One-Click Spine MRI

Alfredo De Goyeneche
HeartVista, Inc.
Los Altos, CA
adg@heartvista.com

Eric Peterson
HeartVista, Inc.
Los Altos, CA

J. Jason He
HeartVista, Inc.
Los Altos, CA

N. Okai Addy
HeartVista, Inc.
Los Altos, CA

Juan Santos
HeartVista, Inc.
Los Altos, CA

1 Context

Spine magnetic resonance (MR) imaging covers over a quarter of the MR imaging market [1] and is the only imaging modality for direct evaluation of the spinal cord [2]. Clinically, it is used to diagnose and monitor recovery for many musculoskeletal pathologies [2; 3]. However, despite the proven clinical impact and many advancements in imaging techniques, one persistent challenge of spine MR data acquisition is to precisely localize the intervertebral discs (IVD) in a consistent and time-efficient manner [4].

IVD localization is challenging because of the long narrow dimensions and highly variable geometry; often involving normal and abnormal curvature in both the sagittal and coronal planes [5]. Technicians manually use a three-plane localizer to prescribe a sagittal stack, and use the sagittal stack to find the angle, position, and extent for each IVD. In a typical lumbar spine MR scan, there are five to seven IVDs. Accuracy is crucial and limits technician efficiency; according to [6], a study of thousands of MR images, almost 75% of IVD acquisitions do not meet diagnostic standards. We believe there is an incredible opportunity to use artificial intelligence to optimize both time between scans and accuracy of IVD axial scans.

While many advancements have been made in applying artificial intelligence to post-processing tasks, very little clinically impactful research towards optimizing MR image acquisition exists because compatibility with existing MR hardware systems is a substantial obstacle. Inspired by existing work to solve the hardware compatibility issue [7], we have created a one-click spine MR imaging software system, using deep learning methods and a vendor-agnostic scanning platform to automate the lumbar sagittal plane and IVD localization to obtain diagnosis-ready data.

The data required to apply AI to the data acquisition process is also far more limited than data for post-processing tasks. Therefore, a major challenged we met was training with limited available data while maximizing generalizability.

2 Methods

Two main components make the one-click automation possible: neural networks for decision-making and a scan control interface to direct their execution on the scanner.

2.1 Networks

There are two automation steps during the axial IVD localization process. The lumbar exam begins with a three-plane localizer, a set of five slices in each of the axial, sagittal, and coronal planes, separated by 10 mm. The localizer is used to assess the position of the patient within the scanner.

3rd Medical Imaging Meets NeurIPS Workshop at the Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

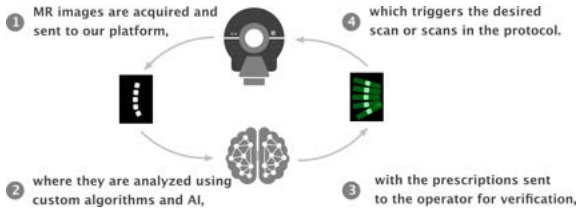

Figure 1: The workflow of the One-Click Scan Control.

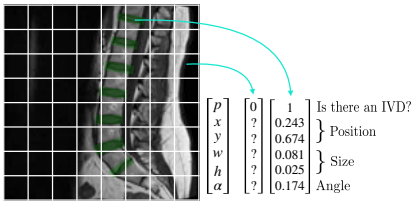

Figure 2: Yolo-based output for the IVD localizer network. Each grid cell has 6 outputs: IVD presence, x and y position, height, width, and angle.

The objective of the first network is to use the three-plane localizer to determine an appropriate multi-slice sagittal stack prescription to cover the lumbar spine. One image from each plane of the three-plane localizer images are concatenated and fed into a CNN based on the Inception ResNet v2 architecture [8]. The outputs of the network were the field of view, number of sagittal slices, and angle with respect to the superior-inferior axis. The training used the momentum optimizer [9] to minimize L1 loss.

Once the sagittal stack has been acquired, the objective of the second network is determining appropriate axial prescriptions for each of the lumbar IVDs. The central image of the sagittal stack is the input to a YOLO-based CNN [10]. The network outputs a 16-by-16 grid, where for each cell the network determines the presence of an IVD, with parameters for the extent, angle, and position relative to the grid cell (Figure 2). Predictions relative to the grid cell are then converted to absolute image coordinates, and up to the top-n most confident proposed IVD boxes are chosen using non-max suppression. Each of the chosen boxes correspond to one IVD axial scan, where the number of slices is based on the height of the box. Cross entropy error was used for training decisions, and mean square error was added for coordinate predictions when an IVD was present in the grid cell. Both losses were minimized using the Adam optimizer [11].

2.1.1 Dataset

We trained both networks on DICOM data of 37 patients from [12], split into 15 training, 4 validation, and 18 test sets. The data included three-plane localizers and T1- and T2-weighted sagittal stack acquisitions. Although our training data was limited, we were able to obtain satisfactory results by augmenting our images with rotations, translations, scaling, clipping, and using non-central sagittal images as the input for the second model. We evaluated our networks quantitatively with data from 18 patients [12] and qualitatively with data acquired *in vivo* with a GE Signa Excite 1.5T MRI Scanner.

2.2 One-Click Scan Control

To run the neural networks during the scan, we developed software to receive DICOM images from the scanner, process the images, verify the new prescriptions, and send prescription commands back to the scanner. It consists of a settings file which defines the scan protocol and network selection, a DICOM listener which listens to all DICOM images generated by the scanner, an interface to display the resulting prescription, and a scanner control module. The entire process is triggered by running a three-plane localizer on the scanner, which begins the chain of processing. This workflow is illustrated in Fig. 1 above.

This allows for fast and automated prescription generation within 10 seconds of the previous scan completion. The scanner can immediately continue with the next scan, or the user can verify and modify the prescription before continuing.

3 Experiments & Results

First, we tested the performance of our neural network models on our test data of 18 patients from [12]. Fig. 3 depicts results of our IVD localizer model for qualitative review. For quantitative analysis,

we calculated the Dice similarity coefficient for the lumbar IVDs and mean absolute error (MAE) of the discs predictions for the 18 labeled examples. The quantitative results are shown in Table 1

	X Position (mm)	Y Position (mm)	Width (mm)	Height (mm)	Angle ($^\circ$)	Dice Score
MAE	3.32 ± 2.11	3.89 ± 2.28	3.96 ± 2.48	1.97 ± 1.44	2.92 ± 2.31	0.839 ± 0.037

Table 1: Mean squared error (MSE) and Dice score for IVD axial predictions on data from [12].

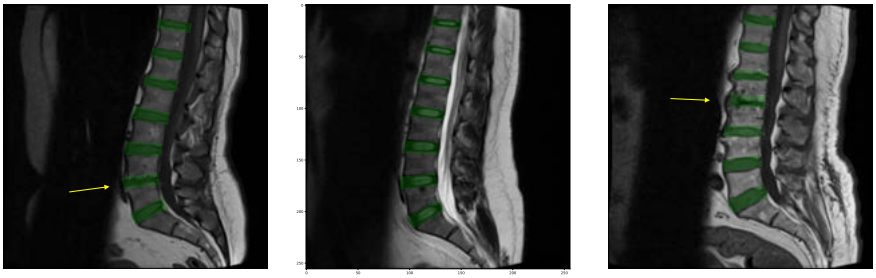

Figure 3: Qualitative prediction examples for the IVD localizer model on data from [12]. The arrows indicate damaged discs.

Finally, we tested the proposed full, one-click process on a GE Signa Excite 1.5T Scanner. The scan protocol consisted of a three-plane localizer scan, a fast spin echo (FSE) T1 fluid-attenuated inversion recovery (FLAIR) sagittal scan, and multiple FSE T1 axial scans. The scanner control software was installed on a laptop next to the scanner console and connected via a network cable to the scanner's network. The *in vivo* results are shown in Figure 4 for qualitative review.

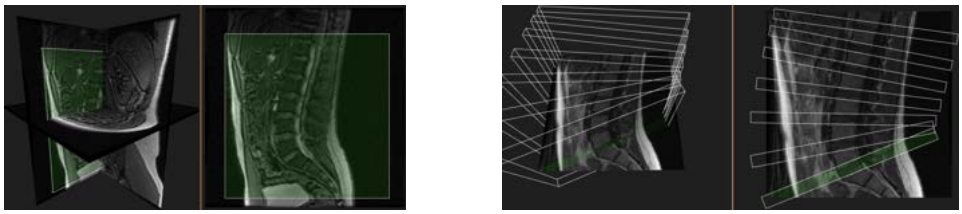

Figure 4: Example of *in vivo* sagittal stack and IVD localization prescriptions.

4 Discussion

From our retrospective experiments on test data, we saw that, although our training set was limited in the variety of pathologies, the model proved to be robust and successfully found discs with irregular geometries, as indicated by the arrows in Figure 3.

In the future, we can evaluate our prospective one-click pipeline and neural networks on a 3T scanner, since that is the preferred field strength for certain spine pathologies [5]. Since our networks were trained on data with a variety of contrasts and augmentations, we believe both networks should generalize well to 3T data. Further evidence of this comes from our prior AI-assisted cardiac MR imaging work, which works at both field strengths. From an anatomical perspective, further work can also be done to test the generalizability of our models on cervical and thoracic spine MR data.

We have shown a software system to utilize deep learning to expedite and increase the accuracy of IVD image acquisition. We have demonstrated our direct interface with GE MR imaging hardware systems to create immediate clinical impact for lumbar spine MR imaging, and our software system is capable of interfacing with scanners from Siemens, too. We propose to evaluate further on different vendor systems. Therefore, our deep learning-based approach to sagittal stack and IVD localization should generalize well to other hardware systems and regions of the spine.

References

[1] P. A. Rinck, *Magnetic Resonance in Medicine*. The Basic Textbook of the European Magnetic Resonance Forum., 11 ed., 2017. Electronic version 11, published 1 June 2017.

[2] J. Elliott, T. Flynn, A. Al-Najjar, J. Press, B. Nguyen, and J. T. Noteboom, "The pearls and pitfalls of magnetic resonance imaging for the spine," *Journal of Orthopaedic and Sports Physical Therapy*, vol. 41, pp. 848–860, 2011 2011.

[3] S. Dutta, A. Bhave, and S. Patil, "Correlation of 1.5 tesla magnetic resonance imaging with clinical and intraoperative findings for lumbar disc herniation," *Asian Spine Journal*, vol. 10, pp. 1115–1121, Dec 2016.

[4] R. Herzog, D. Elgort, A. Flanders, and P. Moley, "Variability in diagnostic error rates of ten mri centers performing lumbar spine mri exams on the same patient within a three week period," *The spine journal : official journal of the North American Spine Society*, vol. 17, 11 2016.

[5] M. I. Vargas, B. M. A. Delattre, J. Boto, J. Gariani, A. Dhouib, A. Fitsiori, and J. L. Dietemann, "Advanced magnetic resonance imaging (mri) techniques of the spine and spinal cord in children and adults," *Insights into Imaging*, vol. 9, pp. 549–557, June 2018.

[6] M. Studin, "Mri spine protocols," *The American Chiropractor*, pp. 54–58, Sep 2013.

[7] J. M. Santos, G. A. Wright, and J. M. Pauly, "Flexible real-time magnetic resonance imaging framework," in *The 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, vol. 1, pp. 1048–1051, Sep. 2004.

[8] C. Szegedy, S. Ioffe, and V. Vanhoucke, "Inception-v4, inception-resnet and the impact of residual connections on learning," in *AAAI*, 2016.

[9] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, "On the importance of initialization and momentum in deep learning," in *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28*, ICML'13, pp. III–1139–III–1147, JMLR.org, 2013.

[10] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *CoRR*, vol. abs/1804.02767, 2018.

[11] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization.," *CoRR*, vol. abs/1412.6980, 2014.

[12] Y. Cai, S. Osman, M. Sharma, M. Landis, and S. Li, "Multi-modality vertebra recognition in arbitrary views using 3d deformable hierarchical model," *IEEE Transactions on Medical Imaging*, vol. 34, p. 1676–1693, Jan 2015.