

Published in final edited form as:

Magn Reson Med. 2022 May 01; 87(5): 2254–2270. doi:10.1002/mrm.29132.

Rigid motion-resolved B_1^+ prediction using deep learning for real-time parallel-transmission pulse design

Alix Plumley¹, Luke Watkins^{1,2}, Matthias Treder³, Patrick Liebig⁴, Kevin Murphy², Emre Kopanoglu¹

¹School of Psychology, Cardiff University Brain Research Imaging Centre (CUBRIC), Cardiff University, Cardiff, United Kingdom

²School of Physics & Astronomy, Cardiff University Brain Research Imaging Centre (CUBRIC), Cardiff, United Kingdom

³School of Computer Science and Informatics, Cardiff University, Cardiff, United Kingdom

⁴Siemens Healthcare, GmbH, Erlangen, Germany

Abstract

Purpose—Tailored parallel-transmit (pTx) pulses produce uniform excitation profiles at 7 T, but are sensitive to head motion. A potential solution is real-time pulse redesign. A deep learning framework is proposed to estimate pTx B_1^+ distributions following within-slice motion, which can then be used for tailored pTx pulse redesign.

Methods—Using simulated data, conditional generative adversarial networks were trained to predict B_1^+ distributions in the head following a displacement. Predictions were made for two virtual body models that were not included in training. Predicted maps were compared with groundtruth (simulated, following motion) B_1^+ maps. Tailored pTx pulses were designed using B_1^+ maps at the original position (simulated, no motion) and evaluated using simulated B_1^+ maps at displaced position (ground-truth maps) to quantify motion-related excitation error. A second pulse was designed using predicted maps (also evaluated on ground-truth maps) to investigate improvement offered by the proposed method.

Results—Predicted B_1^+ maps corresponded well with ground-truth maps. Error in predicted maps was lower than motion-related error in 99% and 67% of magnitude and phase evaluations, respectively. Worst-case flip-angle normalized RMS error due to motion (76% of target flip angle) was reduced by 59% when pulses were redesigned using predicted maps.

Conclusion—We propose a framework for predicting B_1^+ maps online with deep neural networks. Predicted maps can then be used for real-time tailored pulse redesign, helping to overcome head motion-related error in pTx.

1 Introduction

Parallel transmission (pTx) of RF pulses through independently controlled channels can help to overcome B_1 nonuniformity seen in the head at 7 T,^{1, 2} particularly when tailored pulses are used.³ Tailored pulse design incorporates the measured transmit sensitivities (B_1^+) of each pTx channel, achieving a homogeneous flip angle across specified slices or regions.

For optimal tailored pulse performance, the measured B_1^+ distributions must match those present at the time of pulse payout. However, channels' electromagnetic fields (including B_1^+) and their interference patterns depend critically upon the object being imaged (i.e., the coil load), including its position, geometry, and composition.⁴⁻⁶

Geometrical and compositional differences between human subjects are partly addressed in alternative, nontailored approaches such as universal pulses (UPs),^{7, 8} SmartPulse,⁹ and fast online-customized pTx pulses.¹⁰ Intersubject robustness is achieved by designing a UP (offline) to minimize error across a small database of representative subjects. An underlying assumption is that the range in head geometry and composition across human subjects is relatively constrained, implying that B_1^+ distributions are similarly constrained. The designed pulse (a minimum error solution for excitation over multiple B_1^+ distributions) is therefore assumed to work fairly well for any individual subject without the need for B_1^+ mapping. Plug-and-play usability of UPs in pTx has led to the method's growing popularity.

However, the intersubject robustness of UPs comes at a cost to flip-angle uniformity. Tailored pulses typically yield lower normalized RMS error (nRMSE) of flip angle compared with UPs (7% vs 11% in Gras et al⁷). Additionally, the database approach is problematic in cases in which an individual is an outlier with respect to anatomies represented in the database. Moreover, these methods do not address the dependence of B_1^+ on load position, leading to unpredictable pulse performance in cases of different initial subject positioning¹¹ and/or within-scan head motion.¹²⁻¹⁴ The former is often overlooked, whereas the latter is commonly reported.¹⁵ Large head movements (exceeding 20 mm/degree) often occur among certain clinical populations,^{16, 17} elderly,¹⁸ and pediatric^{19, 20} subjects. Because flip angle (and therefore the acquired signal) depends on B_1^+ , displacements of approximately 5° have been found to cause an excitation error of 12%–19% (percent of target flip angle) when using pTx at 7 T,¹² with larger movements causing larger flip angle–related artifacts.

A few approaches have been proposed to correct motion-related RF field changes. Faraji-Dana et al. partially overcame motion-related effects on the (receive) B_1 field by simply reorienting coils' measured sensitivity maps using a Euclidean transformation.²¹ Similarly, Wallace et al. used radial basis functions to extrapolate channel sensitivities to voxel locations outside of the head, providing sensitivity information for all voxels in the FOV, regardless of head position.²² Extrapolated maps were used for retrospective correction. Neither approach considered dynamic motion-related field changes (e.g., changes in coil loading, shifting susceptibility gradients in tissue), as their effects were deemed minimal at 3 T. However, interactions between channels' highly nonuniform transmit fields at 7 T,²³ especially with pTx, indicate that dynamic motion-induced field changes cannot be overlooked. In contrast with these approaches, data-driven approaches inherently incorporate these changes.

Motion artifacts are often addressed through retrospective correction^{22, 24-27}; however, this is problematic for several reasons. First, the issues described previously cannot be corrected retrospectively without motion-resolved B_1^+ maps, which are not available. Second, channels' electric fields depend on the same factors, including load position.

Specific absorption rate (SAR) distribution and associated tissue heating are therefore also sensitive to motion, and are especially so in pTx due to constructive interference between channels' electric fields.²⁸⁻³⁰ Peak local SAR can exceed safety limits when head motion occurs in pTx simulations²⁹—a critical issue that cannot be addressed retrospectively. Conservatively bounded SAR estimates may be used, but this can prevent optimal imaging performance by limiting the RF power.^{2, 11, 31} In this study, the effect of motion on flip angle is the primary focus.

It is therefore desirable to overcome the motion dependence of tailored pTx pulse performance, and to do so using prospective techniques. Real-time pTx pulse design has been proposed as a solution, in which channels' complex coefficients are continuously updated to counteract motion-induced sensitivity changes. Multispoke pTx pulses can be designed in less than 0.5 seconds,³² whereas 2D spatially selective spiral pulses can be estimated in about 9 ms using deep neural networks.³³ With motion detection (e.g., Refs 22, 25, 34, and 35), channel updates could be determined by instantaneous head position, retaining flip-angle uniformity in cases of arbitrary and/or extreme motion. However, the required updates to channel coefficients depend on the motion-related field changes. Because real-time (i.e., motion-resolved) B_1^+ maps are not measurable, this requires that the relationship between head position and B_1^+ distribution to be characterized.

Deep convolutional neural networks have previously been used to estimate (non-pTx) B_1^+ distributions. In Wu *et al.*,³⁶ high-quality maps were predicted from reconstructed T_1 -weighted images, removing the need for B_1^+ mapping, while still allowing retrospective correction of B_1^+ related artifacts in quantitative MRI. This approach was limited to postprocessing; prediction quality deteriorated when B_1^+ was predicted directly from undersampled images. Abbasi-Rad et al. used a convolutional neural network to reconstruct B_1^+ from a localizer scan for the purpose of SAR reduction through pulse scaling based on slice-wise B_1^+ magnitude; however, B_1^+ prediction quality was dependent on head position.³⁷

In this work, we train a system of conditional generative adversarial networks³⁸ to predict pTx B_1^+ distributions (referred to as B_1 maps) following simulated head motion, given the initial B_1 maps at the centered position as input. If used in conjunction with motion detection, this would constitute motion-resolved B_1 -map estimation, and therefore permit real-time tailored pulse design. B_1 -map prediction quality is assessed by comparison with ground-truth (simulation output) B_1 maps following motion. Furthermore, flip-angle distributions of multispoke pTx pulses designed using network-predicted B_1 maps are compared with those produced by tailored pulses designed using the initial subject-specific B_1 maps alone. Finally, we also observe peak 10-g averaged local SAR for both pulses following motion.

2 Methods

2.1 Simulations and data

Dizzy, Billie, Duke, and Ella (Figure 1A-D) of the Virtual Population³⁹ (IT'IS, Zurich, Switzerland) were simulated with a generic 8-channel pTx coil in Sim4Life (ZMT, Zurich,

Switzerland). Each model was simulated at one central, and 32 off-center, positions. Off-center positions included rightward 2, 4, 5, 10, and 20 mm, posterior 2, 4, 5, and 10 mm, and all possible combinations thereof. These 29 positions are hereafter referred to as the R-P grid (Figure 1E). In addition, yaw 5°, 10°, and 15° positions were also simulated (Figure 1F). The Duke model was scaled to 90% of the original size, as the body and coil models intersected at some positions when the model was full-sized. To ensure consistent voxelization (and therefore consistent partial volume effects) in the body model across all simulated positions, the coil array was displaced rather than the body model. Simulations included the head, neck and shoulders,⁴⁰ and were run at 295 MHz following coil tuning to this frequency. Simulation results were normalized to an accepted power of 1 W per channel beyond the input port to the coil elements, to override imperfections in coil matching and any positional dependencies. The simulations were manually checked for input impedance and reflection coefficient as well as field smoothness across positions.

Channels' 3D B_1 , electric field, current density, and SAR distributions were masked to exclude background (air) voxels and exported to MATLAB (The MathWorks, Natick, MA). To incorporate interactions between channels for local SAR evaluations, 10g-averaged Q-matrices were calculated.^{28, 41, 42} Elements of the 8×8 Q-matrices were

$$Q_{ij}(r) = \frac{1}{2\rho(r)} \left[J_{x,j}^H(r) E_{x,i}(r) + J_{y,j}^H(r) E_{y,i}(r) + J_{z,j}^H(r) E_{z,i}(r) \right], \quad (1)$$

where $\rho(r)$ is the tissue mass density (kg/m^3) in voxel r ; j is the complex current density (A/m^2); E is the complex electric field (V/m); x, y and z are the three Cartesian axes; i and j are transmit channel indices; and H denotes Hermitian transpose.

B_1 maps from 51 slices spanning a mid-axial slab with a thickness of 9 cm from the Duke and Ella body models (Figure 1C,D) were prepared for network training by interpolating to 256×256 in-plane resolution. The same preprocessing was applied to the Billie and Dizzy data, but at only six slice locations (Figure 1A,B). Magnitude and phase data were separated and normalized between 0 and 1, where 1 corresponds to the maximum magnitude across all channels, slices, and body models, and to 2π for phase. Random offsets were applied to phase maps so that the phase wrap boundary did not occur at the same location across slices. B_1 maps were input to networks as individual axial slices with size $256 \times 256 \times 8$, where the third dimension is channels. Corresponding B_1 -map slices before (input) and after (ground truth) a given displacement formed the networks' input-target pairs. Note that inputs are not necessarily at the centered position (explained later in Section 2.2).

2.2 Neural networks and network training

Models were implemented in TensorFlow 2.3⁴³ using Python 3.7. Network architecture is summarized in Figure 2. Except where specified, network hyperparameters were the same as those used in the Pix2Pix conditional generative adversarial network.⁴⁴ The generators were U-Net⁴⁵ models with eight convolutional (encoding) and eight deconvolutional (decoding) layers linking the input and output (predicted) B_1 maps, each followed by rectified linear unit activation layers. Filters were 4×4 for magnitude and 8×8 for phase. Although comprehensive hyperparameter optimization was beyond the scope of this project, during

initial testing it was found that phase networks benefited from the large receptive field of 8×8 filters. Conversely, magnitude networks generated smoother maps when more filters were used. To avoid increasing the number of trainable parameters, filters were smaller for magnitude. The number of filters (initially 128 for magnitude, and 64 for phase) increased to a maximum of 1024 (512 for phase) for the middle layers, and stride size was 2. Filters were split into eight groups to facilitate simultaneous processing of all pTx channels. For phase, batch normalization was applied at all layers except the first convolution layer. For magnitude, removing batch normalization resulted in a smoother training curve and higher-quality estimated maps. Skip connections joined each convolution layer to the symmetric deconvolution layer for network stability. The network was regularized through dropout layers following each of the first three deconvolution layers (rate = 0.5).

In contrast to encoder–decoder models that typically rely on minimizing L1 loss between predicted and target images, generative adversarial networks include an additional loss term, which helps to reduce blurring often seen with L1 loss alone.⁴⁴ This is provided through a second convolutional neural network—the discriminator—which is trained to distinguish between generator-predicted and ground-truth distributions. The input B1 maps, concatenated with either ground-truth or generator-predicted B1 maps, served as input to the discriminators, which consisted of five convolution layers. The discriminators used leaky rectified linear unit activation layers ($a = 0.3$) as recommended in Radford et al.⁴⁶ Filter size was the same as that for the generators, and convolution stride was 2 except for the final two layers, where it was 1. A single 2D distribution of probability (entropy) values was output.

The overall conditional generative adversarial network loss function can be expressed as

$$Loss = \operatorname{argmin}_G \max_D \mathcal{L}_{cGAN}(G, D) + \lambda \mathcal{L}_{L1}(G) \quad (2)$$

where G denotes the generator; D is the discriminator; and λ (set to 100) is a scaling parameter acting on the L1-norm between generator-predicted and ground-truth maps. The first term can be further described as

$$\begin{aligned} \mathcal{L}_{cGAN}(G, D) = & E_{B1_{gt}, B1_{predicted}} [\log D(B1_{gt}, B1_{predicted})] \\ & + E_{B1_{gt}, B1_{initial}} [\log(1 - D(B1_{gt}, G(B1_{gt}, B1_{initial})))], \end{aligned} \quad (3)$$

where $B1_{gt}$ are the ground-truth displaced B1 maps; $B1_{predicted}$ are the generator-predicted displaced B1-maps; and $B1_{initial}$ are the pre-displacement B1 maps (network input).

The effect of head motion on B1 depends on the displacement type (i.e., direction, magnitude).^{12, 13} Because data-driven approaches assume that all input-target pairs share a common underlying mapping, separate networks were trained for different displacement types (e.g., rightward vs posterior). Head motion was discretized into large (5 mm) and small (2 mm) displacements in rightward (R) and posterior (P) directions to cover the R-P grid. Additional networks were trained for 5° yaw rotation. Separate networks were trained for magnitude and phase, yielding a total of 10 networks.

The Duke and Ella data were used for training. All available relative displacements were included. For example, to train the R5 mm network, such as (input)–(target) pairs included (R0, P0 mm)–(R5, P0 mm); (R5, P0 mm)–(R10, P0 mm); (R5, P2 mm)–(R10, P2 mm). This yielded a training data set of 1020 unique slices for rightward and leftward networks, and 1224 for each posterior network. The yaw network training database was smaller (306 slices).

The Adam⁴⁷ optimizer was used to train models for 60 epochs. Learning rate was critical during initial testing, so learning rates within the range 5e-5 to 1e-3 were tested. The default value of 2e-4 converged most effectively and was therefore used for all networks. Network weights were saved at the epoch, which yielded the lowest total error across the validation data set (the Billie data) as a form of early stopping to help prevent network overfitting. Networks took approximately 16 hours to train with a batch size of 1 using a standard PC with NVIDIA GeForce GTX 1050 Ti.

2.3 Network evaluation and cascading

Networks were tested using the Billie and Dizzy data at six slice locations (Figure 1A). For Billie, different slices were used compared with those used for early stopping (Dizzy was not involved in the training process at all). Like training, testing was conducted for all available examples of each displacement, yielding test data sets of between 6 and 72 slices. In addition to the positions listed in section 2.1, Billie and Dizzy models were simulated at three combined yaw-rightward positions to test performance for motion involving both rotation and translation. Because networks were only trained for five displacements but evaluated at 35 positions, networks were cascaded where necessary. Starting with the center position's B_1 maps as input, generators were run sequentially, with the output of one generator used as input to the next, until the desired evaluation position was reached. For example, R5 mm, R5 mm, and P2 mm networks were cascaded for evaluation at the (R10, P2 mm) position.

Finally, the Billie model was also simulated at inferior 5, 10, and 15 mm to investigate error for through-plane motion.

Predicted B_1 maps were exported to MATLAB. Voxels with < 1% of the maximum magnitude were smoothed with a Gaussian kernel. Corresponding magnitude and phase network outputs were subsequently combined to form complex predicted maps ($B_{1\text{predicted}}$).

The $B_{1\text{predicted}}$ quality was assessed through voxel-wise correlation (using MATLAB's `corrcoef` function) and prediction error between predicted and ground-truth maps at each position. These values were compared with error and correlation following head motion (i.e., between the two simulated maps). Calculations were performed on the tissue-masked region, with the outermost two voxels excluded to avoid artificial amplification of error due to partially filled voxels. Prediction error for magnitude and phase distributions were assessed through nRMSE and L1 norm, respectively, as follows:

$$nRMSE_{\| prediction \|} = 100 \times \frac{\sqrt{\frac{1}{N_v} \sum_{r=1}^{N_v} \| B1_{gt}(r) - B1_{predicted}(r) \|^2}}{\frac{1}{N_v} \sum_{r=1}^{N_v} \| B1_{gt}(r) \|} \quad (4)$$

$$L1_{\angle prediction} = \frac{1}{N_v} \sum_{r=1}^{N_v} \| \angle e^{j(B1_{gt}(r) - B1_{predicted}(r))} \| \quad (5)$$

where j is $\sqrt{-1}$; and N_v is the number of voxels in a slice, indexed by r . Motion-induced error was calculated analogously, but substituting $B1_{initial}$ for $B1_{predicted}$ in Equations 4 and 5.

2.4 Pulse design and analysis

Outputs from the R-P grid positions were further processed to assess whether predicted maps were of sufficient quality to be used for tailored pTx pulse design. Five-spoke excitation pulses were designed using a small tip-angle spatial domain method,^{3, 48, 49} and two approaches were compared in terms of their performance following motion within the R-P grid. A schematic of the process is shown in Figure 3. First, a conventional tailored pulse (pulse_{initial}) was designed using the subject-specific B_1 maps at the initial position ($B1_{initial}$). A uniform magnitude target excitation profile (target flip angle = 70°) was specified for pulse_{initial}. Pulse coefficients were optimized iteratively to minimize magnitude error, whereas the profile's phase was relaxed.⁵⁰ The resultant complex profile was used as the target profile for a second pulse (pulse_{re-designed}), which was designed without phase relaxation (because magnitude and phase distributions need to be consistent across positions to ensure data consistency for motion occurring mid-acquisition). Pulse_{e-designed} was designed using the network-output $B1_{predicted}$ (the proposed approach).

Both pulses (pulse_{initial} and pulse_{re-designed}) were subsequently evaluated using the ground-truth B_1 maps at the displaced position ($B1_{gt}$) to quantify motion-induced effects on the conventional approach, and improvement provided by the proposed method. Their flip-angle distributions were compared with that of pulse_{initial} without motion in terms of nRMSE, expressed as percent target flip angle as follows:

$$nRMSE_{\theta} = 100 \times \frac{\sqrt{\frac{1}{N_v} \sum_{r=1}^{N_v} \| \theta_{displaced}(r) - \theta_{initial}(r) \|^2}}{\theta_t} \quad (6)$$

where $\theta_{initial}$ is flip angle without motion; $\theta_{displaced}$ is that following motion (with either pulse_{initial} or pulse_{re-designed}); and θ_t is the target flip angle. The nRMSE for pulse_{initial} without motion (i.e., the “gold standard”) was also calculated by substituting $\theta_{displaced}$ for θ_t in Equation 6.

Peak local SAR (psSAR) of both pulses was also evaluated using the 10-g averaged Q-matrices at each position. Because psSAR sensitivity to motion has been reported to exhibit

no slice dependence,²⁹ SAR was evaluated at four target imaging slices (out of the six used for pulse design) (Figure 1A).

3 Results

3.1 B_1 prediction quality

B_1 maps were predicted by networks in about 14 ms using the same PC as used for training. Example B_1 maps, motion-induced error, and prediction error are shown for a 5 mm displacement in Figure 4. Motion-induced error (averaged across channels) for this example was 15.1% (magnitude) and 4.9° (phase), whereas mean prediction error was 3.2% (magnitude) and 3.5° (phase).

Figure 5 shows a summary of error and correlation coefficient for magnitude and phase at each evaluated displacement (averaged across Dizzy and Billie models, slices, channels, and initial positions). Position dependence of prediction quality was minimal compared with motion-related error, as seen by the reduced gradient with respect to displacement norm in all cases. Dizzy and Billie models yielded very similar prediction quality (Supporting Information Figure S1).

Mean motion-induced magnitude error scaled linearly with displacement magnitude at about 3% per millimeter (or 3.2% per degree of rotation), compared with 0.36% per millimeter (0.27% per degree) for prediction error. Prediction error was lower than motion-related error in 99.8% of translation, and 90% of rotation evaluations. Figure 6A shows B_1 magnitude nRMSE for magnitude for all slices and channels for 10 example displacements. Motion caused a worst-case magnitude error of 117% following a displacement of R20, P10 mm, whereas maximum prediction error was 33% (at the $y15^\circ$, R4 mm position). Worst-case prediction error from the R-P grid was lower (20% at the R20, P10 mm position).

Example magnitude correlations are shown in Figure 7A. The lowest observed correlation coefficient between $B_{1\text{initial}}$ and $B_{1\text{gt}}$ magnitudes was 0.79 following a $y15^\circ$, R4 mm displacement. Correlation between $B_{1\text{predicted}}$ and $B_{1\text{gt}}$ did not fall below 0.96.

Motion-induced error and correlation were observed to be slice-dependent and channel-dependent (i.e., the error depended on the displacement relative to each channel's location, as expected). The $B_{1\text{predicted}}$ quality was similar across channels, as demonstrated by the strong correlation across all channels in Figure 7A. However, prediction error was somewhat slice-dependent, with the most inferior slice locations yielding the highest prediction errors (slice information not shown).

Phase error for 10 example displacements are shown in Figure 6B. For phase, maximum observed prediction error (57°) was similar to maximum motion-related error (55°). These worst cases arose in the Dizzy model; for the Billie model, maximum prediction error (29.4°) was lower than that caused by motion (44.2°). Furthermore, prediction error was lower than motion-induced error for phase in 68% and 66% of translation and rotation evaluations, respectively (including both models).

Yaw rotation caused substantially higher error than axial translations; for the R-P grid, maximum prediction-related and motion-related errors were 19.8° and 34.7° , respectively. Mean phase prediction error was less position-dependent than motion-related error, with axial translations causing error of approximately 0.9° per millimeter displacement, compared with 0.4° per millimeter in predicted maps. For displacements including rotation, analogous gradients were 2° and 0.2° per degree of yaw, respectively.

Mean phase correlation coefficient between predicted and ground-truth maps was higher than (or very similar to) that between initial and ground-truth maps for all displacements. Phase correlation examples are shown in Figure 7B. Correlation coefficient between $B1_{\text{predicted}}$ and $B1_{\text{gt}}$ exceeded that between $B1_{\text{predicted}}$ and $B1_{\text{initial}}$ in 69% of cases.

3.2 Parallel-transmit pulse performance

Subsequent analyses were conducted using the Billie model with the R-P grid data only. Five-spokes pTx pulses designed using $B1_{\text{initial}}$ ($\text{pulse}_{\text{initial}}$) yielded uniform flip-angle profiles (mean nRMSE $\sim 1\%$) without motion. However as expected, uniformity was lost following axial translation. Pulses were about 7.7 ms long.

Figure 8 shows that flip-angle nRMSE for $\text{pulse}_{\text{initial}}$ was strongly position-dependent, reaching a maximum of 76% following a displacement of R20, P5 mm. Conversely, pulses redesigned using $B1_{\text{predicted}}$ (pulseredesigned) produced much improved flip-angle profiles when evaluated at the displaced position, yielding nRMSE of 14% for the same displacement. Maximum pulseredesigned nRMSE was 17% (at the R2, P10 mm position), whereas this error value was exceeded by $\text{pulse}_{\text{initial}}$ (i.e., without any correction) after displacements of just 5 mm. The largest errors occurred in inferior slice locations for both pulses (slice information not shown). Maximum motion-related error in the excitation profile's phase (110.4°) was reduced by 7.8° when using pulseredesigned .

Figure 8B shows flip-angle profiles for both pulses following several example displacements. Supporting Information Figure S2A also shows flip-angle nRMSE for nine example displacements. It should be noted that flip-angle uniformity for pulseredesigned could be further improved by including phase relaxation in the design (as was done for $\text{pulse}_{\text{initial}}$); however, this would permit excitation phase to vary throughout the scan, causing reconstruction inconsistencies.

3.3 Cascading

The $B1_{\text{predicted}}$ quality remained high when networks were cascaded multiple times; however, there was a weak linear relationship between prediction error and motion magnitude. To investigate the impact of cascading on prediction quality, we ran secondary analyses for displacements of R0, P10 mm, R-2, P10 mm, and R-5, P10 mm using only the P2 mm network for the posterior component. Running the 2 mm network five times (i.e., four cascades) led to approximate average increases in magnitude and phase error of 1.2% and 1.2° , respectively, compared with running the 5 mm network twice (one cascade). There was also reduced flip-angle uniformity compared with using the 5 mm network. Nevertheless, Figure 9 shows that motion-induced error was appreciably reduced using either approach.

3.4 Specific absorption rate

In addition to flip angle, SAR was also evaluated for the R-P grid positions. Following motion, psSAR produced by pulse_{redesigned} was lower than that of pulse_{initial} in 89% of cases. For pulse_{initial}, motion caused psSAR to increase (relative to that without motion; psSAR_{center}) in 72% of evaluations. When pulse_{redesigned} was used instead, psSAR increased relative to the centered case in only 16% of the cases. The psSAR for nine example displacements is shown in Supporting Information Figure S2B.

Figure 10A shows worst-case psSAR for each evaluated position (relative to psSAR_{center}). Figure 10B shows the same information, separated by slice location. Motion-related SAR change was similar across slices, whereas benefits offered by pulse_{redesigned} were most consistently seen in slices 1 and 6 (furthest from mid-axial locations). Pulse_{redesigned} yielded lower psSAR than pulse_{initial} following large displacements, but sometimes resulted in higher psSAR for small displacements, especially in slice 4 (mid-axial). In the worst observed case, pulse_{redesigned} yielded a 3.1-fold increase in psSAR (slice 4 at the posterior 5-mm position), whereas pulse_{initial} increased by a maximum of 3.3-fold (following the largest axial displacement). It should be noted that SAR was not used as a constraint in either pulse's design.

4 Discussion

As far as we know, this study is the first to demonstrate motion-resolved B₁ map prediction in pTx. We successfully trained a system of deep neural networks to estimate B₁⁺ sensitivity distributions following simulated in-plane head motion. Predicted B₁ maps were of sufficiently high quality to be used for tailored pTx pulse design, and resulting pulses outperformed conventionally designed tailored pulses under conditions of head motion.

Across the four R-P grid magnitude networks, prediction error was $4.5\% \pm 1.5\%$ (mean \pm SD) of the ground-truth B₁ magnitude (tested individually without cascading and error calculated according to the definition in Abbasi-Rad et al³⁷). This is in line with expectations based on previous work, in which magnitude B₁ maps were predicted with mean error of $9.5 \pm 7.1\%$.³⁷ The experimental and 3D nature of data in Abbasi-Rad et al³⁷ may have caused the slightly higher error. We also observe similar B₁_{predicted} correlation coefficients to those reported in Wu et al³⁶ (~0.99), which was retained even when our networks were cascaded multiple times. Yaw rotation caused higher error than translational motion, and prediction error was generally higher than that of translations. We attribute this to the much smaller training database used to train yaw networks.

It was previously reported that excitation profile nRMSE increases by approximately 2.4% per millimeter of axial displacement in simulations,¹² and our results (2.7% per millimeter) are in good agreement. A 12%–22% flip-angle error was observed in vivo following motion in the range of about 5–16 mm in Kopanoglu et al.,¹² and we observed an error of about 11%–35% for a similar range of motion. Error was reduced to approximately 8%–10% using the proposed approach in our study. For larger movements, the benefit of pulse redesign using predicted maps was even greater.

There was some slice dependence for both $B_{1\text{predicted}}$ quality and flip-angle error, with the highest errors observed for inferior slice locations. This is in line with previous research, in which higher motion sensitivity was observed for inferior slices passing through the temporal lobes and the cerebellum.¹² Inferior slices yield lower field magnitude, and $B_{1\text{prediction}}$ quality was lower in regions with very low field magnitude, which may explain the residual slice dependence in predicted maps. An alternative explanation is that there were fewer inferior slices in the training data set compared with mid-axial slices, which could have resulted in better training for mid-axial slices.

Considering that the R5 mm training data included positions up to just 10 mm along the rightward axis, it is noteworthy that the R5 mm network was able to extrapolate beyond this by successfully cascading four times to estimate the fields at the R20 mm position. We attribute this to the global normalization, conducted over all data sets. Results for large displacements could likely be further improved by including extreme positions in the training data set.

Magnitude networks consistently outperformed those of phase. Prediction error exceeded motion-induced error for phase in about a third of evaluations. Improvements to excitation phase were evident but modest. It has been previously acknowledged that phase changes due to motion are difficult to model, and other work on motion-related field changes,²² including B_1 prediction,³⁷ often neglect phase altogether. Most of the phase-prediction error occurred at phase-wrapping boundaries. This was somewhat reduced by applying random phase offsets to pairs of input B_1 maps, but it was not eliminated. In terms of pulse performance, these small, local errors are likely to have less impact than the global changes caused by motion, which could potentially accumulate when channels are superposed. Incidentally, when we instead trained networks using unwrapped phase data, the error in $B_{1\text{predicted}}$ was globally higher and yielded inferior results compared with wrapped data. In contrast, the local prediction errors seen with wrapped data are not structured. We believe that the increase in error seen for unwrapped data was due to the increased dynamic range of unwrapped data, meaning that relative changes due to motion were smaller following normalization.

One limitation of this study is that it deals with simulated data only because of the lack of models for the RF coils at the institution due to proprietary information. In DeepQSM,⁵¹ training data were solely synthetic; the ground truth consisted of overlapping cubes and spheres with known (simulated) susceptibility values. This was convolved with a forward dipole kernel to create the corresponding input. Networks were able to resolve high-quality susceptibility maps for human brains, despite only being trained on simple geometric shapes. Similarly, Meladio et al. demonstrated successful in vivo validation following training with synthetic data.⁵² We believe that the method proposed here would be similarly generalizable if a realistic RF coil model (i.e., a model of the coil to be used) is used for simulations. Moreover, using simulated training data avoids the requirement for choreographed in vivo head movement to be replicated precisely across several subjects to create the training data sets, which would be practically infeasible. Aside from the initial (measured) input B_1 map, networks automatically output maps in patient coordinates, making online registration unnecessary. However, motion tracking is needed to determine

which network(s) are required, and for online corrections to gradient waveforms to update the imaging volume as in Zaitsev et al.⁵³

The minimum motion resolution we consider here is 2 mm. Smaller movements could remain problematic for quantitative MR protocols that rely on signal changes on the scale of 1%–4% of the total signal.²¹ Finer discretization is possible through simulations of smaller displacements. The most appropriate motion discretization will depend on the user's primary aims. While using finer discretization and cascading more (i.e., 5×2 mm vs 2×5 mm) did result in slightly higher error here, motion-induced error was still largely ameliorated.

In this study we considered only positive displacements (e.g., rightward, but not leftward). Because the trained networks are direction-specific, the same networks cannot handle negative displacements. This is due to the nature of inverting a deep neural network (not possible, as reversing convolutional layers results in a highly underdetermined problem). However, the training of networks for positive displacements does not limit generalizability, and additional networks can similarly be trained for negative displacements. In fact, by training a network for positive and negative displacements for each of the 6 degrees of freedom of motion, all rigid-body motion could be covered.

We focused on in-plane motion (axial translations and yaw rotation), as motion-induced error was observed to be more spatially varying. Although some motion-induced B_1 error was observed in peripheral slices following inferior translation (see Supporting Information Figure S3), the error was relatively lower, and importantly, more global (i.e., spatially smoothly varying) within the slice. We attribute this to the fact that relative tissue-channel distances remain constant for through-slice translation. This means that simpler correction methods (e.g., slice-dependent pulse scaling) could feasibly be used to counteract the B_1 effects of through-plane motion.

In contrast to B_1 , through-plane motion was shown to be more disruptive to B_0 than within-plane motion.^{54, 55} Although B_0 off-resonance can be incorporated in tailored pulse design,⁵⁶ this cannot currently be updated in real time with pTx, as it increases the degrees of freedom to be optimized in the pulse, pushing redesign times beyond practically feasible TR values.³² Instead, motion-related effects on B_0 can be corrected retrospectively, such as using data-driven coefficients to link motion with field changes.⁵⁵ Alternatively, real-time B_0 shimming may be possible with multicoil shim arrays by predicting B_0 field changes due to motion in a manner similar to the method proposed here.

The SAR observations reported here are incidental. The focal point of this study was to develop a method to accurately estimate B_1 maps following motion. Using SAR as a pulse design constraint would trade flip-angle homogeneity for reduced SAR, thereby overshadowing B_1 quality. Hence, SAR was not used as a design constraint. Although motion-related SAR increase was generally lower for pulses redesigned with predicted B_1 maps, it was higher for a minority of cases. Motion sensitivity of SAR in pTx has previously been reported to be similar across axial slices,²⁹ and we also did not observe clear slice dependence for SAR motion sensitivity. However, we did observe that inferior-most and superior-most slices benefited from the proposed approach more consistently

than mid-axial slices. Nevertheless, the overall improvement offered by pulseredesigned, especially for larger displacements, is promising for future development of this approach. Neural networks have previously been used to predict B_1 maps for the purpose of SAR reduction.³⁷ This was achieved through slice-wise pulse scaling based on a predicted 3D B_1 magnitude. The entire 9-cm axial slab could be predicted within approximately 0.8 seconds, permitting concatenation into pseudo-3D B_1 maps that could feasibly be used for pulse scaling or similar SAR management here. Pulse scaling based on B_1^+ cannot guarantee SAR compliance, as B_1^+ does not necessarily reflect electric-field distributions. However, SAR compliance could be ensured if (3D) electric fields were also predicted.

5 Conclusions

We have demonstrated a framework for a deep-learning approach for motion-resolved B_1^+ estimation in pTx. Estimated maps can be used for real-time tailored pulse redesign, yielding homogeneous flip-angle profiles in cases of head motion. Importantly, networks can be run sequentially to predict B_1 maps following arbitrary displacements comprising multiple directions. Here, error was reduced for 35 displacements using networks trained for just five displacements. Our findings represent one potential avenue toward user-friendly, optimized pTx at 7 T.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

References

1. Katscher U, Börner P. Parallel RF transmission in MRI. *NMR Biomed.* 2006; 19: 393–400. [PubMed: 16705630]
2. Deniz CM. Parallel transmission for ultrahigh field MRI. *Top Magn Reson Imaging.* 2019; 28: 159–171. [PubMed: 31188274]
3. 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. [PubMed: 16894579]
4. Katscher U, Röhrs J, Börner P. Basic considerations on the im-pact of the coil array on the performance of Transmit SENSE. *Magn Reson Mater Phys, Biol Med.* 2005; 18: 81–88.
5. Deniz CM, Vaidya MV, Sodickson DK, Lattanzi R. Radiofrequency energy deposition and radiofrequency power requirements in parallel transmission with increasing distance from the coil to the sample. *Magn Reson Med.* 2016; 75: 423–432. [PubMed: 25752250]
6. De Greef M, Ipek O, Raaijmakers AJE, Crezee J, Van Den Berg CAT. Specific absorption rate intersubject variability in 7T parallel transmit MRI of the head. *Magn Reson Med.* 2013; 69: 1476–1485. [PubMed: 22760930]
7. Gras V, Vignaud A, Amadon A, Le Bihan D, Boulant N. Universal pulses: a new concept for calibration- free parallel transmission. *Magn Reson Med.* 2017; 77: 635–643. [PubMed: 26888654]
8. Gras V, Mauconduit F, Vignaud A, et al. Design of universal parallel- transmit refocusing kT-point pulses and application to 3D T2-weighted imaging at 7T. *Magn Reson Med.* 2018; 80: 53–65. [PubMed: 29193250]
9. 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. [PubMed: 31257612]

10. Herrler J, Liebig P, Gumbrecht R, et al. Fast online-customized (FOCUS) parallel transmission pulses: a combination of uni-versal pulses and individual optimization. *Magn Reson Med*. 2021; 85: 3140–3153. [PubMed: 33400302]
11. Kopanoglu, E. Patient specific parallel transmit pulses are patient position dependent while safety models are fixed: safety implications; Proceedings of the 29th Annual Meeting of ISMRM [Virtual]; 2021. Abstract #2299
12. Kopanoglu, E; Plumley, A; Erturk, MA; Deniz, C; Wise, R. Implications of within-scan patient head motion on B1+ homogeneity and specific absorption rate at 7T; Proceedings of the 27th Annual Meeting of ISMRM; Montréal, Canada. 2019. Abstract #4686
13. Bammer, R; Zhang, B; Deng, W; Wiggins, G; Stenger, AV; Sodickson, D. Impact of motion on parallel transmission; Proceedings of the 19th Annual Meeting of ISMRM; Stockholm, Sweden. 2010. Abstract #4590
14. Schön, N; Petzold, J; Seifert, F; , et al. Impact of respiration on B1+ field and SAR distribution at 7T using a novel EM simulation setup; Proceedings of the 28th Annual Meeting of ISMRM [Virtual]; 2020. Abstract #1120
15. Zaitsev M, Maclaren JH, Herbst M. Motion artifacts in MRI: a complex problem with many partial solutions. *J Magn Reson Imaging*. 2015; 42: 887–901. [PubMed: 25630632]
16. Greene DJ, Church JA, Dosenbach NU, et al. Multivariate pattern classification of pediatric Tourette syndrome using functional connectivity MRI. *Developmental Science*. 2016; 19: 581–598. [PubMed: 26834084]
17. Schwarz ST, Afzal M, Morgan PS, Bajaj N, Gowland PA, Auer DP. The ‘swallow tail’ appearance of the healthy nigrosome— a new accurate test of Parkinson’s disease: a case-control and retrospective cross-sectional MRI study at 3T. *PLoS One*. 2014; 9: 1–8.
18. Saccà V, Sarica A, Quattrone A, Rocca F, Quattrone AN, Novellino F. Aging effect on head motion: a machine learning study on resting state fMRI data. *J Neurosci Methods*. 2021; 352: 109084 [PubMed: 33508406]
19. Chen KT, Salcedo S, Chonde DB, et al. MR- assisted PET motion correction in simultaneous PET/MRI studies of dementia sub-jects. *J Magn Reson Imaging*. 2018; 48: 1288–1296. [PubMed: 29517819]
20. Kecskemeti S, Samsonov A, Velikina J, et al. Robust motion correction strategy for structural MRI in un-sedated children demonstrated with three-dimensional radial MPnRAGE. *Radiology*. 2018; 289: 509–516. [PubMed: 30063192]
21. Faraji-Dana Z, Tam F, Chen JJ, Graham SJ. A robust method for suppressing motion- induced coil sensitivity variations during prospective correction of head motion in fMRI. *Magn Reson Imaging*. 2016; 34: 1206–1219. [PubMed: 27451407]
22. Wallace TE, Afacan O, Waszak M, Kober T, Warfield SK. Head motion measurement and correction using FID navigators. *Magn Reson Med*. 2019; 81: 258–274. [PubMed: 30058216]
23. Vaughan J, Garwood M, Collins C, et al. 7T vs. 4T: RF power, homogeneity, and signal- to-noise comparison in head images. *Magn Reson Med*. 2001; 46: 24–30. [PubMed: 11443707]
24. Bazin P-L, Nijisse HE, Van Der Zwaag W, et al. Sharpness in motion corrected quantitative imaging at 7T. *NeuroImage*. 2020; 222: 117227 [PubMed: 32781231]
25. Afacan O, Wallace TE, Warfield SK. Retrospective correction of head motion using measurements from an electromagnetic tracker. *Magn Reson Med*. 2020; 83: 427–437. [PubMed: 31400036]
26. An H, Shin H-G, Ji S, et al. DeepResp: deep learning solution for respiration- induced B0 fluctuation artifacts in multi-slice GRE. *NeuroImage*. 2021; 224: 117432 [PubMed: 33038539]
27. Gallichan D, Marques JP, Gruetter R. Retrospective correction of involuntary microscopic head movement using highly accelerated fat image navigators (3D FatNavs) at 7T. *Magn Reson Med*. 2016; 75: 1030–1039. [PubMed: 25872755]
28. Deniz CM, Alon L, Brown R, Sodickson DK, Zhu Y. Specific absorption rate benefits of including measured electric field interactions in parallel excitation pulse design. *Magn Reson Med*. 2012; 67: 164–174. [PubMed: 22135040]
29. Kopanoglu E, Deniz CM, Erturk MA, Wise RG. Specific absorption rate implications of within-scan patient head motion for ultra- high field MRI. *Magn Reson Med*. 2020; 84: 2724–2738. [PubMed: 32301177]

30. Ajanovic, A; Hajnal, J; Malik, S. Positional sensitivity of specific absorption rate in head at 7T; Proceedings of the 28th Annual Meeting of ISMRM [Virtual]; 2020. Abstract #4251
31. Melià EF, Sbrizzi A, Van Den Berg CAT, Steensma BR, Luijten PR, Raaijmakers AJE. Conditional safety margins for less conservative peak local SAR assessment: a probabilistic approach. *Magn Reson Med.* 2020; 84: 3379–3395. [PubMed: 32492249]
32. Kopanoglu, E. Near real- time parallel- transmit pulse design; Proceedings of the 26th Annual Meeting of ISMRM; Paris, France. 2018. Abstract #3392
33. Vinding MS, Aigner CS, Schmitter S, Lund TE. DeepControl:2DRF pulses facilitating inhomogeneity and B0 off- resonance compensation in vivo at 7T. *Magn Reson Med.* 2021; 85: 3308–3317. [PubMed: 33480029]
34. DiGiacomo P, Maclaren J, Aksoy M, et al. A within-coil optical prospective motion-correction system for brain imaging at 7T. *Magn Reson Med.* 2020; 84: 1661–1671. [PubMed: 32077521]
35. Bortolotti, L; Mougin, OB. R. Measurement of head motion using a field camera in a 7T scanner; Proceedings of the 28th Annual Meeting of ISMRM [Virtual]; 2020. Abstract #464
36. Wu Y, Ma Y, Du J, Xing L. Accelerating quantitative MR imaging with the incorporation of B1 compensation using deep learning. *Magn Reson Imaging.* 2020; 72: 78–86. [PubMed: 32610065]
37. Abbasi-Rad S, O'Brien K, Kelly S, et al. Improving FLAIR SAR efficiency at 7T by adaptive tailoring of adiabatic pulse power through deep learning estimation. *Magn Reson Med.* 2021; 85: 2462–2476. [PubMed: 33226685]
38. Goodfellow, IJ; Pouget-Abadie, J; Mirza, M; , et al. Generative adversarial nets; NIPS' 14: Proceedings of the 27th International Conference on Neural Information Processing Systems; 2014. 2672–2680.
39. Gosselin M-C, Neufeld E, Moser H, et al. Development of a new generation of high- resolution anatomical models for medical device evaluation: the Virtual Population 3.0. *Phys Med Biol.* 2014; 59: 5287–5303. [PubMed: 25144615]
40. Wolf S, Diehl D, Gebhardt M, Mallow J, Speck O. SAR simulations for high- field MRI: how much detail, effort, and accuracy is needed? *Magn Reson Med.* 2013; 69: 1157–1168. [PubMed: 22611018]
41. Graesslin I, Homann H, Biederer S, et al. A specific absorption rate prediction concept for parallel transmission MR. *Magn Reson Med.* 2012; 68: 1664–1674. [PubMed: 22231647]
42. Standard III. Determining the peak spatial- average specific absorption rate (SAR) in the human body from wireless com-munications devices, 30 MHz to 6 GHz— Part 1: General requirements for using the finite-difference time-domain (FDTD) method for SAR calculations. IEC/IEEE. 2017; 1: 1–86. 62704-
43. Abadi M, Agarwal A, Barham P. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. 2015.
44. Isola, P; Zhu, J-Y; Zhou, T; Efros, AA. Image-to- image translation with conditional adversarial networks. arXiv; IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2017. 5967–5976.
45. Ronneberger O, Fischer P, Brox T. U-net: convolutional networks for biomedical image segmentation. *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015* arXiv. 2015; 9351: 234–241. 1505.04597v1
46. Radford A, Metz L, Chintala S. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. arXiv. 2016. 1511.06434v06432
47. Kingma, D; Lei, BJ. Adam: a method for stochastic optimization. arXiv; Int Conf Learning Representations; 2015. 1412.6980v9
48. Grissom WA, Setsompop K, Hurley SA, Tsao J, Velikina JV, Samsonov AA. Advancing RF pulse design using an open- competition format: report from the 2015 ISMRM challenge. *Magn Reson Med.* 2017; 78: 1352–1361. [PubMed: 27790754]
49. Kopanoglu E, Constable RT. Radiofrequency pulse design using nonlinear gradient magnetic fields. *Magn Reson Med.* 2015; 74: 826–839. [PubMed: 25203286]
50. Setsompop K, Wald L, Alagappan V, Gagoski B, 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. [PubMed: 18383281]

51. Bollmann S, Rasmussen KGB, Kristensen M, et al. DeepQSM— using deep learning to solve the dipole inversion for quantitative susceptibility mapping. *NeuroImage*. 2019; 195: 373–383. [PubMed: 30935908]
52. Melià E, Raaijmakers A, Sbrizzi A, et al. A deep learning method for image- based subject-specific local SAR assessment. *Magn Reson Med*. 2020; 83: 695–711. [PubMed: 31483521]
53. Zaitsev M, Dold C, Sakas G, Hennig J, Speck O. Magnetic resonance imaging of freely moving objects: prospective real- time motion correction using an external optical motion tracking system. *NeuroImage*. 2006; 31: 1038–1050. [PubMed: 16600642]
54. Liu J, de Zwart JA, van Gelderen P, Murphy-Boesch J, Duyn JH. Effect of head motion on MRI B0 field distribution. *Magn Reson Med*. 2018; 80: 2538–2548. [PubMed: 29770481]
55. Brackenier, Y; Lucilio, C-G; Tomi-Tricot, R; , et al. Data-driven motion-corrected brain MRI incorporating pose dependent B0 fields; Proceedings of the 29th Annual Meeting of ISMRM [Virtual]; 2021. Abstract #0068
56. Wu X, Adriany G, Ugurbil K, Van de Moortele P-F. Correcting for strong eddy current induced B0 modulation enables two- spoke RF pulse design with parallel transmission: demonstration at 9.4T in the human brain. *PLoS One*. 2013; 8 e78078 [PubMed: 24205098]

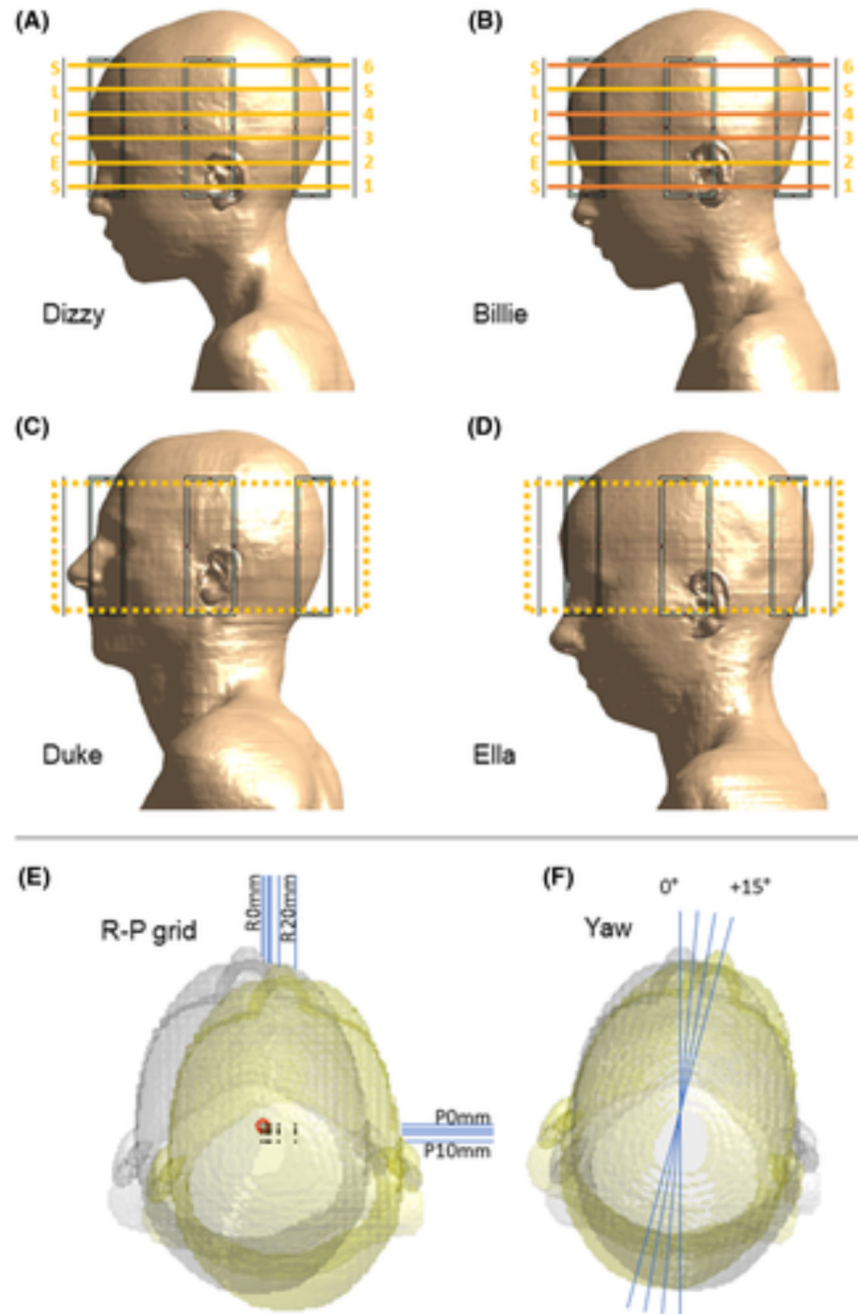
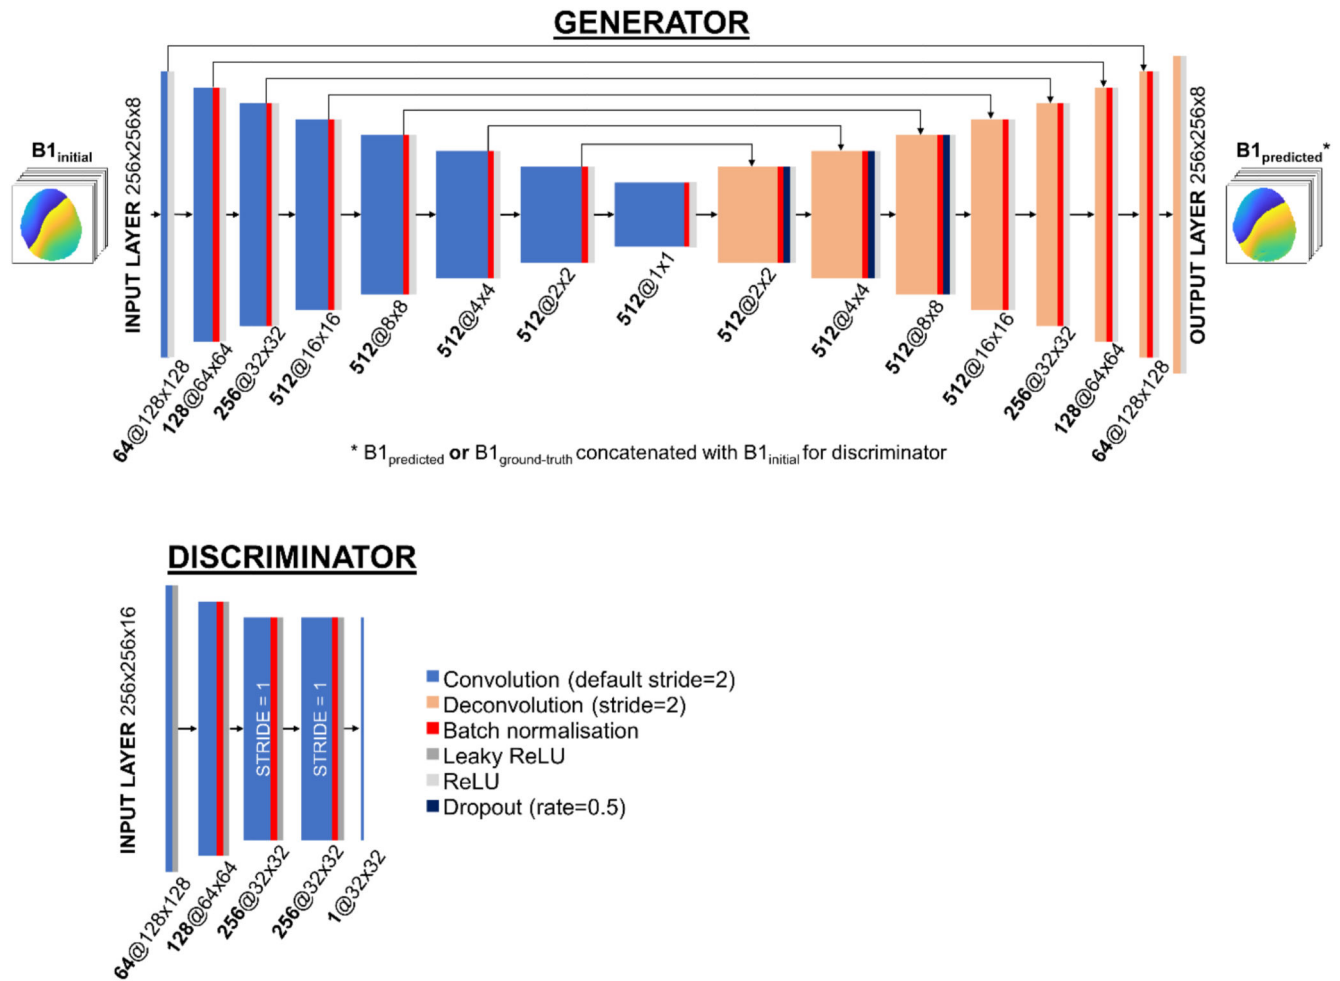


Figure 1. Simulation model setup.

(A–D) The four body models used in Sim4Life simulations. Ella and Duke (C, D) were used to generate training data for networks, Billie (B) was used for network validation and testing, and Dizzy (A) was used for testing only. Testing (including pulse design) was conducted at the six indicated slice locations. Validation slices were offset by about 4 mm from these, but within the same axial range. Slices in orange were also used for specific absorption rate (SAR) evaluations. All axial slices within the dashed slab were used for training. (E) Positions simulated for the R-P grid. The origin of the central position is

indicated with a red circle, whereas all other positions' origins are indicated with black dots. Axial displacements were all possible combinations of rightward (R) 0, 2, 4, 5, 10, and 20 mm and posterior (P) 0, 2, 4, 5, and 10 mm. (F) Yaw rotations were 5°, 10°, and 15°. The head at the central position (gray isosurface) and most extreme displaced position (yellow isosurface) are shown.

**Figure 2.**

Conditional generative adversarial network (cGAN) architecture. Generators were U-Nets with eight convolution and eight deconvolution layers, each with rectified linear unit (ReLU) activation. Discriminators consisted of five convolutional layers with ReLU activation. Square matrix size and number of filters (initially 64 for phase networks) are indicated beneath the layers. Convolution stride was 2 except where specified. Skip connections are shown with arrows. Dropout was applied at indicated layers (dark blue). Batch normalization (red) was used for phase networks, but not for magnitude networks. Filters for phase networks were 8×8 . Magnitude networks used double the number of filters, with filter size = 4×4 .

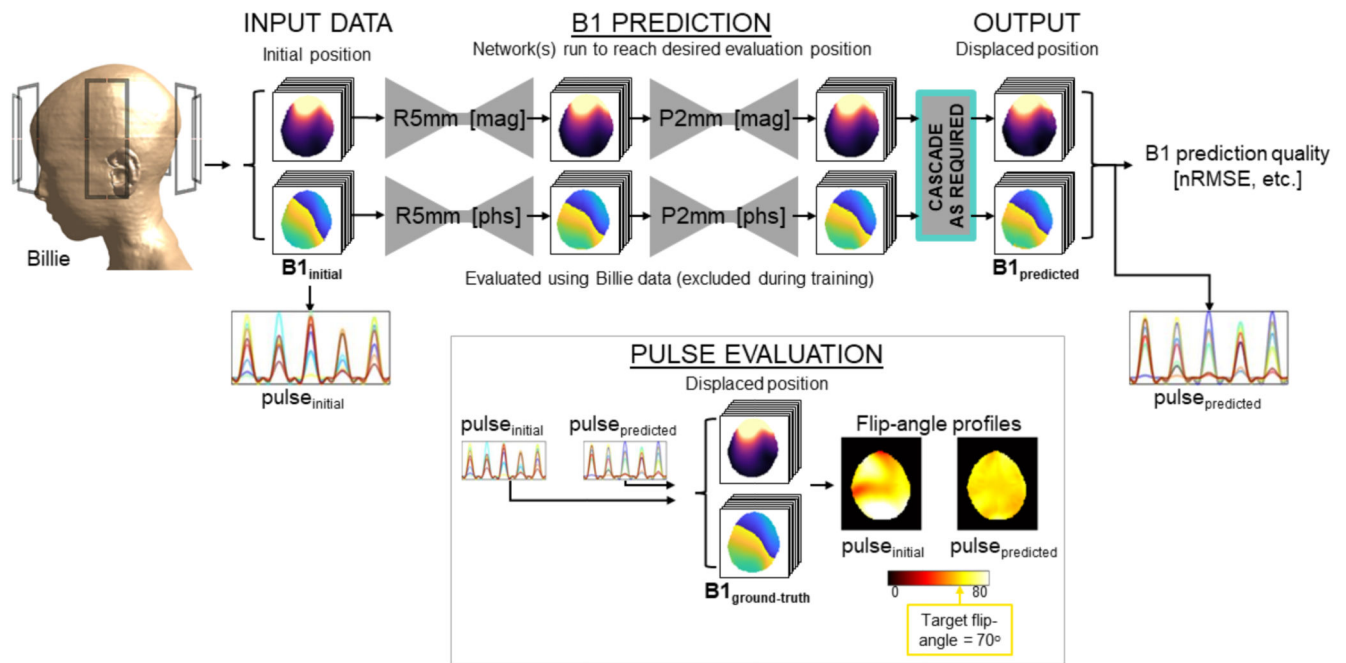


Figure 3.

Outline of the testing workflow. Simulated B_1^+ maps from the center position are input to the first trained generator. Generators were trained for small displacements but can be run sequentially (cascaded) until the desired off-central (displaced) position is reached in evaluations. Prediction quality is assessed by normalized RMS error (nRMSE) and voxelwise correlation with respect to the ground-truth (simulation output) displaced B_1^+ map. In addition, pulses designed using the initial B_1^+ map are compared with those designed using predicted maps, in terms of their excitation profiles following head motion.

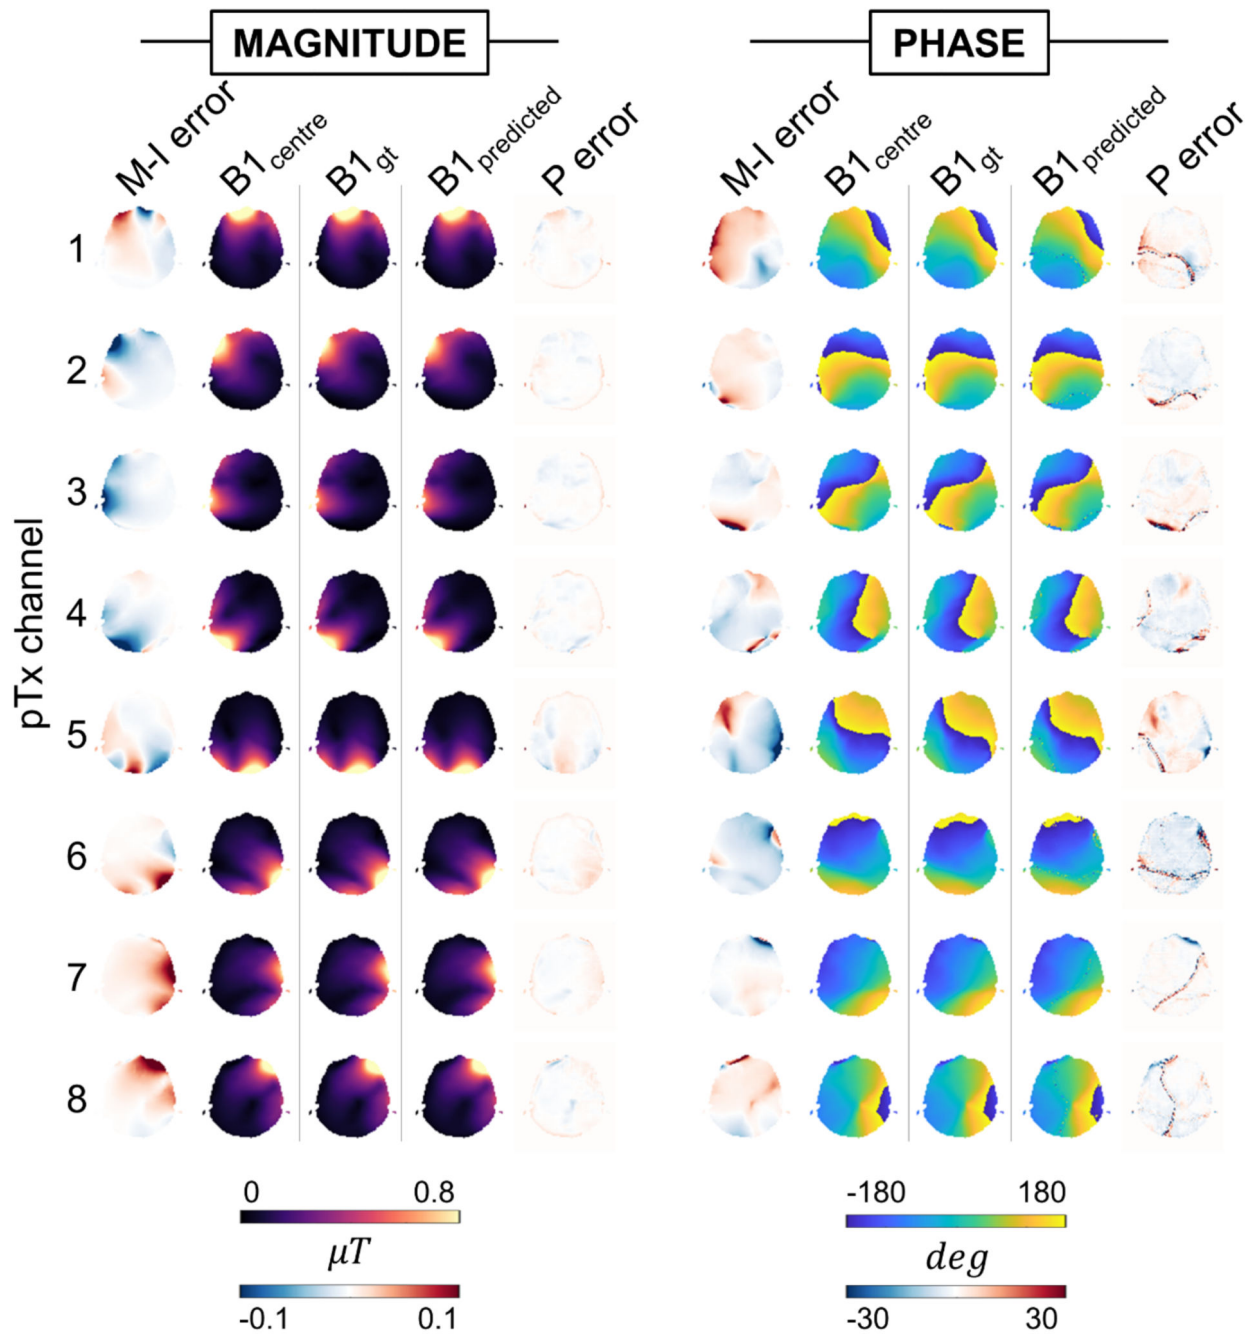


Figure 4.

Example magnitude and phase B_1^+ maps and error following a rightward displacement of 5 mm (slice location = 2). Motion-induced (M-I) error shows difference between simulation-output B_1 at the centered and displaced positions ($B_{1_{\text{initial}}}$ and $B_{1_{\text{gt}}}$, respectively). Prediction (P) error shows the difference between simulation-output $B_{1_{\text{gt}}}$ and generator-predicted B_1 ($B_{1_{\text{predicted}}}$). Motion-induced error (averaged across channels) for this example was 15.1% (magnitude) and 4.9° (phase), whereas mean prediction error was 3.2% (magnitude) and 3.5° (phase). Abbreviation: pTx, parallel transmission.

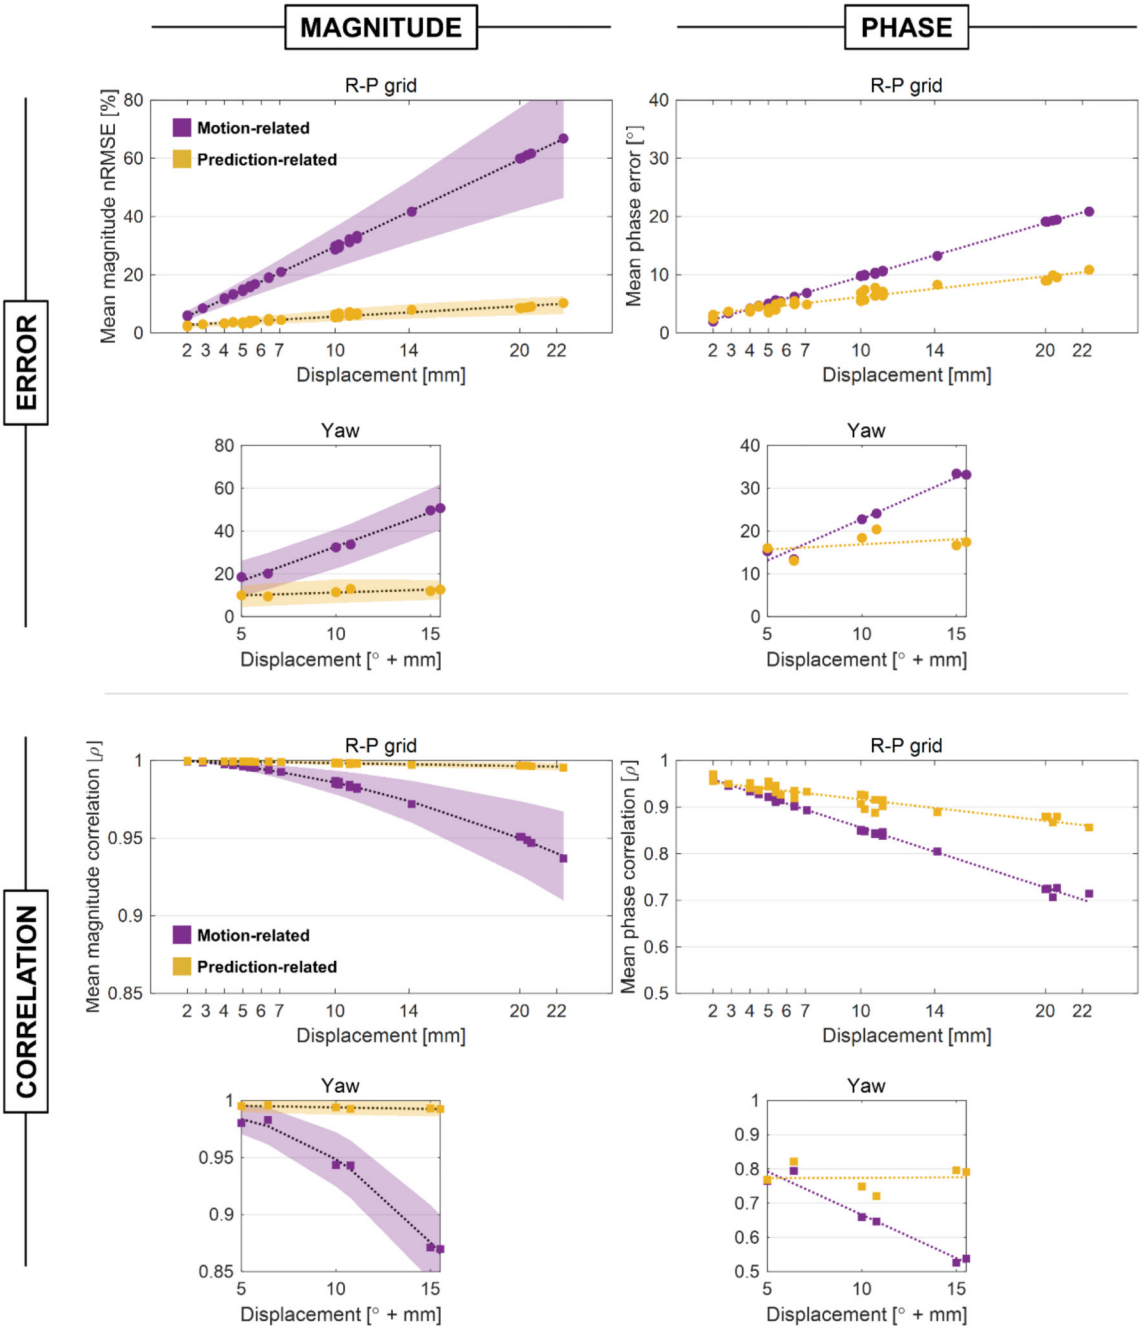


Figure 5. Error (nRMSE for magnitude, L1 norm for phase) and correlation coefficient (ρ) shown for magnitude and phase, averaged over Dizzy and Billie body models, channels, slices, and initial positions for each evaluated displacement. Translational displacements (the R-P grid) are shown in the large panels, while rotations (yaw) and combined rotation-translations (yaw plus a 4-mm translation) are shown in the smaller panels below (for the purpose of the x-axis, the amount of yaw rotation is treated as magnitude displacement; for example, yaw 5° plus 4-mm translation is shown at x = 6). The effects of motion are shown in purple,

while network-related prediction error is shown in yellow. The SD is shown as shaded regions for magnitude but is omitted for phase for clarity, as values were similar.

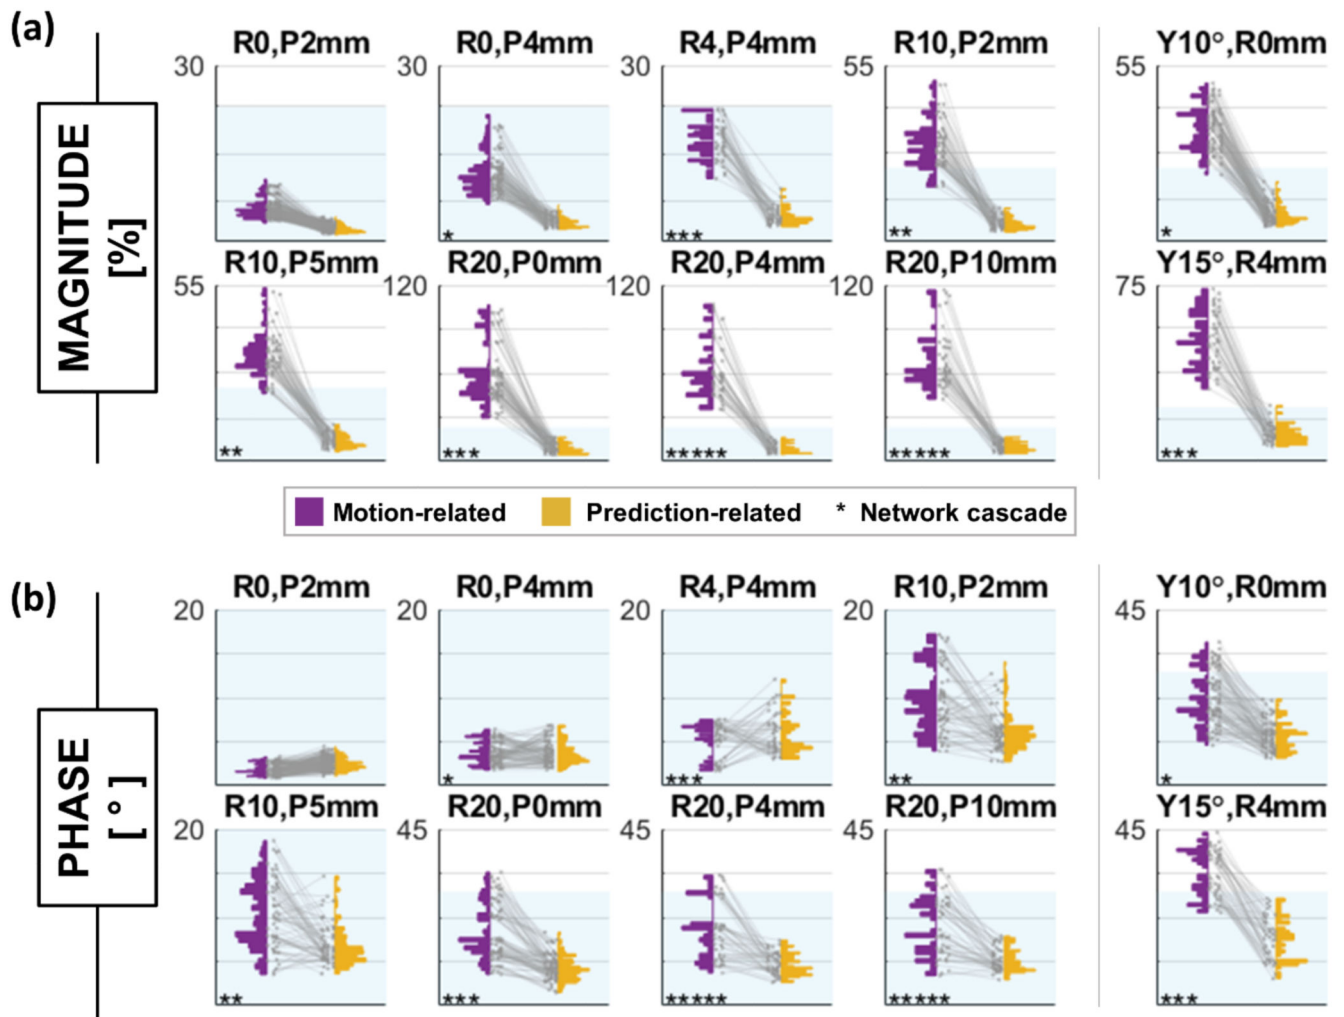


Figure 6.

B_1 error (nRMSE for magnitude [A], L1 norm for phase [B]) for all evaluations with the Billie model following 10 example displacements. Motion-related error is shown in purple, while error for predicted maps is in yellow. Asterisks indicate the number of network cascades required for evaluation. The blue-shaded region shows the maximum observed prediction error across all 35 displacements for the Billie model (consistent across panels).

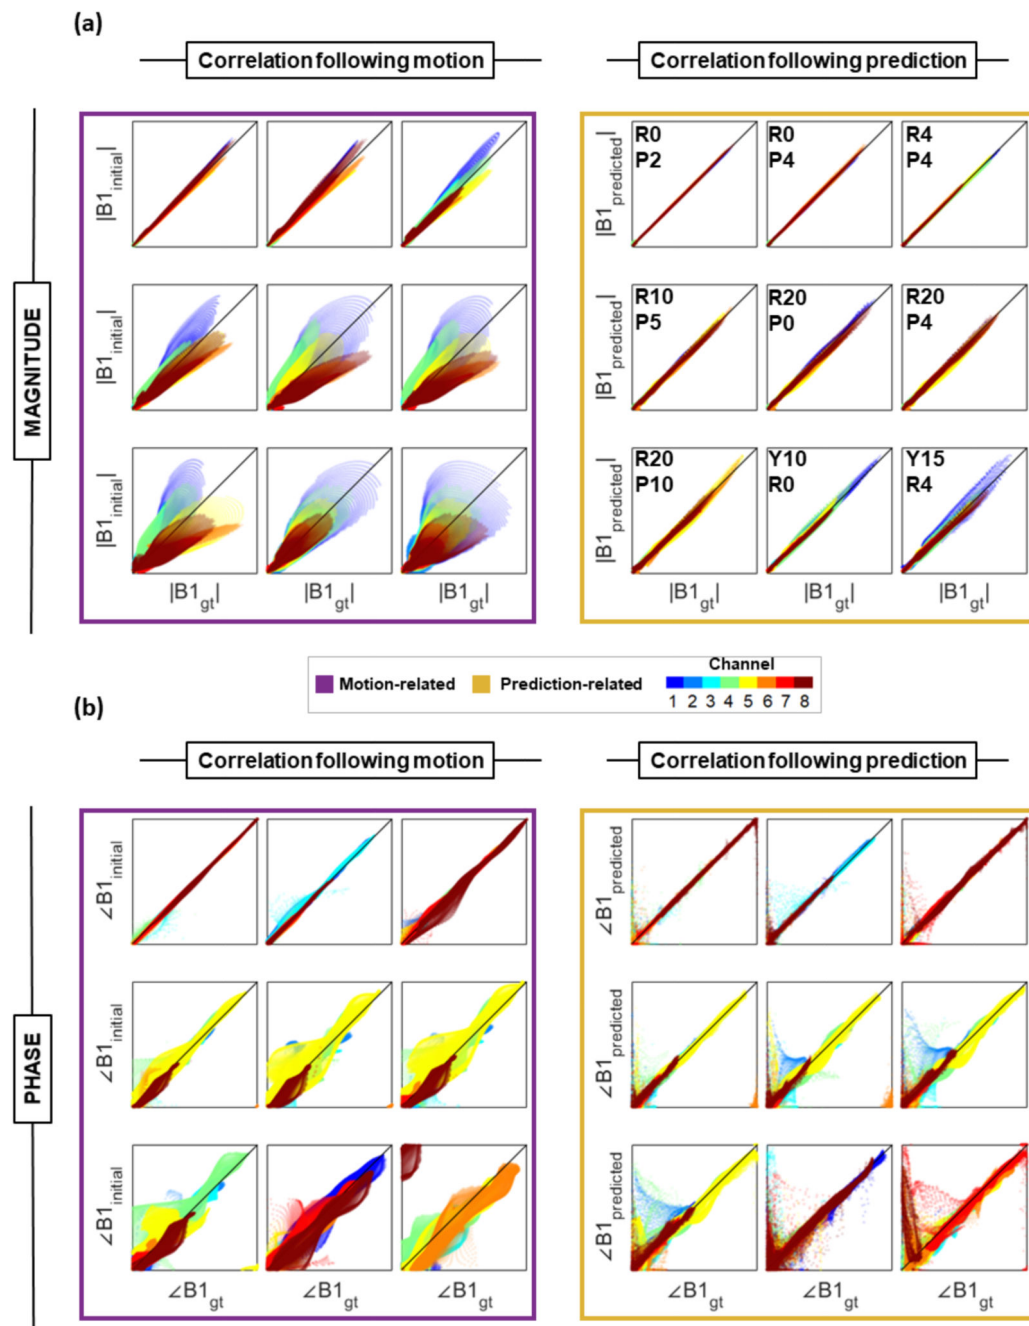


Figure 7.

Example voxel-wise correlations between $B1_{\text{initial}}$ and $B1_{\text{gt}}$ (left) and $B1_{\text{predicted}}$ and $B1_{\text{gt}}$ (right) for nine example displacements. The pTx channels are indicated by color. The x and y axes range between 0 and 3 μT for magnitude (A), and 0 and 2π for phase (B).

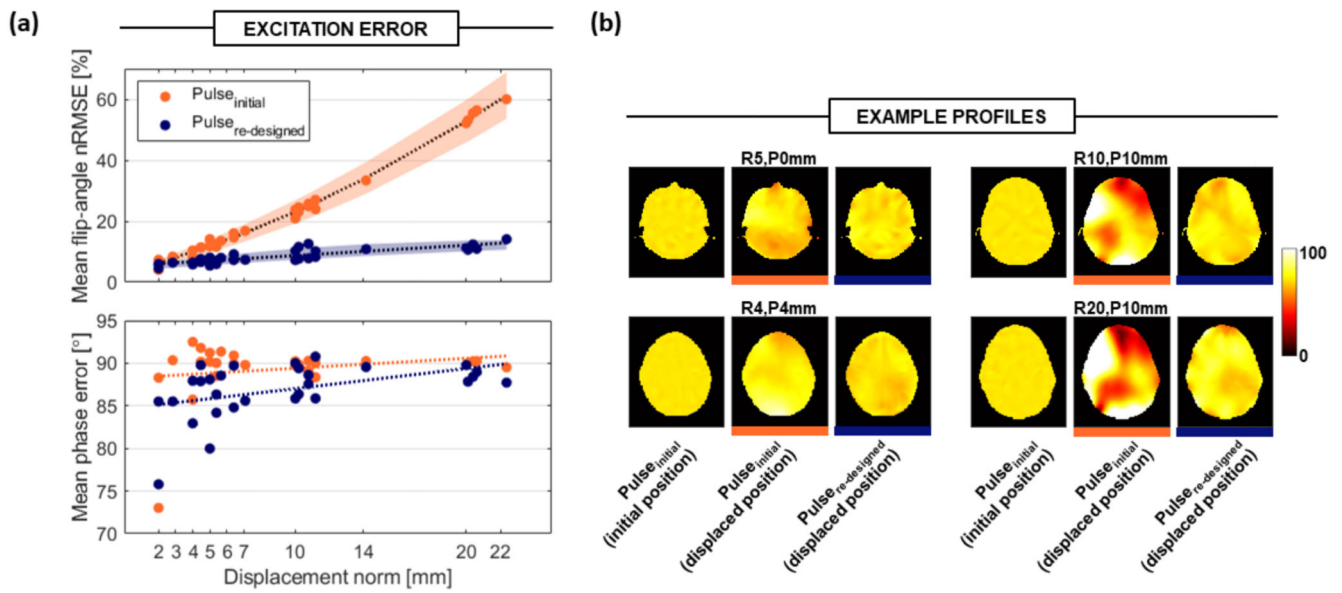


Figure 8. Excitation profile results for five-spoke pTx pulses following head motion.

(A) Mean flip-angle nRMSE (above) and phase RMSE (below) for excitation profiles, averaged over slices and initial positions for each evaluated displacement. Excitation pulses were five-spoke pTx pulses designed using either the initial position (pulse_{initial}) or predicted (pulseredesigned) B₁ maps. The SD is shown as shaded regions for magnitude but is omitted for phase for clarity, as values were similar. (B) Example flip-angle profiles produced by pulse_{initial} at the initial position, by pulse_{initial} at the displaced position, and by pulseredesigned at the displaced position.

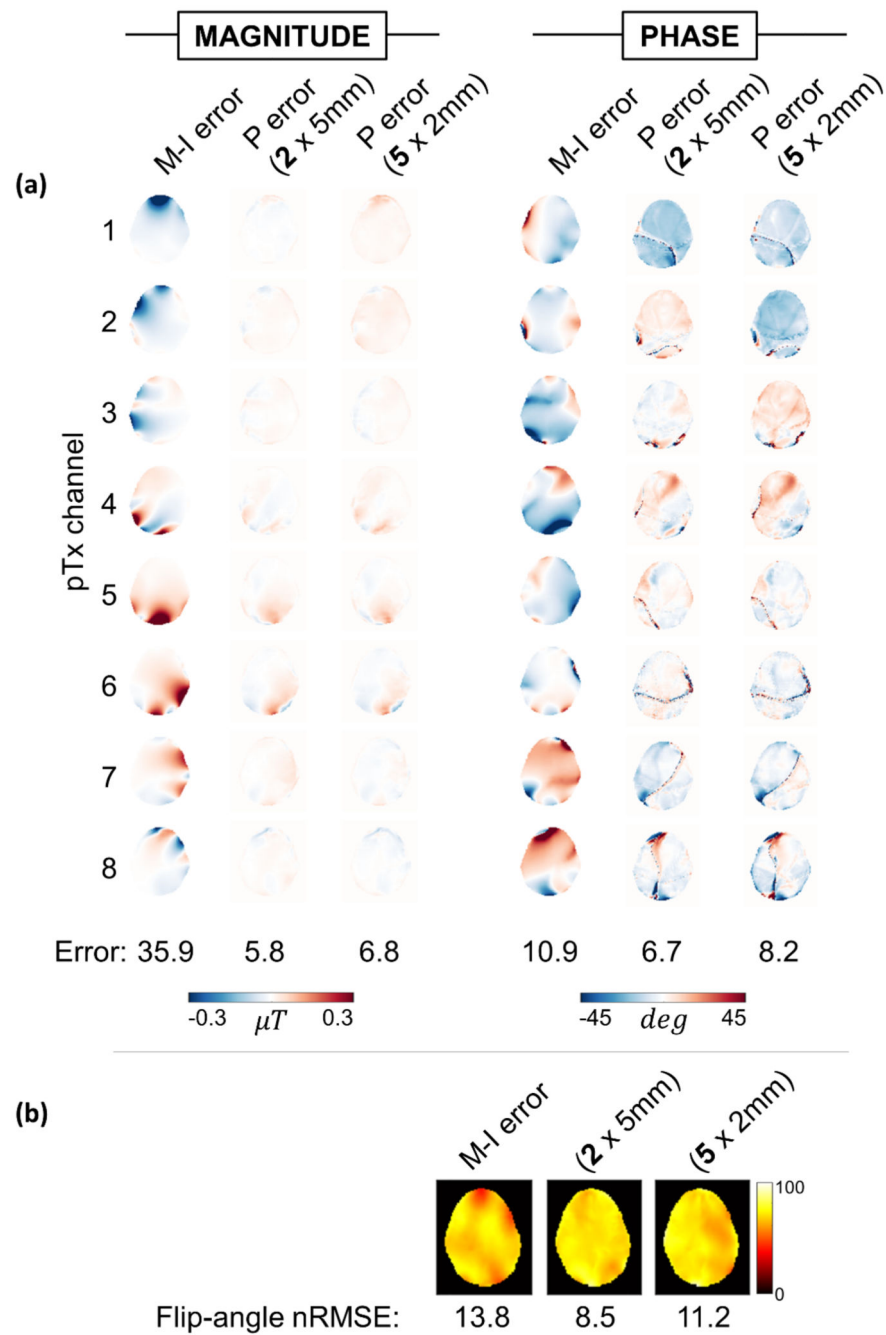


Figure 9. Effect of cascading the P2 mm network four times compared with cascading the P5 mm network once for evaluation at the R5, P10 mm position (along with the R5 mm network for the rightward component).

(A) Example motion-induced (M-I) error and prediction (P) error for both cascade approaches for magnitude (left) and phase (right). Error shown below maps is nRMSE (%) for magnitude and L1 norm ($^{\circ}$) for phase, both averaged over channels. (B) Comparison of flip-angle profiles and nRMSE for pulses designed using initial (left) and predicted maps using both cascade regimes (center and middle). Target flip angle is 70° .

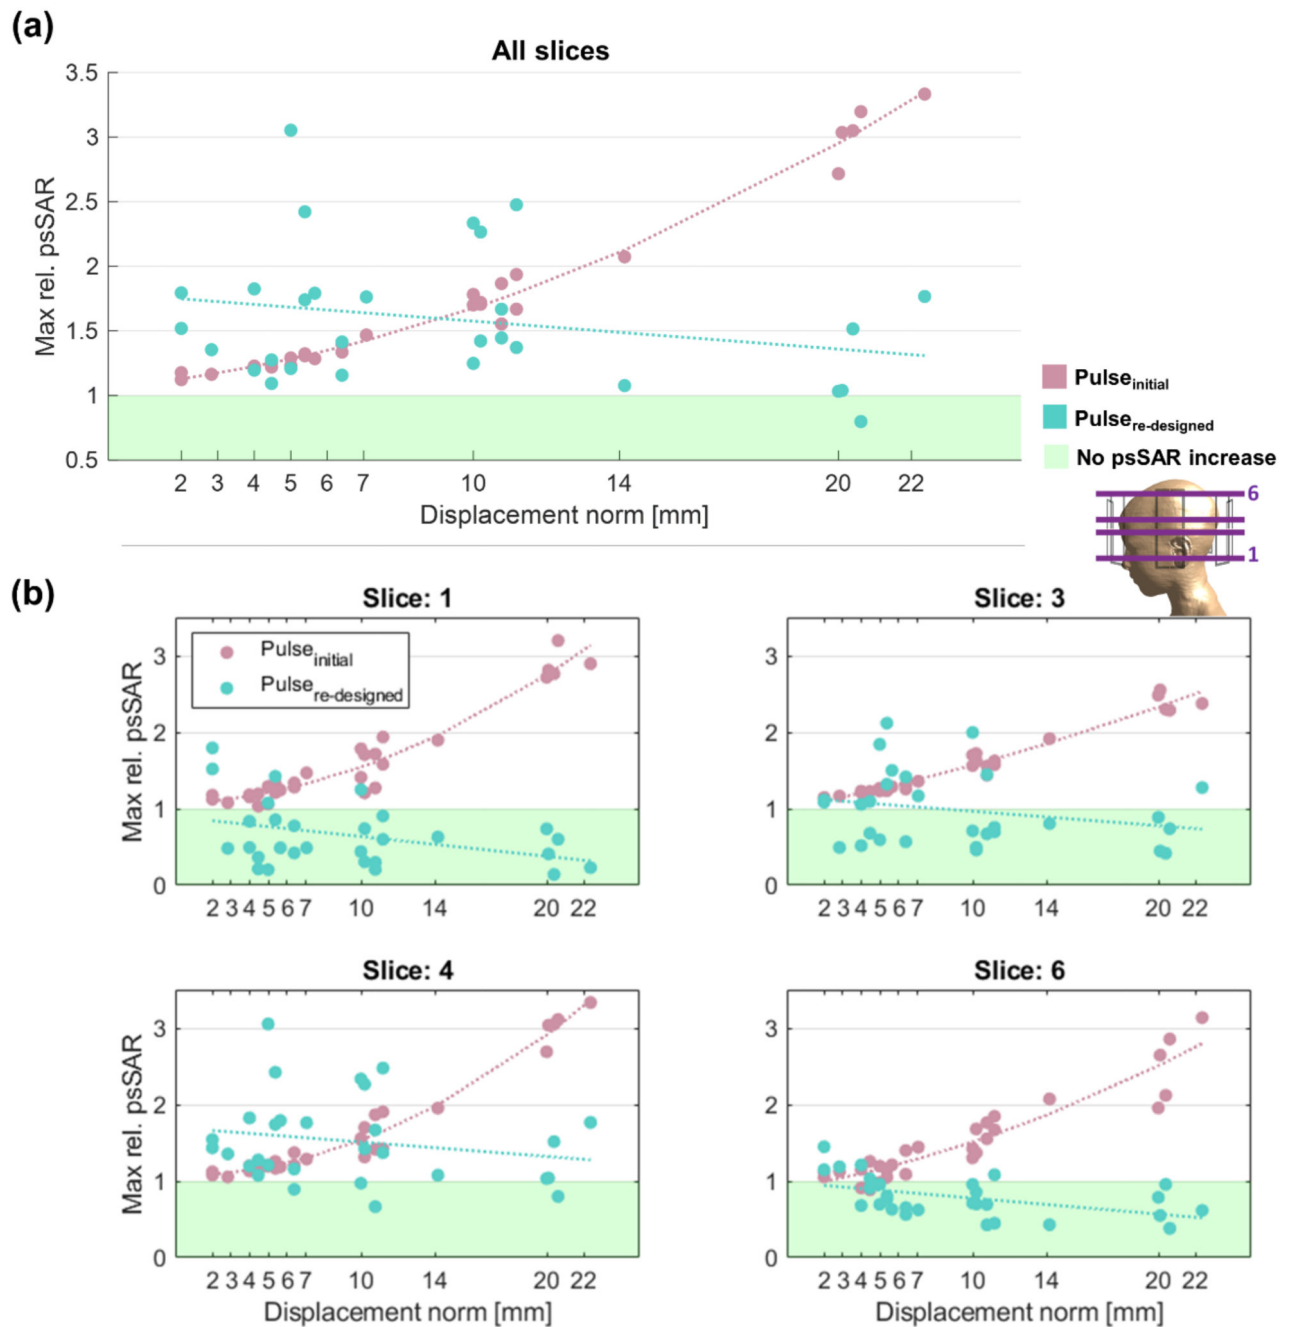


Figure 10. Peak 10g-averaged local SAR (psSAR) for pulse_{initial} and pulse_{re-designed} following motion

(A) Worst-case psSAR for both pulses across all observed slice locations. (B) Slice-wise worstcase psSAR for both pulses. Slice locations are indicated in the inset on the right. Vertical axes show relative psSAR, calculated as psSAR as a factor of that without motion using pulse_{initial} (psSAR_{center}). The green-shaded region shows psSAR at or below psSAR_{center} (i.e., indicating that psSAR did not increase following motion). Neither pulse included SAR constraint in the design.