

4276

# Single board computer as a satellite-linked, deep learning capable pocket MR workstation: a feasibility study

Keerthi Sravan Ravi<sup>1,2</sup>, John Thomas Vaughan Jr.<sup>2</sup>, and Sairam Geethanath<sup>2</sup><sup>1</sup>Biomedical Engineering, Columbia University, New York, NY, United States, <sup>2</sup>Columbia Magnetic Resonance Research Center, New York, NY, United States

## Synopsis

This work shows the feasibility of employing a Raspberry Pi (RPI) single-board computer as a deep learning capable MR workstation. RPI's ability to run a Tensorflow-Lite optimized brain tumour segmentation model is demonstrated. A comparison of data upload across fixed broadband, cellular broadband (LTE, 3G) and satellite terminal methods of internet access is presented. Finally, the setup described in this work is compared with a conventional fixed MRI workstation and a portable MRI workstation.

## Introduction

The number of MR scanners per million people (pmp), referred to as MR scanner density, quantifies and compares access to MR technology across geographies [1]. MR scanner density is heterogeneous globally varying from >35 pmp in high-income countries such as the USA, Japan while it is <1 pmp in underserved countries such as Africa and India. Lack of access to educational tools contributes to inefficient use of the technology. Additionally, the penetration of fixed broadband and/or cellular broadband (3G/LTE) services is low in such regions. However, satellite-linked internet access is available globally (Figure 1). In Autonomous MRI (AMRI) [2], we demonstrated that this shortage of expertise can be augmented by leveraging edge-based and cloud-based deep learning. Recent commercial single-board computers are capable of running deep learning models; one example of such a device is the Raspberry Pi (Raspberry Pi Foundation, UK). In this work, we study the feasibility of using a Raspberry Pi Model 4B (RPI) as an affordable pocket MR workstation, called the AMRI-Pi. The RPI's significantly small form factor allows radiologists to carry only the RPI device and leverage its plug-and-play capabilities to interface with input/output and display peripherals. We demonstrate running a brain tumour segmentation model on-device and present a comparison of data upload times to a cloud service via fixed broadband, cellular broadband and satellite terminal internet access methods.

## Methods

The hardware setup in this work included a Satmodo Explorer 710 (Satmod, USA) terminal for satellite internet access and a Raspberry Pi (RPI) Model 4B as a single board computer. Input/output peripherals included a keyboard, mouse and a monitor. For internet connectivity, three access methods were trialed. These were: wireless fixed broadband, cellular (3G/LTE) and satellite-link (Figure 2). We performed two experiments to demonstrate the feasibility of transforming a RPI into an affordable, open-source MR workstation with deep learning capabilities. In the first experiment, we leveraged an open-source implementation of the fully automatic brain tumour segmentation model developed by Kermi et al [3]. We utilized Tensorflow-Lite [4] to optimize the model's disk space via post-training quantization. Subsequently, we compared its performance with the primary model on the BraTS 2018 dataset [4-9]. The goal of this experiment was to demonstrate running a deep learning model on an edge-device. In the second experiment, the brain tumour segmentation maps (155 slices) were uploaded to the Google Drive cloud-storage service. The times elapsed across wireless fixed broadband (100 Mbps), cellular broadband (LTE, 3G) and satellite internet were measured. Cellular broadband was leveraged via USB tethering to a LTE-capable Android smartphone. Satellite internet was leveraged via wireless access. The satellite terminal was carefully placed with a clear line-of-sight to the sky. This experiment was designed to explore the feasibility of leveraging satellite internet for cloud-connectivity. Finally, we leveraged the free, open-source Aeskulap DICOM viewer ([www.nongnu.org/aeskulap/](http://www.nongnu.org/aeskulap/)) to visualize DICOM images on the RPI. Aeskulap can load DICOM images for review and provides functionalities such as querying and fetching DICOM images from PACS over the network, panning, zooming, scrolling and modifying window-level.

## Results

In the first experiment, the primary deep learning model occupied 38.9 MB disk space whereas the post-training quantized Tensorflow-Lite model occupied 9.73 MB. This was a 75% reduction in disk space. The mean inference time did not vary significantly - increasing from 11.26 seconds per slice on the primary model to 12.69 seconds per slice on the quantized model. In the second experiment, slice-wise uploads of the brain tumour segmentation masks required 1.756±1.368 seconds (fixed broadband), 2.025±2.180 seconds (LTE), 4.786±7.336 seconds (3G) and 19.69±14.72 seconds (satellite-link). A graph illustrating these results is presented in Figure 3. Finally, Figure 4 presents a screenshot of the Aeskulap DICOM viewer utilized for visualizing a brain volume. This DICOM viewer allows basic functionalities such as querying and fetching DICOM images from archive nodes over the network, real-time image manipulation, user-defined window-level presets and supports compressed (lossless, lossy), uncompressed and multi-frame DICOM images. Figure 5 also compares a fixed workstation, a commercially available portable workstation and Raspberry Pi-based portable workstation (this work). The three types of workstations are compared across dimensions such as cost, physical dimensions, power source and functionalities offered. As per the comparison, AMRI-Pi is the most portable, most affordable and least power consuming solution. The constraint on computational power can be complemented by leveraging cloud computing facilities.

## Discussion and conclusion

The RPI has a significantly small form factor and its plug-and-play capabilities allow easy interfacing with input/output and display peripherals. These features transform AMRI-Pi into a highly portable pocket MR workstation allowing radiologists to only carry the RPI device. Figure 5 does not account for the cost of the satellite-terminal rental for the RPI-based workstation. However, future work involves exploring SpaceX's Starlink low-latency broadband mode of internet access for a faster and more reliable connection. Leveraging additional components to enable spectrometer-interfacing capability for the RPI will allow remote MR exams [2]. RPI's software development support can possibly enable the implementation of application-specific deep learning models and image-processing methods. Source code will be made available on request. In conclusion, this work presents a low-cost, deep learning capable MR workstation that can be deployed in resource-constrained geographies.

## Acknowledgements

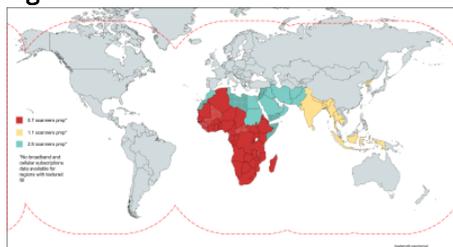
This study was funded [in part] by the Seed Grant Program for MR Studies and the Technical Development Grant Program for MR Studies of the Zuckerman Mind Brain Behavior Institute at Columbia University and Columbia MR Research Center site.

## References

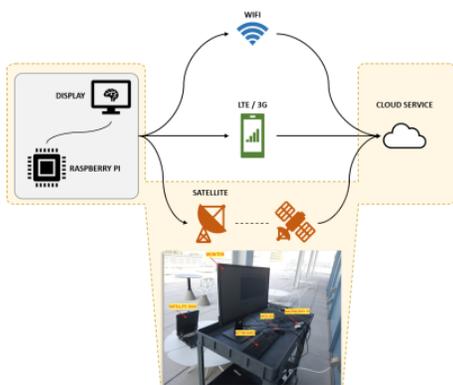
1. Geethanath, Sairam, and John Thomas Vaughan Jr. "Accessible magnetic resonance imaging: A review." *Journal of Magnetic Resonance Imaging* 49.7 (2019): e65-e77.
2. Ravi, Keerthi Sravan and Sairam Geethanath. "Autonomous Magnetic Resonance Imaging." *Magnetic Resonance Imaging* 73 (2020): p177-185.
3. Kermi, Adel, Issam Mahmoudi, and Mohamed Tarek Khadir. "Deep convolutional neural networks using U-Net for automatic brain tumor segmentation in multimodal MRI volumes." *International MICCAI Brainlesion Workshop*. Springer, Cham, 2018.
4. Abadi, Martín, et al. "Tensorflow: A system for large-scale machine learning." 12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16). 2016.
5. Menze, Bjoern H., et al. "The multimodal brain tumor image segmentation benchmark (BRATS)." *IEEE transactions on medical imaging* 34.10 (2014): 1993-2024.

6. Bakas, Spyridon, et al. "Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features." *Scientific data* 4 (2017): 170117.
7. Bakas, Spyridon, et al. "Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the BRATS challenge." *arXiv preprint arXiv:1811.02629* (2018).
8. Bakas, Spyridon, et al. "Segmentation labels and radiomic features for the pre-operative scans of the TCGA-LGG collection." *The cancer imaging archive* 286 (2017).
9. Bakas, Spyridon, et al. "Segmentation labels and radiomic features for the pre-operative scans of the TCGA-LGG collection." *The cancer imaging archive* 286 (2017).

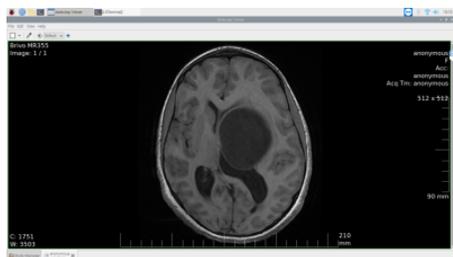
**Figures**



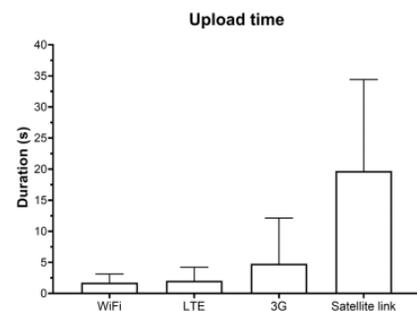
**Figure 1. Motivation for a low-cost, satellite-linked DICOM workstation.** Countries coloured red, yellow and blue have scanner densities of 0.7, 1.1 and 2.8 respectively. MR scanner density is defined as the number of scanners per million people (pmp). These countries in color also have a low number (<25) of broadband subscriptions and/or cellular subscriptions per 100 people. The red dashed line indicates satellite connectivity coverage. Broadband and cellular subscription data are not available for dotted countries.



**Figure 2. Illustration of the three hardware setups discussed in this work.** The Raspberry Pi was used in combination with keyboard/mouse and display monitor input/output peripherals. The three Internet connectivity modalities studied in this work are fixed broadband (WiFi), cellular broadband (3G/LTE) and satellite-link. Picture at the bottom is an example of the hardware setup highlighted in yellow.



**Figure 3. Comparison of upload durations across Internet access methods.** The durations to upload 155 slices of a single pathological brain volume from the BraTS 2018 dataset via fixed broadband, LTE, 3G and satellite-link were recorded. As expected, upload via fixed broadband was fastest ( $1.756 \pm 1.368$  seconds), while satellite-link was slowest ( $19.69 \pm 14.72$  seconds).



**Figure 4.** The remote MR workstation - screenshot of viewing a DICOM image via the Aeskulap medical image viewer. Aeskulap is a Raspberry Pi (RPI)-friendly free to use, open-source DICOM image viewer. It enables viewing reconstructions of acquired data, visualizing output of deep learning models, etc. before being uploaded to the cloud for downstream applications on the RPi.

	Raspberry Pi mobile workstation	Portable MRI workstation	Fixed MRI workstation
Hardware	Single-board computer, monitor, keyboard/mouse	Smart device (example – iPad*)	Workstation, monitor, keyboard/mouse
Cost	\$75	~\$700	\$50,000
Power source	Battery, power outlet	Battery, power outlet	Power outlet
Open-source	Yes	No	No
Dimensions	3.37" x 2.22"	9.87" x 6.85"	Approx. 24" x 8" x 18"
MR functionalities	View DICOM; deep-learning inference	Invoke MR exams; View DICOM; image analysis and visualization	Invoke MR exams; View DICOM; image analysis and visualization

**Figure 5.** Comparing mobile and fixed workstations across six dimensions. The Raspberry Pi based mobile workstation (this work, first column) is the most affordable and most portable. \*An example of a commercially available portable MRI workstation is from Hyperfine (<https://hyperfine.io/>).