

Unsupervised deep learning with convolutional neural networks for static parallel transmit design: A retrospective study

Toygan Kilic^{1,2}  | Patrick Liebig³  | Omer Burak Demirel^{1,2}  |
Jürgen Herrler³  | Armin M. Nagel^{4,5}  | Kamil Ugurbil² | Mehmet Akçakaya^{1,2} 

¹Electrical and Computer Engineering,
University of Minnesota, Minneapolis,
Minnesota, USA

²Center for Magnetic Resonance
Research, University of Minnesota,
Minneapolis, Minnesota, USA

³Siemens Healthcare, Erlangen, Germany

⁴Institute of Radiology, University
Hospital Erlangen,
Friedrich-Alexander-Universität
Erlangen-Nürnberg (FAU), Erlangen,
Germany

⁵Division of Medical Physics in Radiology,
German Cancer Research Centre (DKFZ),
Heidelberg, Germany

Correspondence

Mehmet Akçakaya, University of
Minnesota, 200 Union Street S.E.,
Minneapolis, MN, 55455, USA.
Email: akcakaya@umn.edu

Funding information

National Science Foundation,
Grant/Award Number: CCF-1651825;
National Institute of Biomedical Imaging
and Bioengineering, Grant/Award
Numbers: P41EB027061, R01EB032830

Abstract

Purpose: To mitigate B_1^+ inhomogeneity at 7T for multi-channel transmit arrays using unsupervised deep learning with convolutional neural networks (CNNs).

Methods: Deep learning parallel transmit (pTx) pulse design has received attention, but such methods have relied on supervised training and did not use CNNs for multi-channel B_1^+ maps. In this work, we introduce an alternative approach that facilitates the use of CNNs with multi-channel B_1^+ maps while performing unsupervised training. The multi-channel B_1^+ maps are concatenated along the spatial dimension to enable shift-equivariant processing amenable to CNNs. Training is performed in an unsupervised manner using a physics-driven loss function that minimizes the discrepancy of the Bloch simulation with the target magnetization, which eliminates the calculation of reference transmit RF weights. The training database comprises 3824 2D sagittal, multi-channel B_1^+ maps of the healthy human brain from 143 subjects. B_1^+ data were acquired at 7T using an 8Tx/32Rx head coil. The proposed method is compared to the unregularized magnitude least-squares (MLS) solution for the target magnetization in static pTx design.

Results: The proposed method outperformed the unregularized MLS solution for RMS error and coefficient-of-variation and had comparable energy consumption. Additionally, the proposed method did not show local phase singularities leading to distinct holes in the resulting magnetization unlike the unregularized MLS solution.

Conclusion: Proposed unsupervised deep learning with CNNs performs better than unregularized MLS in static pTx for speed and robustness.

KEYWORDS

7 T, convolutional neural networks, deep learning, parallel excitation, RF inhomogeneity mitigation, unsupervised learning

1 | INTRODUCTION

Ultra-high field (UHF) MRI offers higher SNR^{1,2} with potential for better spatiotemporal resolutions and has found extensive use in a numerous applications including functional neuroimaging.^{3,4} However, the usage of higher static magnetic fields (B_0) leads to a RF magnetic (B_1^+) field with a lower wavelength, which then becomes comparable to the size of the imaging targets.⁵ This results in greater non-uniformity of the flip angle distributions at UHF.⁶ Image artifacts are caused by this non-uniformity and must be eliminated for better image evaluation. B_1^+ fields are subject-dependent and therefore subject-specific correction is required.^{7,8}

Parallel transmission (pTx) has been instrumental for handling B_1^+ inhomogeneities at UHF.^{9–13} Prior to scanning, many pTx methods solve an optimization problem to generate the associated RF pulse for a target magnetization magnitude, which builds on a magnitude least squares (MLS) objective function.^{3,9,14,15} However, since this calculation is time-consuming and needs to be done when the patient already lies in the scanner, there is interest in speeding up this process.¹⁶ One such approach is the use of subject-independent methods, such as the universal pulses.¹⁷ In this setup, the optimization problem is solved for a target magnetization over an existing database of B_1^+ maps from distinct subjects. Thus, this strategy aims to produce sufficient B_1^+ homogeneity for the training set of subjects, but its performance may be sub-optimal compared to subject-specific methods.¹⁸ Hybrid methods have been proposed to improve this strategy by taking subject-specific information into account while using universal pulses as a starting point for optimization.^{19,20} More recently, machine learning methods have been proposed as an alternative to these optimization techniques.^{21–23} A specific absorption rate (SAR)-efficient method uses kernelized ridge regression to learn B_1^+ shimming weights.²¹ In another work, a fully-connected neural network is proposed to find B_1^+ shimming weights.²² However, fully-connected neural networks require more tunable parameters than convolutional neural networks (CNNs) of similar depth and also cannot handle different input sizes. Therefore, CNNs are used in a subject-specific pulse design that predicts 2D spatially selective RF (2DRF) pulses for a single channel B_1^+ map.²³ Other deep learning (DL) methods for indirectly enhancing pTx RF pulse design include generative adversarial networks to predict B_1^+ distributions in the head following displacement.²⁴ Multi-channel B_1^+ maps have been utilized for a classification type approach with CNNs²⁵ for better RF pulse initialization. Additionally, reinforcement learning has been

applied to RF pulse design.²⁶ However, to the best of our knowledge, there are no works that use CNNs to directly output RF pulses with multi-channel B_1^+ maps as inputs. Furthermore, the aforementioned methods have been trained using supervised learning, where a reference pTx RF waveform was calculated using computationally costly optimization algorithms across the whole training database.

In this work, we propose a strategy for unsupervised DL for pTx using CNNs with multi-channel B_1^+ maps as input. We evaluate its feasibility in static pTx design at 7T. The proposed method outperforms the traditional unregularized MLS method quantitatively in terms of RMS error (RMSE), and coefficient of variation (CoV). Our method has comparable excitation energy demand. Furthermore, the proposed method improves upon unregularized MLS in reducing nulls in the flip angle maps.

2 | METHODS

2.1 | Imaging data and pre-processing

B_1^+ maps were obtained at University Hospital (Universitätsklinikum) in Erlangen, Germany, with approval from the local Ethical Review Board. In vivo brain imaging was performed on 143 healthy subjects utilizing a 7T MAGNETOM Terra whole-body MR system (Siemens Healthcare, Erlangen, Germany) with an 8Tx/32Rx head coil (Nova Medical, Wilmington, MA), which led to 3824 2D sagittal slices that were used for the training database.

Several standard pre-processing steps were applied to B_1^+ maps to preserve lower and peripheral regions (including the upper cervical spine, skull, nose, and jaw) and excluding erroneous measurements which correspond to low SNR, such as in the presence of air, or by phase differences or motion between the prepared and unprepared gradient echo sequences acquired during B_1^+ mapping.²⁰ First, a binary mask was generated for each sagittal slice by interpolating it $\mathbf{W}_{\text{ind}}$ to a target FOV with a normalized threshold of 0.5. Then, a universal binary mask, $\mathbf{W}_{\text{uni}}$, was obtained to eliminate possible outlier effect from B_1^+ maps by thresholding the sum of the individual masks with a normalized threshold of 0.4. $\mathbf{W}_{\text{ind}}$ and $\mathbf{W}_{\text{uni}}$ are binary matrices representing all voxels.²⁰ Concurrently, the B_1^+ maps were spatially interpolated such that all subjects have the same FOV = $156 \times 176 \text{ mm}^2$ and in-plane resolution = $4 \times 4 \text{ mm}^2$. Finally, B_1^+ maps were normalized by the 98.5th percentile among all B_1^+ maps from all subjects to avoid possible outlier effects.²⁰

2.2 | Proposed processing for multi-channel B_1^+ maps and unsupervised training

We propose to concatenate multichannel B_1^+ maps, $B_1^+(x, y, c)$, along a spatial (y) dimension to yield 2D data, $B_{1c}^+(x, y)$, transforming the problem for shift-equivariant processing, amenable to CNNs. Once this input is generated, the real and imaginary parts of the complex maps are given as different channels leading to two input channels. The network itself is a feed-forward CNN, depicted in Figure 1. Convolutions and max-pool operations use 5×5 and 2×2 kernels, respectively. ReLU is utilized for activation. The real and imaginary parts of the complex shimming weights for each coil are obtained at the output of this network.

The network is trained in an unsupervised manner with RMSE loss:

$$J = \frac{1}{m} \sum_{i=1}^m \sqrt{\frac{\| |\mathbf{A}^{(i)} \mathbf{b}^{(i)} | - \mathbf{a}^{(i)} \|^2_{\mathbf{W}^{(i)}}}{N^{(i)}}} \quad (1)$$

where i denotes the training sample, m is the number of datasets in the database, $\mathbf{A}^{(i)}$ is the system matrix generated via the B_1^+ maps for each transmit coil at each

spatial location vectorized as column vectors, $\mathbf{a}^{(i)}$ is the target flip angle map, the diagonal matrix $\mathbf{W}^{(i)}$ obtained by elements-wise multiplication of $\mathbf{W}_{ind}^{(i)}$ and $\mathbf{W}_{uni}^{(i)}$. $N^{(i)}$ is the number of elements in $\mathbf{W}^{(i)}$, and $\mathbf{b}^{(i)}$ is the B_1^+ complex shimming weights for each coil, that is the output of the neural network. The weighted norm $\| |\mathbf{A}^{(i)} \mathbf{b}^{(i)} | - \mathbf{a}^{(i)} \|^2_{\mathbf{W}^{(i)}}$ is defined as $(|\mathbf{A}^{(i)} \mathbf{b}^{(i)} | - \mathbf{a}^{(i)})^T \mathbf{W}^{(i)} (|\mathbf{A}^{(i)} \mathbf{b}^{(i)} | - \mathbf{a}^{(i)})$.

The database was randomly split into 80% training, 10% validation, and 10% testing. The CNN was trained using stochastic gradient descent (SGD) with a learning rate of 10^{-3} with batch size of 1 over 100 epochs. All training and testing were performed using PyTorch and processed on a workstation with the following specifications: Intel E5-2640V3 CPU (2.6GHz and 256GB memory) and an NVIDIA Tesla V100 GPU with 32GB memory. Implementation of the proposed method will be provided online (<https://imagine.umn.edu/research/software>).

2.3 | B_1^+ shimming experiments

Target flip angle maps, $\mathbf{a}^{(i)}$, were obtained by multiplying the $\mathbf{W}^{(i)}$ with desired flip angle α . Hard constraints were not considered. Soft constraints or regularization terms were also not included to avoid confounding factors due to




FIGURE 1 Flowchart of the proposed method. B_1^+ maps of different coils are concatenated in the y dimension for shift-equivariant processing. Real and imaginary parts are concatenated in the channel dimension, as usual. The neural network layers and corresponding layer output sizes are depicted.

tuning of weight hyperparameters. Thus, only the unregularized MLS objective function was used. For this study, the target flip angle was arbitrarily chosen as $\alpha = 5.7^\circ$, lending itself to the small flip angle regime.

For comparison, the unregularized MLS optimization problem,⁹ formulated as $\min_{\mathbf{b}^{(i)}} \left\| |\mathbf{A}^{(i)} \mathbf{b}^{(i)}| - \alpha^{(i)} \right\|_{\mathbf{w}^{(i)}}^2$, was solved separately for each slice i from the database to find the complex shimming weights for 100 distinct initializations with magnitudes ranging from $5.7/\sqrt{8}$ to $5.7 \cdot \sqrt{8}$, and with phases spanning $-\pi$ to π , along with the default CP initialization in the vendor-supplied software implemented in MATLAB. Additionally, to provide a fair assessment of inference times, a PyTorch version of MLS was also implemented based on the MATLAB code.

The methods were quantitatively evaluated for each slice using

$$\text{RMSE} = \sqrt{\frac{\left\| |\mathbf{A}^{(i)} \mathbf{b}^{(i)}| - \alpha^{(i)} \right\|_{\mathbf{w}^{(i)}}^2}{N^{(i)}}} \quad (2)$$

$$\text{Coefficient of variation (CoV)} = \frac{\text{std}\left(|\mathbf{A}^{(i)} \mathbf{b}^{(i)}|\right)}{\text{mean}\left(|\mathbf{A}^{(i)} \mathbf{b}^{(i)}|\right)} \quad (3)$$

and energy demand

$$E_{\text{total}} = \frac{\mathbf{b}^H \mathbf{b}}{a^2}. \quad (4)$$

Statistical differences in RMSE, CoV, and E_{total} were assessed for normality using a one-sample Kolmogorov-Smirnov test with a significance level of 5%. Subsequently, a paired t-test was used for Gaussian distributions, while a Wilcoxon signed-rank test was employed for non-Gaussian distributions, both with a significance level of $p < 0.05$.

Additionally, an analysis of the null problem in flip angle maps^{21,27} was also made. In order to detect nulls, that is areas of severely low flip angle, the following steps were applied: (1) A binary mask was generated by thresholding the normalized flip angle maps above level t_d , which was empirically set to 0.6 (Figure S1). (2) Difference images were calculated between this mask and the normalized target magnetization, which is 1 in all the pixels of interest and 0 otherwise. (3) Image opening, which performs image erosion followed by dilation, was applied to remove spurious pixels in the difference image. (4) The number of nonzero pixels in this image was used as a surrogate for the number of pixels corresponding to nulls, with 0 corresponding to no nulls in the flip angle maps.

Finally, to test the robustness and performance of the proposed training strategy to the size of the training database, a stress test was conducted. To this end, the model's performance was evaluated using varying numbers of training samples on an approximately logarithmic scale {1, 2, 6, 15, 35, 86, 211, 514, 1254, 3059}. These training samples corresponded to {1, 1, 1, 1, 2, 3, 9, 21, 48, 115} subjects, respectively.

3 | RESULTS

Figure 2 depicts best, median, and worst-case flip angle maps, assessed through the RMSE metric, for both the MLS and the proposed method. Figure 2A shows the best, median and worst cases for MLS, along with the corresponding DL images. These have RMSE values of 0.49° , 0.94° , and 1.52° for MLS, and 0.53° , 0.61° , and 0.46° for DL, respectively. The corresponding CoV values are 0.08, 0.17, and 0.27 for MLS and 0.09, 0.11, and 0.08 for DL. Similarly, Figure 2B shows the best, median and worst cases for DL, along with the corresponding MLS images. The corresponding RMSE values are 0.88° , 1.26° , and 0.63° for MLS, and 0.33° , 0.71° , and 1.11° for DL, respectively, while the CoV values are 0.12, 0.16, and 0.11 for MLS and 0.06, 0.13, and 0.20 for DL. We also note that the inference times for MLS and proposed DL method, both implemented in PyTorch, among test samples were 129.47 ± 93.25 and 0.75 ± 0.04 ms for single initialization,




FIGURE 2 Example of best, median, and worst-case flip angle maps (based on the RMSE metric) for the MLS (A) and proposed method (B). The corresponding results for each image are also provided below for comparison.




FIGURE 3 CoV, RMSE, and energy demand results for MLS and the proposed method are shown in the violin plots. According to these plots, the proposed method performs better than the MLS method in terms of RMSE and CoV, highlighting improved field homogeneity, while having similar energy consumption. * indicates that results are significant with $p < 0.05$.

respectively. The training process itself, over the dataset of 3059 samples, required approximately 11 min for 100 epochs for the proposed method.

Figure 3 shows violin plots for the RMSE, CoV, and energy demand metrics. The average RMSE, CoV, and energy demand across test subjects for the proposed method and the best cases for all initializations from MLS for $\alpha = 5.7^\circ$ are $0.94^\circ \pm 0.24^\circ$ (mean $\pm$ std), 0.14 ± 0.03 , and 17.22 ± 9.89 for MLS, and $0.71^\circ \pm 0.20^\circ$, 0.13 ± 0.04 and 17.82 ± 5.84 for the proposed method, respectively. The differences in RMSE, CoV, and energy demand are statistically significant ($p < 0.05$). We note that the energy demand was not Gaussian-distributed, hence the Wilcoxon signed-rank test was used. Further evaluation of the correlations among these quantitative measures, including COV, RMSE, and energy demand, for both MLS and proposed method are provided in Figure S2 using a correlation matrix plot.

Figure 4 depicts representative images that highlight an instance of the nulls appearing in flip angle maps generated using the MLS technique. The proposed method with unsupervised training is able to mitigate this issue. Quantitative assessment reveals that a null is observed with a 13.84% and 4.70% frequency for MLS and the proposed method, respectively, among testing datasets.

Figure 5 illustrates results of the stress test conducted on the proposed method by varying the number of training samples on an approximately logarithmic scale. While the network underperforms with a very limited number of samples as expected, its performance starts to improve

with as few as 86 samples, corresponding to three subjects. With only 514 training samples, corresponding to 21 subjects, it matches the performance of the full database consisting of 3059 samples from 115 subjects. These findings highlight the efficiency of the proposed DL training in this scenario, as it achieves noteworthy results with a substantially reduced number of training samples, making it a promising solution even in resource-constrained scenarios.

4 | DISCUSSION

In this study, we proposed an unsupervised DL method with CNNs for multi-channel pTx design. We demonstrated its feasibility through static B_1^+ shimming at 7T, employing a slice-by-slice shimming technique,^{28,29} which offered an additional degree of freedom compared to regional approaches.³⁰ Our method enhanced the uniformity of flip angle profiles quantitatively, displaying favorable results compared to the unregularized MLS technique in terms of RMSE and CoV metrics while maintaining a comparable energy consumption level.

DL methods have received attention for speeding up pTx design while maintaining subject-specific processing. However, to the best of our knowledge, previous methods for directly estimating RF pulses did not handle multi-channel B_1^+ maps when using CNNs. A previous work that used neural networks with multi-channel maps relied on fully-connected neural networks,²² which




FIGURE 4 Representative flip angle (FA) for MLS and proposed methods showing the null problem in the MLS method. The difference images are produced by subtracting the target magnetization, as described in (2), from the solution obtained using a threshold. This figure shows that the proposed method successfully finds complex shimming parameters so that there are no significantly large nulls in the resulting image. The green circles on the obtained MLS solution point the null which does not exist in the proposed method. The proposed method is close to the desired magnetization with $\text{CoV} = 0.11$, $\text{RMSE} = 0.60^\circ$ compared to the MLS solution with $\text{CoV} = 0.26$, $\text{RMSE} = 1.47^\circ$.




FIGURE 5 A stress test to evaluate the performance of the neural network trained on different numbers of samples on an approximately logarithmic scale. Interestingly, the network starts performing reasonably with as few as 86 samples and achieves comparable performance to the full database with only 514 samples.

require more parameters and cannot work across different input sizes, as opposed to CNNs. On the other hand, previous DL approaches that used CNNs for RF pulse estimation have used single coil B_1^+ maps,²³ and concatenation along the channel dimension with B_0 maps. Here, building on our insights from multi-slice image reconstruction,^{31,32}

we concatenated multichannel B_1^+ maps along a spatial dimension, which enabled shift-equivariant processing suitable for CNNs. We also note that multi-coil B_1^+ input has been used for a classification-type task previously.²⁵

Supervised training has been popular for applications of DL to pTx.^{21–24} However, this necessitates additional

computational resources and time to find reference RF waveforms. In this study, this extra computation is eliminated via unsupervised learning, speeding up training. Furthermore, supervised training implicitly ties the performance of the trained network to the conventional method used to generate the reference RF waveform. In particular, generating a reference label in this study involves solving a nonconvex optimization problem based on B_1^+ maps, specifically the MLS method. Consequently, supervised training would learn to predict the MLS solution, albeit using a CNN for improved efficiency. Thus, the performance of supervised learning would be inherently limited by the MLS solutions. On the other hand, unsupervised learning bypasses this issue by directly minimizing the difference to target magnetization through a loss function. By avoiding reliance on MLS solutions as reference, unsupervised learning holds potential for superior performance in B_1^+ shimming. This is similar to the use of DL in image reconstruction when acquisition of ground-truth data is infeasible, in which unsupervised methods outperform other conventional methods such as compressed sensing or parallel imaging.^{32–34} This strength of unsupervised learning is best exemplified in its ability to tackle the null problem associated with MLS.²⁷ Better results are obtained using the proposed unsupervised learning, whose performance is not limited by the conventional approach that is used to generate a reference RF waveform, since the latter itself suffers from this issue. Conversely, it is worth noting that, while unsupervised learning incorporates physics information in the loss function, it does not solve an explicit objective function as the MLS algorithm, making it potentially more challenging to interpret, particularly for failure modes.

One of the drawbacks of the proposed method is the need for a large training database. In vivo acquisition of such data containing more than a thousand slices is time-consuming and resource-intensive. While the results of the stress test on the number of training samples suggest fewer datasets may be sufficient, this may still put constraints on implementing similar DL techniques for new acquisition schemes. We note that synthetic B_1^+ maps have also been proposed in the literature, which would eliminate the curation of B_1^+ map databases.^{23,35} We also note newer efforts for circumventing the acquisition of B_1^+ maps in the pTx pipeline by estimating it from B_1^- maps.³⁶ The synergistic combination of our approach with such techniques warrants further investigation.

During initial experiments, several different network architectures of varying depths, including residual networks and U-nets, were studied. However, their performance did not substantially differ from the simple feed-forward CNN employed in this study. As a result, the current CNN architecture was chosen due to its expedited

training speed and inherent simplicity. The 100-to-200 epochs were determined to be sufficient for training (Figure S3). Opting for the lower number reduced the training time and resources.

Our current implementation does not have explicit constraints on power, voltage, or SAR. Among these, voltage constraint is the simplest to enforce through the addition of an appropriate nonlinear activation at the network output, such as sigmoid or ReLU6, as in.³⁷ However, power or SAR constraints are substantially more difficult. To the best of our knowledge, inclusion of hard quadratic constraints with DL methods is an open question, although progress has been made in the broader artificial intelligence community.^{38–43} One step in this direction would be to use a soft constraint on power or SAR by incorporating an additional term into the loss function. However, the performance and robustness will substantially depend on tuning the weight for this soft constraint term. Nonetheless, hard constraints on SAR or power remain an open question that is critical for pTx problems.

The results in the article used a target $\alpha = 5.7^\circ$, which is in the small flip-angle regime, where the forward operator is described by a linear system matrix $\mathbf{A}$. To make the model applicable to higher flip angles, a Bloch simulation-based system function needs to be implemented in the loss function. There are existing publicly available implementations in PyTorch⁴⁴ and in MATLAB,⁴⁵ which will be investigated in future studies. We also note that the incorporation of Bloch simulation-based methods in training would increase the total computational time. Thus, further studies are needed for the large flip angle regime.

While our study concentrated on single slice acquisitions, there is also a demand for shimming in non-selective volumetric or multi-slice scenarios.⁴⁶ To this end, we adopted our method to both of these setups by replacing the 2D CNNs with their 3D counterparts. In the first case of volumetric 3D B_1^+ shimming (Figure S4), MLS shows better RMSE and CoV compared to DL, although the difference is non-significant ($p = 0.63$ for RMSE and 0.06 for CoV, Wilcoxon signed-rank test). Furthermore, DL has significantly improved energy demand compared to MLS ($p < 0.05$, Wilcoxon signed-rank test). In the second case of multi-slice imaging with two slices, where 1313 training samples are available, proposed DL significantly improves on MLS in all metrics (Figure S5). Addressing the challenges associated with insufficient training database size⁴⁷ and fine-tuning the CNN model parameters may lead to further improvements for 3D B_1^+ shimming scenarios, although this was beyond the scope of our study.

Throughout this study, testing was retrospectively performed on in vivo data that were acquired previously.

This allowed us to test the methods thoroughly on a large database and isolate cases with adverse failures, such as the null issue in the MLS method, instead of a small number of examples acquired post-training. Thus, an inherent limitation of our study is the lack of new acquisitions to compare the shimming performance of DL and MLS prospectively.

The pTx design pipeline presents time-consuming challenges beyond just RF pulse design, notably in B_1^+ mapping, potentially B_0 mapping, ROI selection, masking, and sequence preparation. Our proposed approach tackles only the RF pulse design part, offering enhanced speed and generalizability, and may potentially find applications for real-time pTx design.⁴⁸

We note that static B_1^+ shimming was used to establish proof-of-concept. In this setting, optimization-based methods are still relatively fast and accurate. Thus, the benefits of DL approaches may not be as pronounced. Nonetheless, further investigations for more complicated pTx pulse designs are warranted to fully harness the potential of our approach. Future studies will aim to extend the work to more complicated pTx pulse designs and incorporate additional constraints, such as SAR or peak power.

5 | CONCLUSIONS

The proposed unsupervised DL strategy for CNNs with image domain concatenated inputs enables fast pTx pulse design, outperforming the unregularized MLS method in terms of homogeneity for static B_1^+ shimming in the 2D small flip angle regime without constraints in this proof-of-principle study.

CONFLICT OF INTEREST STATEMENT

Patrick Liebig and Jürgen Herrler are employees of Siemens Healthineers AG (Germany).

DATA AVAILABILITY STATEMENT

Implementation of the proposed method will be provided online (<https://imagine.umn.edu/research/software>).

ORCID

Toygan Kilic  <https://orcid.org/0000-0001-7969-5738>

Patrick Liebig  <https://orcid.org/0000-0001-7342-3715>

Omer Burak Demirel  <https://orcid.org/0000-0003-4726-0590>

Jürgen Herrler  <https://orcid.org/0000-0002-4620-8216>

Armin M. Nagel  <https://orcid.org/0000-0003-0948-1421>

Mehmet Akçakaya  <https://orcid.org/0000-0001-6400-7736>

REFERENCES

1. Pohmann R, Speck O, Scheffler K. Signal-to-noise ratio and MR tissue parameters in human brain imaging at 3, 7, and 9.4 tesla using current receive coil arrays. *Magn Reson Med*. 2016;75:801-809.
2. Cao Z, Park J, Cho Z-H, Collins CM. Numerical evaluation of image homogeneity, signal-to-noise ratio, and specific absorption rate for human brain imaging at 1.5, 3, 7, 10.5, and 14T in an 8-channel transmit/receive array. *J Magn Reson Imaging*. 2015;41:1432-1439.
3. Padormo F, Beqiri A, Hajnal JV, Malik SJ. Parallel transmission for ultrahigh-field imaging. *NMR in Biomedicine*. Vol 29. John Wiley and Sons Ltd; 2016:1145-1161.
4. Vargas MI, Martelli P, Xin L, et al. Clinical neuroimaging using 7 T MRI: challenges and prospects. *J Neuroimaging*. 2018; 28:5-13.
5. Collins CM, Smith MB. Calculations of B1 distribution, SNR, and SAR for a surface coil adjacent to an anatomically-accurate human body model. *Magn Reson Med*. 2001;45:692-699.
6. Ibrahim TS, Robert L, Abduljalil AM, Baertlein BA, Robitaille PML. Dielectric resonances and B1 field inhomogeneity in UHFMRI: computational analysis and experimental findings. *Magn Reson Imaging*. 2001;19:219-226.
7. Van De Moortele PF, Akgun C, Adriany G, et al. B1 destructive interferences and spatial phase patterns at 7 T with a head transceiver array coil. *Magn Reson Med*. 2005;54:1503-1518.
8. Yang QX, Wang J, Zhang X, et al. Analysis of wave behavior in lossy dielectric samples at high field. *Magn Reson Med*. 2002;47:982-989.
9. Setsompop K, Alagappan V, Gagoski B, et al. Slice-selective RF pulses for in vivo B1+ inhomogeneity mitigation at 7 tesla using parallel RF excitation with a 16-element coil. *Magn Reson Med*. 2008;60:1422-1432.
10. Katscher U, Börner P. Parallel RF transmission in MRI. *NMR Biomed*. 2006;19:393-400.
11. Ullmann P, Junge S, Wick M, Seifert F, Ruhm W, Hennig J. Experimental analysis of parallel excitation using dedicated coil setups and simultaneous RF transmission on multiple channels. *Magn Reson Med*. 2005;54:994-1001.
12. Setsompop K, Wald LL, Alagappan V, et al. Parallel RF transmission with eight channels at 3 Tesla. *Magn Reson Med*. 2006;56:1163-1171.
13. Grissom W, Yip C-Y, Zhang Z, Stenger VA, Fessler JA, Noll DC. Spatial domain method for the design of RF pulses in multicoil parallel excitation. *Magn Reson Med*. 2006;56:620-629.
14. Grissom WA, Khalighi M-M, Sacolick LI, Rutt BK, Vogel MW. Small-tip-angle spokes pulse design using interleaved greedy and local optimization methods. *Magn Reson Med*. 2012;68:1553-1562.
15. Setsompop K, Wald LL, Alagappan V, Gagoski BA, Adalsteinsson E. Magnitude least squares optimization for parallel radio frequency excitation design demonstrated at 7 Tesla with eight channels. *Magn Reson Med*. 2008;59:908-915.
16. Cao Z, Donahue MJ, Ma J, Grissom WA. Joint design of large-tip-angle parallel RF pulses and blipped gradient trajectories. *Magn Reson Med*. 2016;75:1198-1208.
17. Gras V, Vignaud A, Amadon A, Bihan D, Boulant N. Universal pulses: a new concept for calibration-free parallel transmission. *Magn Reson Med*. 2017;77:635-643.

18. Aigner CS, Dietrich S, Schaeffter T, Schmitter S. Calibration-free pTx of the human heart at 7T via 3D universal pulses. *Magn Reson Med*. 2022;87:70-84.
19. Tomi-Tricot R, Gras V, Thirion B, et al. SmartPulse, a machine learning approach for calibration-free dynamic RF shimming: preliminary study in a clinical environment. *Magn Reson Med*. 2019;82:2016-2031.
20. Herrler J, Liebig P, Gumbrecht R, et al. Fast online-customized (FOCUS) parallel transmission pulses: a combination of universal pulses and individual optimization. *Magn Reson Med*. 2021;85:3140-3153.
21. Ianni JD, Cao Z, Grissom WA. Machine learning RF shimming: prediction by iteratively projected ridge regression. *Magn Reson Med*. 2018;80:1871-1881.
22. Mirfin C, Glover P, Bowtell R. Optimisation of parallel transmission radiofrequency pulses using neural networks. International Society for Magnetic Resonance in Medicine. Paris, France. 2018;3388.
23. Vinding MS, Aigner CS, Schmitter S, Lund TE. DeepControl: 2DRF pulses facilitating B1+ inhomogeneity and B0 off-resonance compensation in vivo at 7 T. *Magn Reson Med*. 2021;85:3308-3317.
24. Plumley A, Watkins L, Treder M, Liebig P, Murphy K, Kopanoglu E. Rigid motion-resolved prediction using deep learning for real-time parallel-transmission pulse design. *Magn Reson Med*. 2022;87:2254-2270.
25. Herrler J, Liebig P, Majewski K, et al. *Neural network-supported fast online-customized (FOCUS) parallel transmit (pTx) pulses for slice-selective, large flip angle excitation*. International Society for Magnetic Resonance in Medicine; London, England: 2022;3312.
26. Shin D, Kim Y, Oh C, et al. Deep reinforcement learning-designed radiofrequency waveform in MRI. *Nat Mach Intell*. 2021;3:985-994.
27. Paez A, Gu C, Cao Z. Robust RF shimming and small-tip-angle multispoke pulse design with finite-difference regularization. *Magn Reson Med*. 2021;86:1472-1481.
28. Mao W, Smith MB, Collins CM. Exploring the limits of RF shimming for high-field MRI of the human head. *Magn Reson Med*. 2006;56:918-922.
29. Curtis AT, Gilbert KM, Klassen LM, Gati JS, Menon RS. Slice-by-slice B1+ shimming at 7 T. *Magn Reson Med*. 2012;68:1109-1116.
30. Berrington A, Považan M, Mirfin C, et al. Calibration-free regional RF shims for MRS. *Magn Reson Med*. 2021;86:611-624.
31. Demirel OB, Yaman B, Shenoy C, Moeller S, Weingärtner S, Akçakaya M. Signal intensity informed multi-coil encoding operator for physics-guided deep learning reconstruction of highly accelerated myocardial perfusion CMR. *Magn Reson Med*. 2022;89:308-321.
32. Akçakaya M, Yaman B, Chung H, Ye JC. Unsupervised deep learning methods for biological image reconstruction and enhancement: An overview from a signal processing perspective. *IEEE Signal Process Mag*. 2022;39:28-44.
33. Yaman B, Hosseini SAH, Moeller S, Ellermann J, Uğurbil K, Akçakaya M. Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data. *Magn Reson Med*. 2020;84:3172-3191.
34. Yaman B, Gu H, Hosseini SAH, et al. Multi-mask self-supervised learning for physics-guided neural networks in highly accelerated magnetic resonance imaging. *NMR Biomed*. 2022;35:e4798.
35. Eberhardt B, Poser BA, Shah NJ, Felder J. B1 field map synthesis with generative deep learning used in the design of parallel-transmit RF pulses for ultra-high field MRI. *Z Med Phys*. 2022;32:334-345.
36. Krueger F, Aigner CS, Hammernik K, et al. Rapid estimation of 2D relative B1+-maps from localizers in the human heart at 7T using deep learning. *Magn Reson Med*. 2023;89:1002-1015.
37. Vinding MS, Lund TE. Clipped DeepControl: deep neural network two-dimensional pulse design with an amplitude constraint layer. *Artif Intell Med*. 2023;135:102460.
38. Pan X, Zhao T, Chen M. DeepOPF: Deep Neural Network for DC Optimal Power Flow. 2020.
39. Zhao T, Pan X, Chen M, Venzke A, Low SH. DeepOPF+: A Deep Neural Network Approach for DC Optimal Power Flow for Ensuring Feasibility. 2020. <https://arxiv.org/abs/2009.03147>
40. Huang W, Pan X, Chen M, Low SH. Deepopf-v: solving ac-opf problems efficiently. *IEEE Trans Power Syst*. 2021;37:800-803.
41. Pan X, Chen M, Zhao T, Low SH. DeepOPF: a feasibility-optimized deep neural network approach for AC optimal power flow problems. *IEEE Syst J*. 2023;17:673-683.
42. Donti PL, Rolnick D, Kolter JZ. DC3: A learning method for optimization with hard constraints. arXiv preprint arXiv:2104.12225 2021.
43. Márquez-Neila P, Salzmann M, Fua P. Imposing hard constraints on deep networks: Promises and limitations. arXiv preprint arXiv:1706.02025 2017.
44. Luo T, Noll DC, Fessler JA, Nielsen J-F. Joint design of RF and gradient waveforms via auto-differentiation for 3D tailored excitation in MRI. *IEEE Trans Med Imaging*. 2021;40:3305-3314.
45. Maximov II, Vinding MS, Desmond HY, Nielsen NC, Shah NJ. Real-time 2D spatially selective MRI experiments: comparative analysis of optimal control design methods. *J Magn Reson*. 2015;254:110-120.
46. Guérin B, Setsompop K, Ye H, Poser BA, Stenger AV, Wald LL. Design of parallel transmission pulses for simultaneous multislice with explicit control for peak power and local specific absorption rate. *Magn Reson Med*. 2015;73:1946-1953.
47. Yaman B, Hosseini SAH, Akçakaya M. Zero-shot self-supervised learning for MRI reconstruction. *International Conference on Learning Representations*. 2022. <https://openreview.net/forum?id=085y6YPaYjP>
48. Schmitter S, Schnell S, Uğurbil K, Markl M, Van de Moortele PF. Towards high-resolution 4D flow MRI in the human aorta using kt-GRAPPA and B1+ shimming at 7T. *J Magn Reson Imaging*. 2016;44:486-499.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

Figure S1. The percent null generation results for threshold values on the flip angle maps are shown for both magnitude least-squares (MLS) and the proposed method. A threshold for both methods was chosen to be 0.6, indicated by the green dashed line in the figure. The dashed black line indicates that the MLS method produces nulls in the image even though the threshold is set to the desired

flip angle. This is because some flip angle maps generated by the MLS method overestimate the desired flip angle.

Figure S2. Correlation matrix plot for the quantitative measures, revealing the correlations between various variable pairs. Non-relevant correlations are not shown, and the upper triangular section represents the significance level of these correlations (r : Pearson correlation coefficient).

Figure S3. The figure illustrates the training and test loss, indicating that the model does not exhibit signs of overfitting.

Figure S4. Performance of our approach for 3D B_1^+ shimming. In this setting, the MLS method performs better than our approach in terms of CoV and RMSE, but our method has significantly better energy demand. Note that in this case, the training database size comprises 98 subjects, which is lower than the size indicated in our stress test (Figure 5). Thus, if 3D shimming is the target application, further performance gains may be achieved for the DL approach using a larger training database, which warrants further investigation. * and NS indicates that results are significant with $p < 0.05$

and not significant, respectively (The distributions were assessed for normality. All distributions were found to be non-Gaussian and the Wilcoxon signed-rank test was employed for all of them).

Figure S5. Performance of our approach for multi-slice B_1^+ shimming for two slices using a 3D CNN. In this setting, the proposed method outperforms MLS in terms of RMSE, CoV, and energy. * indicates that results are significant with $p < 0.05$ (The distributions were assessed for normality. CoV and energy distributions were found to be non-Gaussian and the Wilcoxon signed-rank test was employed for these distributions. RMSE distributions were found to be Gaussian and paired t-test was employed).

How to cite this article: Kilic T, Liebig P, Demirel OB, et al. Unsupervised deep learning with convolutional neural networks for static parallel transmit design: A retrospective study. *Magn Reson Med*. 2024;91:2498-2507. doi: 10.1002/mrm.30014