

# A Patch-Wise Deep Learning Approach for Myocardial Blood Flow Quantification with Robustness to Noise and Nonrigid Motion

Khalid Youssef, *Member, IEEE*, Bobby Heydari, Luis Zamudio Rivero, Taylor Beaulieu, Karandeep Cheema, Rohan Dharmakumar, Behzad Sharif, *Member, IEEE*

**Abstract**— Quantitative analysis of dynamic contrast-enhanced cardiovascular MRI (cMRI) datasets enables the assessment of myocardial blood flow (MBF) for objective evaluation of ischemic heart disease in patients with suspected coronary artery disease. State-of-the-art MBF quantification techniques use constrained deconvolution and are highly sensitive to noise and motion-induced errors, which can lead to unreliable outcomes in the setting of high-resolution MBF mapping. To overcome these limitations, recent iterative approaches incorporate spatial-smoothness constraints to tackle pixel-wise MBF mapping. However, such iterative methods require a computational time of up to 30 minutes per acquired myocardial slice, which is a major practical limitation. Furthermore, they cannot enforce robustness to residual nonrigid motion which can occur in clinical stress/rest studies of patients with arrhythmia. We present a non-iterative patch-wise deep learning approach for pixel-wise MBF quantification wherein local spatio-temporal features are learned from a large dataset of myocardial patches acquired in clinical stress/rest cMRI studies. Our approach is scanner-independent, computationally efficient, robust to noise, and has the unique feature of robustness to motion-induced errors. Numerical and experimental results obtained using real patient data demonstrate the effectiveness of our approach.

**Clinical Relevance**— The proposed patch-wise deep learning approach significantly improves the reliability of high-resolution myocardial blood flow quantification in cMRI by improving its robustness to noise and nonrigid myocardial motion and is up to 300-fold faster than state-of-the-art iterative approaches.

## I. INTRODUCTION

Dynamic contrast-enhanced cardiovascular MRI (cMRI) is an attractive alternative to nuclear imaging for evaluation of myocardial perfusion as it avoids exposure to ionizing radiation and provides higher resolution images allowing for greater sensitivity for detection of ischemic heart disease [1]. However, the state-of-the-art perfusion cMRI protocols are challenging since they require quantification of absolute myocardial blood flow (MBF) to objectively assess the disease burden. To quantify MBF, time-intensity curves are extracted from dynamic contrast-enhanced cMRI images and are used in a nonlinear fitting process [1-5]. In the setting of high-resolution pixel-wise MBF quantification, the resulting maps are highly sensitive to noise and motion-induced errors. In fact,

conventional deconvolution-based quantification may yield unreliable MBF values in real-world patient studies where the image quality is often subpar in the setting of vasodilator-stress cMRI exams. Recently various groups have addressed this problem by using Tikhonov regularization or Bayesian inference [6, 7] to incorporate prior knowledge in form of spatial/smoothness constraints based on neighboring-pixel similarity in the unknown MBF maps. Such approaches significantly improve the resulting “noise level” in the estimated MBF map but with the practically-important tradeoff of requiring up to 30 minutes per slice due to the need for employing iterative algorithms [6]. Furthermore, such methods cannot enforce robustness to residual nonrigid motion which can occur in clinical stress/rest studies of patients with arrhythmia. While several deep learning approaches have been proposed and successfully applied to myocardial segmentation [3, 4, 8], the quantification of MBF still relies on these iterative solutions of physical and mathematical models. This poses a bottleneck for the full automation of quantitative analysis of perfusion cMRI datasets.

Herein, we present a new end-to-end patch-wise deep learning approach that enables fast and reliable quantification of MBF maps (computational time < 2 seconds per slice). Our proposed approach utilizes a multi-stage neural network architecture trained to quantify pixel-wise MBF values by incorporating information from a patch of neighboring pixels. Furthermore, our approach offers the unique feature of robustness to motion-induced errors, which is achieved in a data-driven manner through extensive data augmentation. Our training dataset was acquired using two different scanners from two medical centers. Our testing dataset includes images acquired using an additional scanner that was not involved in training. Experimental results obtained with numerical phantoms as well as clinical cMRI patient data demonstrate the effectiveness of our approach in the setting of high-resolution MBF mapping.

## II. METHODS

### A. Training/Testing Dataset and cMRI protocol

Stress/rest perfusion images from 180 volunteer patients who underwent vasodilator stress cMRI exam for evaluation of suspected ischemic heart disease were analyzed. Data was acquired using three different 3T MRI scanners from two

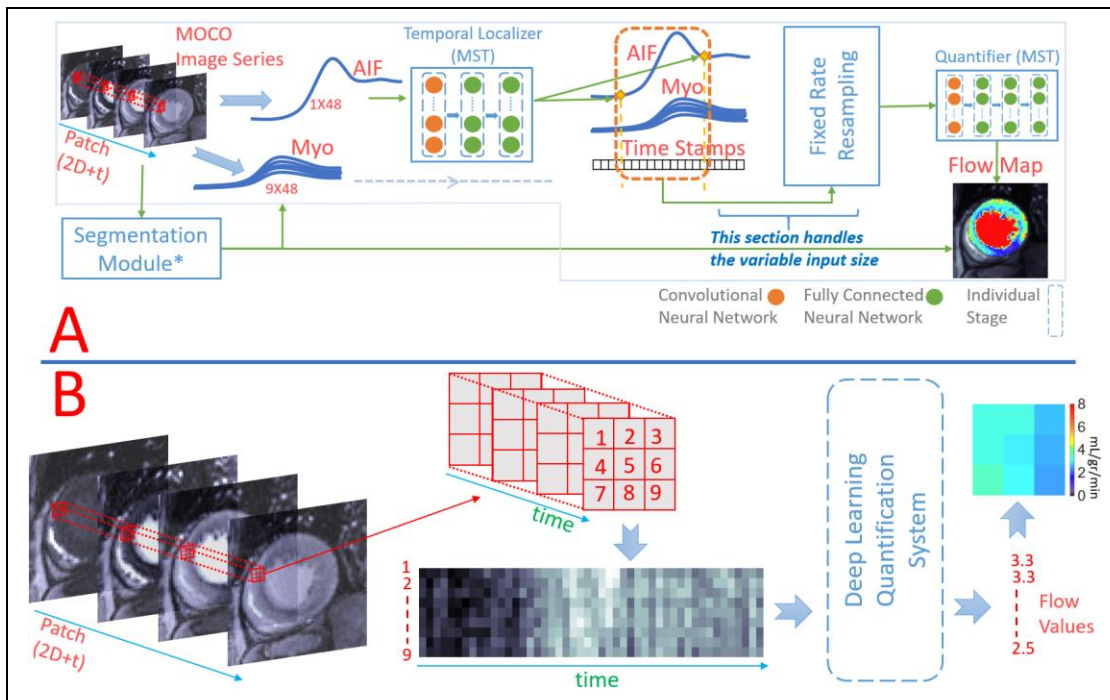
\* Research supported by National Heart, Lung and Blood Institute through National Institutes of Health (NIH) grants R00-HL124323 and R01-HL153430 (PI: B. Sharif) and R01-HL136578 and R01-HL148788 (PI: R. Dharmakumar).

K. Youssef, L. Zamudio, T. Beaulieu, K. Cheema, and B. Sharif are with Cedars-Sinai Medical Center, UCLA Dept. of Bioengineering, and the Laboratory for Translational Imaging of Microcirculation, Krannert

Cardiovascular Research Center, Indiana University School of Medicine, Indianapolis (corresponding author: Behzad Sharif, PhD; e-mail: behzad.sharif@cshs.org, bzsharif@gmail.com).

B. Heydari is with the Cumming School of Medicine, University of Calgary, Alberta, Canada.

R. Dharmakumar is with the Krannert Cardiovascular Research Center, Dept. of Medicine, Indiana University School of Medicine, Indianapolis.



**Figure 1. Processing pipeline.** (A) Patch-wise end-to-end deep-learning approach for quantification of myocardial blood flow from first-pass perfusion cMRI datasets, without the need for using iterative data-fitting. (B) Illustration of the patch extraction step. Each patch (2D+time) contains time curves corresponding to one pixel (patch center) and its eight neighboring pixels. Herein, the input to our end-to-end deep learning system is a perfusion image series. The myocardium and blood pool pixels are identified in each frame using the segmentation module. The arterial input function curve (AIF) is generated using the blood pool pixels and is used by the temporal localizer module to locate the first and last frames of the AIF first-pass. The image series is converted into overlapping spatiotemporal patches each containing the time curves of a center pixel and its neighboring pixels. The time curves in addition to the time stamps are resampled into a fixed length and processed by the quantifier deep learning module which evaluates the blood flow value in the center pixel of the spatiotemporal patch. All slice patches are processed simultaneously in a fast vectorized manner.

different medical centers. For all of the volunteer and patient imaging studies, local Institutional Review Board (IRB) approval and written informed consent was obtained before each imaging exam/study. Data from one of the three scanners was not included in the training data and was only used for further testing of the generalizability of our approach in the testing stage. It is worth mentioning that the total number of images in our training dataset was over 50,000 since each patient study for a complete stress/rest dynamic contrast-enhanced cMRI exam yields over 300 time-resolved images. Our training data augmentation utilized an additional publicly available dataset of 1500 AIF time-curves [9] and, thanks our proposed patch-wise deep learning approach, a total of 10 million training samples were generated as described in Section II.D below.

All subjects underwent high-resolution free-breathing vasodilator-stress cMRI at 3T using two different pulse sequences: either saturation-recovery spoiled gradient-recalled echo with in-plane resolution of  $1.6 \times 1.6 \text{ mm}^2$  or saturation-recovery balanced steady-state free precession with in-plane resolution of  $1.9 \times 1.9 \text{ mm}^2$ . Data was acquired over 48-60 heartbeats in 3 short-axis slices with a slice thickness of 8 mm [10].

### B. Processing Pipeline

Our data-driven approach comprises multiple deep learning modules each specializing in a specific task and incorporates spatial information wherein neighboring pixels contribute to the estimation of myocardial blood flow. The processing pipeline is illustrated in Fig. 1. The input to our

end-to-end deep learning system is a perfusion image series. A segmentation deep learning module developed by our group in previous work [11,12] is used to identify the myocardium and blood pool pixels in each frame in order to extract the myocardial and AIF time curves [1].

A temporal localizer deep learning module is trained to determine the temporal locations of the first and last frames in the first-pass of the AIF time-curve. Only the first-pass frames of the image series are included from this point forward. All time curves are normalized by subtracting the base line and dividing by the area under the curve of the AIF first-pass.

The image series is converted into overlapping spatiotemporal patches (2D + time) each containing the time curves of a center pixel and its neighboring pixels. The time curves are resampled into a fixed length, and together with the time stamps serve as the input of the quantifier deep learning module which is trained to evaluate the blood flow value in the center pixel of the spatiotemporal patch. All patches of a slice are processed simultaneously in a fast vectorized manner yielding a pixel-wise blood flow map as the final output.

### C. Deep Learning Approach

The deep learning modules in our approach utilize the multi-stage training technique (MST). MST was originally designed for high-fidelity denoising of MRI data with non-additive noise [13, 14]. It is a highly distributable structure-parallel deep learning approach typically comprising a system of hundreds of small networks, each specializing in a partial solution of the problem at hand. This allows efficient utilization of Newton based second-order optimization which

offers several advantages in robustness to hyper parameters, accuracy, convergence speed and efficiency [15-21].

A detailed description of MST implementation is presented in [22, 23]. Training is performed in multiple stages, where each stage consists of multiple neural networks. The hierarchical strategy drastically increases the efficiency of training. The burden of reaching a more global solution of a complex model that can perform well on all variations of input data is divided into multiple simpler models such that each simple model performs well only on a subset of the data. Using subsequent stages, the area of specialization of each model is gradually broadened. The process of reaching a more global minimum becomes more feasible at subsequent stages, since models search for combinations of partial solutions of the problem rather than directly finding a complete solution using the original data.

Additionally, the multistage approach allows “very early stopping” [24] at each individual stage where the target error is gradually reduced in subsequent stages. By systematically assigning specific stopping criteria to each stage, we gain a level of control over how fast the overall system fits the data, yielding a better overall performance and generalization. For example, an MST can be designed with few stages where a small error is chosen as a stopping criterion in the first stage and is drastically decreased at successive stages. Alternatively, a larger target error can be chosen and slowly decreased over more stages depending on the complexity of the problem. This data-driven regularization is highly effective and yields a minimal error at the final stage while minimizing the risk of over-fitting.

The MST used for the temporal localizer module is illustrated in Fig. 2(A), and the MST used for the quantifier module is illustrated in Fig. 2(B). Using a Titan W375 workstation with dual AMD EPYC Rome 96 core/192 thread CPU, the time required for training the modules was approximately 20 hours.

#### D. Training Data Augmentation

A time curve extracted from a 48-frame perfusion image series can be defined by a vector:

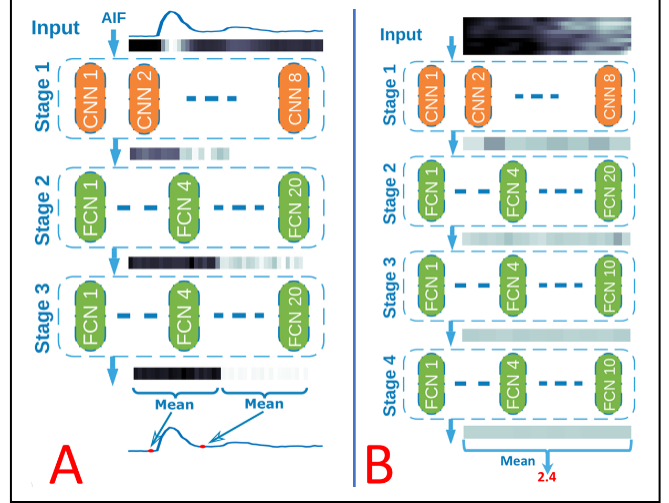
$$C = \{c_1, c_2, c_3, \dots, c_{46}, c_{47}, c_{48}\}$$

where  $c_i$  is the curve value at frame  $i$ . Let  $\tilde{C}$  be the result of modifying  $C$  after an augmentation process. The augmentation types used are defined in Table I. The training data of the temporal localizer module was augmented by adding synthetic noise and shifts as illustrated in Fig. 3. For the quantifier module, robustness to noise and motion induced artifacts were achieved through meticulous data augmentation. The perfusion image series in the training dataset were heavily denoised, and conventional analysis was performed using Fermi-constrained deconvolution of the time-curves to obtain the ground truth labels.

The data was then augmented by systematically adding different levels of synthetic noise to the denoised images, as well as simulated motion errors which manifest as spike noise in the tissue time-curves [25]. Shifts were also added to the myo-patch time curves as illustrated in Fig. 4. A total of 10

million training samples were generated and used for training our quantifier deep learning module, which was optimized to estimate the ground truth labels of the denoised images even when it receives the synthetically distorted images as an input.


Figure 2. Proposed multi-stage training technique (MST) architecture for




the deep-learning modules. (A) Illustration of the temporal localizer architecture. Stage 1 is a convolutional front-end consisting of 8 convolutional neural networks (CNNs). The task for each CNN is to perform temporal localization, i.e. to estimate the start/end truncation indices for the first-pass in the AIF time-curve. Each CNN outputs two real-valued scalars (regressed time indices). The outputs from the 8 CNNs in Stage 1 are concatenated into a 16x1 vector and this forms the input for Stage 2, which consists of 20 fully connected networks (FCNs). The estimation of start/end indices are refined in the subsequent two stages (Stage 2 and Stage 3). Each of the 20 FCNs in Stage 2 produces 2 outputs which are concatenated into a 40x1 vector to form the input to Stage 3. At the output of Stage 3, the first/second outputs from all FCNs are averaged and rounded to obtain the start/end index for the AIF time-curve. The progression of a sample input pixel is demonstrated. The input corresponding to one pixel (an AIF curve) is shown. A gray-scale image corresponding to the input is shown as well. The output of each stage is shown as a gray scale image. The estimated locations of the start/end points are shown at the output. (B) Illustration of the quantifier architecture. The mechanism here is similar to the temporal localizer. The input to this module is a (2D + time) patch. The output of this module is a blood flow value corresponding to the pixel in the patch center. The input to the first stage also includes the arterial input function and time stamps corresponding to the pixels included in the patch (omitted in the figure for visual clarity). The progression of a sample input patch is demonstrated. The output of each stage is shown as a gray scale image. The final stage FCN outputs are averaged to obtain the MBF value.

TABLE I. AUGMENTATION TYPES

Type	Description
(I)	1 frame left: $\tilde{C} = \{c_2, c_3, c_4, \dots, c_{47}, c_{48}, c_{48}\}$
(I)	2 frames left: $\tilde{C} = \{c_3, c_4, c_5, \dots, c_{48}, c_{48}, c_{48}\}$
(I)	1 frame right: $\tilde{C} = \{0, c_1, c_2, \dots, c_{45}, c_{46}, c_{47}\}$
(I)	2 frames right: $\tilde{C} = \{0, 0, c_1, \dots, c_{44}, c_{45}, c_{46}\}$
(II)	$\tilde{C} = C \pm offset$ , a small scalar within 3% of the curve's peak.
(III)	$\tilde{C} = C + noise$ , a vector of gaussian noise with the same size as $C$ .
(IV)	$\tilde{C} = C + spike$ , a vector of spike noise with the same size as $C$ .
(V)	$\tilde{C} = C \times s$ , a scalar value between 0.8 and 1.2



**Figure 3. Training data augmentation for the temporal localizer module.** AIFs were extracted from patient data and the first and last frame locations of the first-pass (foot and valley) were manually labeled to obtain the ground truth. The AIFs are normalized by subtracting the baseline and dividing by the AIF average peak value. The data was augmented by randomly shifting the AIFs and adding noise. Hundreds of random realizations were created for each AIF in order to obtain robustness to noise as well as base line.



**Figure 4. Training data augmentation for the quantifier module.** AIFs and myo-patches of time curves extracted from patient data were placed in separate sets. For each training sample, a combination of an AIF and a myo-patch are randomly selected. The myo-patch time curves are randomly shifted left/right and up/down, randomly scaled and contaminated with noise. Spike noise is added to some of the samples to simulate motion induced errors. Gaussian noise is also added to the AIF curve. The curves are truncated to the first pass using the temporal localizer module and normalized by subtracting the base line and dividing by the AIF first-pass area under the curve. Time stamps are randomly generated based on the range of time durations calculated from our patient dataset. Ground truth labels for each sample were generated using Fermi deconvolution prior to applying the augmentation processes. Samples with a fitting error larger than 20% were considered infeasible and were discarded. A total of 10 million training samples were generated and used for training our quantifier deep learning module.

### III. RESULTS


We validated our approach by comparing it to the established (conventional) approach for MBF quantification, i.e., Fermi-constrained deconvolution [1], through an array of experiments that included real-world clinical patient data as well as numerical phantoms wherein the exact ground truth is known.

#### A. Numerical Experiments

In the first set of experiments, we evaluated the effect of SNR degradation in synthetic data (realistic time-curves generated based on the approach in Section II). As shown in Fig. 5, in high-SNR regime the performance of our proposed patch-wise deep learning is similar to the conventional constrained deconvolution approach. However, our proposed method clearly outperforms the conventional approach as SNR degrades to the realistic levels seen in clinical datasets (below 20 dB). As shown in Fig. 5, we used two different metrics

following the approach described in [26]. Further details are provided in the caption of Fig. 5.

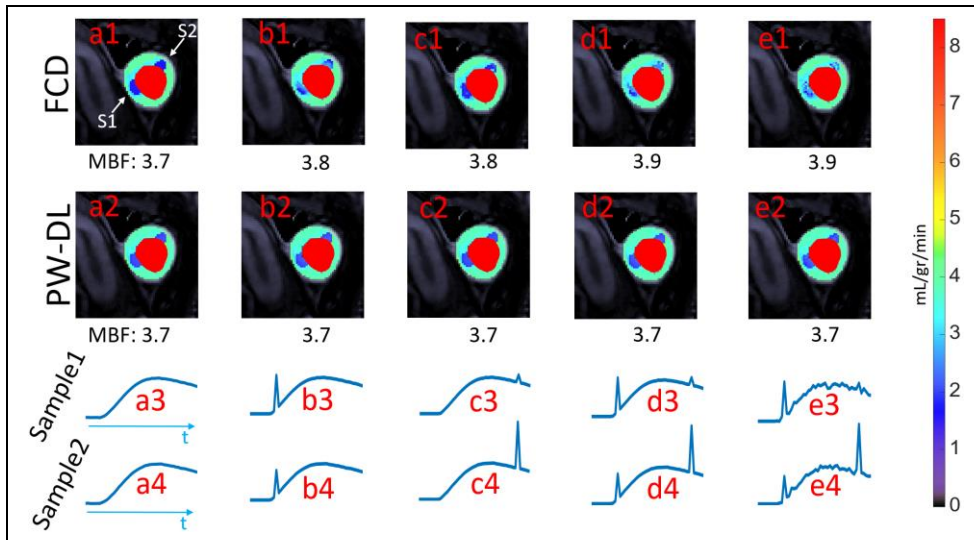
Next, we evaluated the level of robustness our proposed method to residual myocardial motion (errors induced by nonrigid motion) as described in Fig 6, in addition to random noise as described in Fig. 7. A consistent behavior can be observed in both Figs. 6 and 7: our proposed deep-learning approach is capable of “seeing through” the motion errors and random noise, while the conventional approach tends to overestimate MBF values in presence of such motion- or noise-induced time-curve errors. These experiments confirm the ability of our approach to maintain a steadier performance in practical clinical scenarios where the image quality (level of noise and cardiac motion) can vary significantly from patient to patient, or from one medical center to another. It is worth noting that the behavior of the conventional (constrained deconvolution) approach was consistent with previous work [27], where it has been shown that MBF values tend to be overestimated with lower image quality.



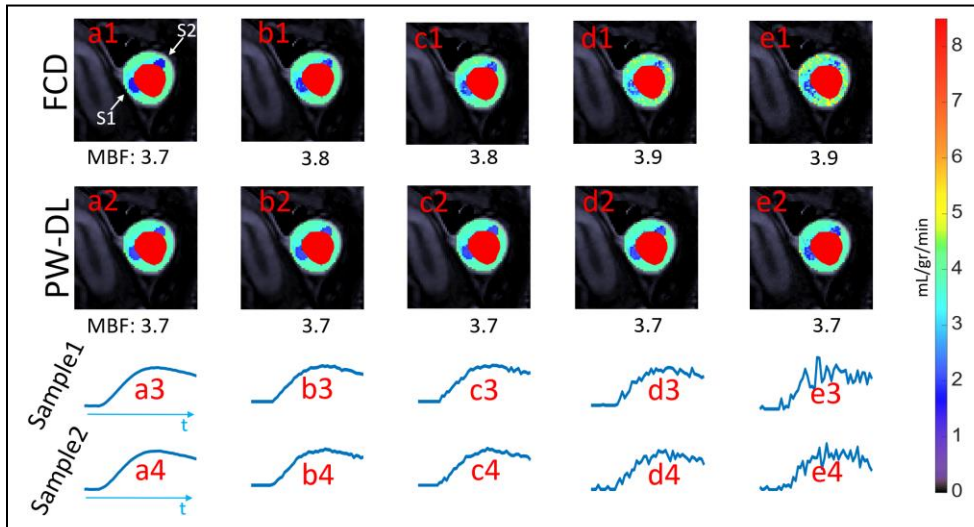
**Figure 5. Comparison of noise effect on performance for our deep learning approach and Fermi constrained deconvolution.** FCD: Fermi Constrained Deconvolution. DL: Patch-wise end-to-end deep learning. Two accuracy metrics are used: (A) Coefficient of repeatability, and (B) Normalized Root Mean Square Error. Each datapoint was evaluated using hundreds of synthetic time curves with known ground truths, with a total of 30 random noise realizations applied to each curve. Both metrics reveal a steadier performance for our approach under low SNR as compared to the conventional approach.

#### B. Clinical Patient Studies

For our in-vivo tests, dynamic contrast-enhanced cMRI data from a total of 29 patients from two different medical centers were analyzed. None of the patients used in the analysis were included in the training dataset. As described in Fig. 8, our results show a high level of agreement versus the conventional Fermi-constrained deconvolution technique for MBF quantification when the goal is relaxed to estimating the mean MBF value for each myocardial slice or myocardial segment (effectively “averaging out” some of the noise-induced errors in the deconvolution process). High agreement between our approach and the conventional method on average MBF values per slice and per segment is observed through Bland-Altman analysis shown in Fig. 8(A), slice-averaged MBF correlation ( $R^2 > 0.9$ ,  $p < 10^{-4}$ ), and American Heart Association (AHA) 6-segment model MBF correlation ( $R^2 > 0.9$ ,  $p < 10^{-4}$ ). However, as demonstrated in Fig. 8(B), the level of agreement between the two methods deteriorates as the number of myocardial sectors (sub-divisions of each myocardial segment) is increased due to the poor robustness of the conventional method in the high-resolution setting to noise/motion effects.



**Figure 6. Demonstration of robustness to motion-induced errors in a numerical phantom.** FCD: Fermi Constrained Deconvolution. PW-DL: Patch wise end to end Deep Learning. S1/S2: Sample time curves extracted from the regions indicated by the white arrows. (A) Results corresponding to the original phantom. (B) Results when a motion-induced error occurs in the upslope. (C) Results when a motion-induced error occurs in the downslope. (D) Result when a motion-induced error occurs both in the upslope and downslope. (E) Results when a motion-induced error occurs both in the upslope and downslope and moderate noise is added. Our method exhibits robustness to motion errors, while the conventional approach produces erroneous results that can lead to misdiagnosis.

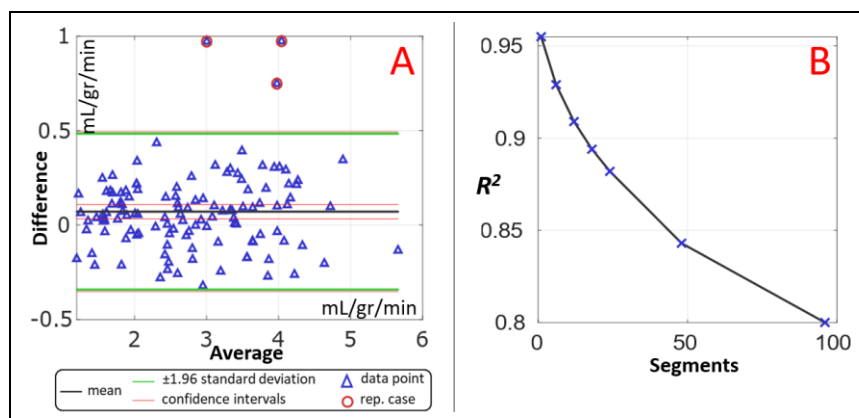


**Figure 7. Demonstration of robustness to noise in a numerical phantom.** FCD: Fermi Constrained Deconvolution. PW-DL: Patch wise end to end Deep Learning. S1/S2: Sample time curves extracted from the regions indicated by the white arrows. (A) Results corresponding to the original phantom. (B-E) Results when different levels of synthetic noise are added. Noise levels increase from left to right. Our methods exhibit graceful degradation with decreasing SNR, while the conventional approach performance deteriorates.

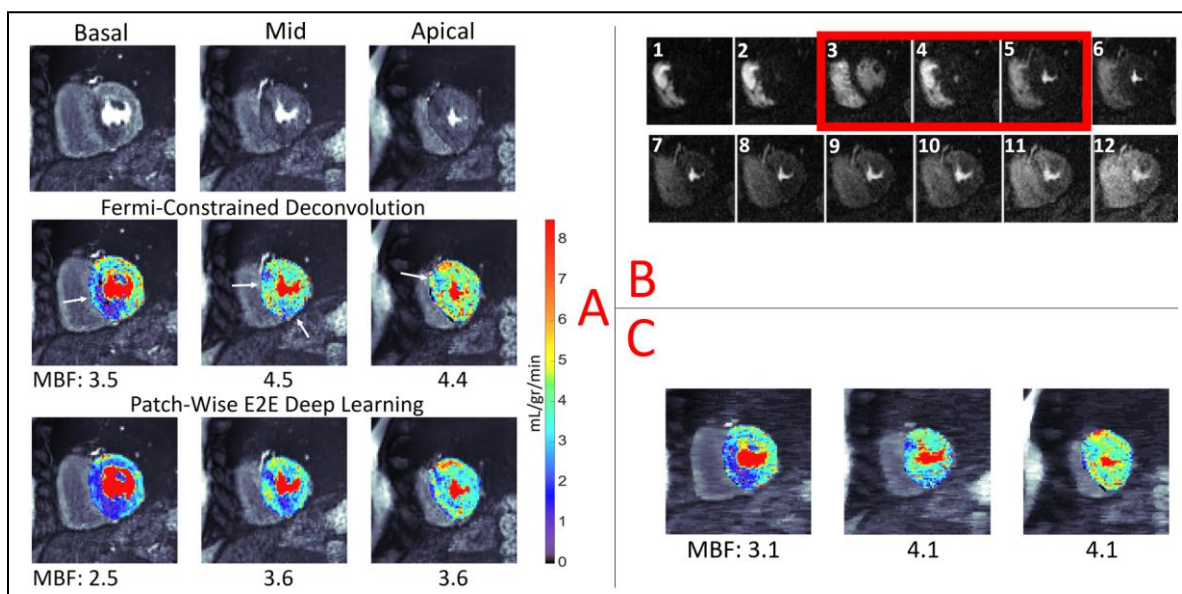
Finally, in Fig 9 we have shown a representative patient study (cMRI with vasodilator stress) in which the MBF maps generated by the conventional approach is compared to those generated by our proposed patch-wise end-to-end deep learning method. Consistent with the numerical results in Figs. 6 and 7, the MBF maps in Fig. 9 show that our method creates a more accurate pixel-wise MBF values which are more consistent with the invasive (gold standard) diagnosis in this patient (see caption). In contrast, the conventional method (even with denoising as shown in Panel C) tends to overestimate MBF values, hence underestimating level of the disease (blood flow deficit).

#### IV. DISCUSSION

In this work, we presented a novel spatially-constrained deep learning approach for quantification of myocardial blood flow from high-resolution dynamic contrast-enhanced cMRI datasets. Since our technique requires less than 2 seconds per myocardial slice (less than 6 seconds per patient) to infer the MBF map, it enables up to 300-fold speedup over recently proposed iterative methods [6,7]. This level of computational speedup has meaningful implications in clinical settings where rapid diagnosis of patients with suspected ischemic heart disease is required. Moreover, our results in clinical



**Figure 8. Comparing high-resolution quantification using our deep learning approach and Fermi-constrained deconvolution in 29 clinical patients.** Data from a total of 29 patients from 2 different medical centers were analyzed. None of the patients used in the analysis were included in the training. Red circles correspond to the representative case from Fig. 9. (A) Bland-Altman plot for average myocardial blood flow (MBF) for each acquired myocardial slice; 97.5% of the samples fall inside the region between the upper and lower limits. (B) Comparison of correlation coefficients ( $R^2$  values) when the myocardium is sub-divided into different number of sectors. As expected, the agreement between the two methods deteriorates as the number of sectors increases due to the poor performance of the conventional method in low-SNR settings.



**Figure 9. Demonstration of robustness to noise and motion-induced errors in a clinical patient scan.** (A) MBF map comparison for a patient with known multi-vessel coronary artery disease based on gold-standard invasive coronary angiography. (B) Adjoining image series of the mid-cavity. ECG mis-triggering highlighted in frames 3-5 causes the conventional approach to underestimate the septal MBF deficit. On the other hand, our deep learning approach is capable of “seeing through” the motion-correction error. (C) Conventional approach results after heavily denoising the image series. The septal MBF deficit is still underestimated due to ECG mis-triggering. This indicates that our deep learning approach can actively detect motion induced errors rather than simply learning to perform a denoising functionality. Furthermore, this example demonstrates the scanner-independence of our approach as it was acquired using a separate scanner that was not included in our training dataset.

studies aimed at high-resolution MBF mapping demonstrated a significant gain over the state-of-the-art Fermi-constrained deconvolution technique in terms of robustness to poor image quality, specifically in the presence of high level of image noise and/or residual myocardial motion. Future work will include testing/validation of the developed framework in large multi-center and multi-vendor studies.

#### REFERENCES

[1] M. Jerosch-Herold, "Quantification of myocardial perfusion by cardiovascular magnetic resonance," *Journal of Cardiovascular Magnetic Resonance*, vol. 12, no. 1, pp. 1-16, 2010.

[2] P. Kellman et al., "Myocardial perfusion cardiovascular magnetic resonance: optimized dual sequence and reconstruction for quantification," *Journal of Cardiovascular Magnetic Resonance*, vol. 19, no. 1, p. 43, 2017.

[3] C. M. Scannell et al., "Deep-Learning-Based Preprocessing for Quantitative Myocardial Perfusion MRI," *J Magn Reson Imaging*, vol. 51, no. 6, pp. 1689-1696, 2020.

[4] H. Xue et al., "Automated inline analysis of myocardial perfusion MRI with deep learning," *arXiv preprint arXiv:1911.00625*, 2019.

[5] D. C. Lee and N. P. Johnson, "Quantification of Absolute Myocardial Blood Flow by Magnetic Resonance Perfusion Imaging," *JACC: Cardiovascular Imaging*, vol. 2, no. 6, pp. 761-770, 2009.

[6] C. M. Scannell, A. Chiribiri, A. D. M. Villa, M. Breeuwer, and J. Lee, "Hierarchical Bayesian myocardial perfusion quantification," *Medical Image Analysis*, vol. 60, 2020.

- [7] J. Lehnert et al., "Pixel-wise quantification of myocardial perfusion using spatial Tikhonov regularization," *Physics in Medicine & Biology*, vol. 63, no. 21, 2018.
- [8] W. Bai et al., "Automated cardiovascular magnetic resonance image analysis with fully convolutional networks," *Journal of Cardiovascular Magnetic Resonance*, vol. 20, no. 1, p. 65, 2018.
- [9] H. Xue et al., "Automated detection of left ventricle in arterial input function images for inline perfusion mapping using deep learning: A study of 15,000 patients," *Magnetic Resonance in Medicine*, vol. 84, no. 5, pp. 2788-2800, 2020.
- [10] Z. Zhou et al., "First-pass myocardial perfusion MRI with reduced subendocardial dark-rim artifact using optimized Cartesian sampling," *Journal of Magnetic Resonance Imaging*, vol. 45, no. 2, pp. 542-555, 2017. doi: 10.1002/jmri.25400.
- [11] Sandoval Z, Van Dyke J, Dharmakumar R, Sharif B. Rapid Automatic Quantification of Myocardial Blood Flow in Free-breathing Myocardial Perfusion MRI without the Need for Motion Correction: A Novel Spatio-temporal Deep Learning Approach. *Proceedings of ISMRM 2019*;27:12.
- [12] Yalcinkaya D, Youssef K, Heydari B, Zamudio L, Dharmakumar R, Sharif B, "Deep Learning-Based Segmentation and Uncertainty Assessment for Automated Analysis of Myocardial Perfusion MRI Datasets Using Patch-Level Training and Advanced Data Augmentation," *Proceedings 43rd IEEE Intern. Conf. Engineering in Medicine and Biology (EMBC)*, In Press.
- [13] K. Youssef, N. N. Jarenwattananon, and L.-S. Bouchard, "Feature-preserving noise removal," *IEEE Transactions on Medical Imaging*, vol. 34, no. 9, pp. 1822-1829, 2015.
- [14] L. Bouchard and K. Youssef, "Feature-preserving noise removal," ed: Google Patents, 2018.
- [15] R. Battiti, "First-and second-order methods for learning: between steepest descent and Newton's method," *Neural computation*, vol. 4, no. 2, pp. 141-166, 1992.
- [16] B. M. Wilamowski and H. Yu, "Improved computation for Levenberg-Marquardt training," *IEEE transactions on neural networks*, vol. 21, no. 6, pp. 930-937, 2010.
- [17] Q. V. Le, J. Ngiam, A. Coates, A. Lahiri, B. Prochnow, and A. Y. Ng, "On optimization methods for deep learning," in *ICML*, 2011.
- [18] H. Yu and B. M. Wilamowski, "Levenberg-marquardt training," *Industrial electronics handbook*, vol. 5, no. 12, 2011.
- [19] G. Montavon, G. Orr, and K.-R. Müller, "Neural networks: tricks of the trade," *springer*, 2012.
- [20] B. M. Ozyildirim and M. Kiran, "Do optimization methods in deep learning applications matter?," *arXiv preprint arXiv:2002.12642*, 2020.
- [21] H. H. Tan and K. H. Lim, "Review of second-order optimization techniques in artificial neural networks backpropagation," in *IOP Conference Series: Materials Science and Engineering*, 2019.
- [22] K. Youssef, L. Bouchard, K. Haigh, J. Silovsky, B. Thapa, and C. Vander Valk, "Machine learning approach to RF transmitter identification," *IEEE Journal of Radio Frequency Identification*, vol. 2, no. 4, pp. 197-205, 2018.
- [23] K. Youssef et al., "Scalable End-to-End RF Classification: A Case Study on Undersized Dataset Regularization by Convolutional-MST," *arXiv preprint arXiv:2104.12103*, 2021.
- [24] R. M. Zur, Y. Jiang, L. L. Pesce, and K. Drukker, "Noise injection for training artificial neural networks: A comparison with weight decay and early stopping," *Medical physics*, vol. 36, no. 10, pp. 4810-4818, 2009.
- [25] C. M. Scannell et al., "Feasibility of free-breathing quantitative myocardial perfusion using multi-echo Dixon magnetic resonance imaging," *Scientific Reports*, vol. 10, no. 1, 2020.
- [26] T. A. Goldstein, H. Zhang, B. Misselwitz, R. G. Gropler, and J. Zheng, "Improvement of quantification of myocardial first-pass perfusion mapping: a temporal and spatial wavelet denoising method," *Magn Reson Med*, vol. 56, no. 2, pp. 439-45, 2006.
- [27] R. McKinley, F. Hung, R. Wiest, D. S. Liebeskind, and F. Scalzo, "A Machine Learning Approach to Perfusion Imaging With Dynamic Susceptibility Contrast MR," *Front Neurol*, vol. 9, pp. 717-717, 2018.