface when the interventionalist is moving the catheter throughout the patient's body. Additionally, another goal was to write a script that would send an alert when the catheter left the imaging plane on these images. Since there were problems with training the model from scratch, leftover time was spent finding a transfer learning model that could work for this dataset, and training that model while fine-tuning parameters until the best model was found. Hopefully these goals can be met in the future.

Using an object detection model to track catheter tips under iMRI is but one step towards adoption of iMRI. One possible application of the Faster R-CNN would be to not only track a passive catheter, but also expand to tracking resonant catheters. Another future direction would be to expand on the architecture of the object detection and tracking model, for example using the Mask R-CNN model, which draws a bounding box over a detected object and additionally draws a segmentation mask over the object. Another possibility is to track and segment the catheter tip, the entire catheter length, and even other objects, like the guidewires used prior to catheter deployment. Another area in which the model could be used is to provide feedback to

robotic systems for automatic catheterization. Improving the visibility and detection of catheters would make the robotic systems more accurate and safer to use.

Some limitations of the Faster R-CNN model itself includes the lack of a large amount of data and lack of time needed to train the model. Researchers typically spend many months training and retraining models, with all of the dataset collection and preparation already performed. In this case, there was only about three months available to collect the data, process the data, create the model, and train the model. Limitations arising due to application of an object detection model like that of the Faster R-CNN includes a possible lack of adoption of the technology since deployment would necessitate use of additional software and processing power, which many facilities do not currently own.

## **Conclusion**

In this study, we proposed a novel MR-compatible catheter detection model using Faster R-CNN architecture to detect and track the tips of a passive catheter in MR images. The mean average precision was 0.594 and the overall average recall at 100 maximum detections was 0.661. The system would be generalizable to include other types of catheter markers, although the system's overall precision would need to be further studied in order to be used on a clinical basis. It is our hope that this system, when improved, will increase operator confidence and pave the way for adoption of interventional MRI.

## References

1. Saver JL. Time is brain--quantified. *Stroke*. 2006;37(1):263-266.  
doi:10.1161/01.STR.0000196957.55928.ab
2. Virani SS, Alonso A, Benjamin EJ, et al. Heart Disease and Stroke Statistics-2020 Update: A Report From the American Heart Association. *Circulation*. 2020;141(9):e139-e596.  
doi:10.1161/CIR.0000000000000757
3. Kleindorfer D, Lindsell CJ, Brass L, Koroshetz W, Broderick JP. National US estimates of recombinant tissue plasminogen activator use: ICD-9 codes substantially underestimate. *Stroke*. 2008;39(3):924-928. doi:10.1161/STROKEAHA.107.490375
4. Agarwal S, Warburton EA, Baron J-C. From Time is brain to Physiology is brain: a case for reflection in acute stroke treatment decisions. *Brain*. 2015;138(7):1768-1770.  
doi:10.1093/brain/awv120
5. Albers GW, Marks MP, Kemp S, et al. Thrombectomy for Stroke at 6 to 16 Hours with Selection by Perfusion Imaging. *New England Journal of Medicine*. 2018;378(8):708-718.  
doi:10.1056/NEJMoa1713973
6. Saeed M, Hetts SW, English J, Wilson M. MR fluoroscopy in vascular and cardiac interventions (review). *Int J Cardiovasc Imaging*. 2012;28(1):117-137.  
doi:10.1007/s10554-010-9774-1
7. Ozturk C, Guttman M, McVeigh ER, Lederman RJ. Magnetic Resonance Imaging-guided Vascular Interventions. *Top Magn Reson Imaging*. 2005;16(5):369-381.

8. Martin AJ, Lillaney P, Saeed M, et al. Digital Subtraction MR Angiography Roadmapping for Magnetic Steerable Catheter Tracking. *J Magn Reson Imaging*. 2015;41(4):1157-1162. doi:10.1002/jmri.24651
9. Jordan CD, Thorne BRH, Wadhwa A, et al. Wireless Resonant Circuits Printed Using Aerosol Jet Deposition for MRI Catheter Tracking. *IEEE Trans Biomed Eng*. 2020;67(3):876-882. doi:10.1109/TBME.2019.2922879
10. Clogenson H, Dobbelsteen J. Catheters and Guidewires for Interventional MRI: Are We There Yet? *Journal of Imaging and Interventional Radiology*. 2016;2:28.
11. Lundervold AS, Lundervold A. An overview of deep learning in medical imaging focusing on MRI. *Zeitschrift für Medizinische Physik*. 2019;29(2):102-127. doi:10.1016/j.zemedi.2018.11.002
12. Ambrosini P, Ruijters D, Niessen WJ, Moelker A, van Walsum T. Fully Automatic and Real-Time Catheter Segmentation in X-Ray Fluoroscopy. *arXiv:170705137 [cs]*. Published online July 17, 2017. Accessed May 25, 2020. <http://arxiv.org/abs/1707.05137>
13. Yang H, Shan C, Kolen AF, de With PHN. Improving Catheter Segmentation & Localization in 3D Cardiac Ultrasound Using Direction-Fused FCN. *arXiv:190205582 [cs]*. Published online February 14, 2019. Accessed May 28, 2020. <http://arxiv.org/abs/1902.05582>
14. Ren S, He K, Girshick R, Sun J. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In: Cortes C, Lawrence ND, Lee DD, Sugiyama M, Garnett R, eds. *Advances in Neural Information Processing Systems* 28. Curran Associates, Inc.;

- 2015:91–99. Accessed May 29, 2020. <http://papers.nips.cc/paper/5638-faster-r-cnn-towards-real-time-object-detection-with-region-proposal-networks.pdf>
15. Settecase F, Martin AJ, Lillaney P, Losey A, Hetts SW. Magnetic Resonance-Guided Passive Catheter Tracking for Endovascular Therapy. *Magn Reson Imaging Clin N Am*. 2015;23(4):591-605. doi:10.1016/j.mric.2015.05.003
  16. Huang J, Rathod V, Sun C, et al. Speed/accuracy trade-offs for modern convolutional object detectors. *arXiv:161110012 [cs]*. Published online April 24, 2017. Accessed June 3, 2020. <http://arxiv.org/abs/1611.10012>
  17. tomahim. Data augmentation in few lines with skimage. Gist. Accessed July 21, 2020. <https://gist.github.com/tomahim/9ef72befd43f5c106e592425453cb6ae>
  18. Moftakhar P, Lillaney P, Losey AD, et al. New-Generation Laser-lithographed Dual-Axis Magnetically Assisted Remote-controlled Endovascular Catheter for Interventional MR Imaging: In Vitro Multiplanar Navigation at 1.5 T and 3 T versus X-ray Fluoroscopy. *Radiology*. 2015;277(3):842-852. doi:10.1148/radiol.2015142648
  19. Fan Q, Brown L, Smith J. A closer look at Faster R-CNN for vehicle detection. In: *2016 IEEE Intelligent Vehicles Symposium (IV)*. ; 2016:124-129. doi:10.1109/IVS.2016.7535375
  20. Simonyan K, Zisserman A. Very Deep Convolutional Networks for Large-Scale Image Recognition. Published online September 4, 2014. Accessed August 27, 2020. <https://arxiv.org/abs/1409.1556v6>

21. Lin T-Y, Maire M, Belongie S, et al. Microsoft COCO: Common Objects in Context.  
*arXiv:14050312 [cs]*. Published online February 20, 2015. Accessed September 1, 2020.  
<http://arxiv.org/abs/1405.0312>
  
22. Zhang C. *Tony607/Object\_detection\_demo.*; 2020. Accessed August 1, 2020.  
[https://github.com/Tony607/object\\_detection\\_demo](https://github.com/Tony607/object_detection_demo)

## Publishing Agreement

It is the policy of the University to encourage open access and broad distribution of all theses, dissertations, and manuscripts. The Graduate Division will facilitate the distribution of UCSF theses, dissertations, and manuscripts to the UCSF Library for open access and distribution. UCSF will make such theses, dissertations, and manuscripts accessible to the public and will take reasonable steps to preserve these works in perpetuity.

I hereby grant the non-exclusive, perpetual right to The Regents of the University of California to reproduce, publicly display, distribute, preserve, and publish copies of my thesis, dissertation, or manuscript in any form or media, now existing or later derived, including access online for teaching, research, and public service purposes.

DocuSigned by:

*Casey McAntar*

3AD6E8CF3167416...

Author Signature

9/4/2020

Date