

Clipped DeepControl: Deep neural network two-dimensional pulse design with an amplitude constraint layer

Mads Sloth Vinding*, Torben Ellegaard Lund

Center of Functionally Integrative Neuroscience (CFIN), Department of Clinical Medicine, Faculty of Health, Aarhus University, Denmark

ARTICLE INFO

Keywords:

MRI
DeepControl
Clipping
2D RF
Optimal control

ABSTRACT

Advanced radio-frequency pulse design used in magnetic resonance imaging has recently been demonstrated with deep learning of (convolutional) neural networks and reinforcement learning. For two-dimensionally selective radio-frequency pulses, the (convolutional) neural network pulse prediction time (a few milliseconds) was in comparison more than three orders of magnitude faster than the conventional optimal control computation. The network pulses were from the supervised training capable of compensating scan-subject dependent inhomogeneities of B_0 and B_1^+ fields. Unfortunately, the network presented with a small percentage of pulse amplitude overshoots in the test subset, despite the optimal control pulses used in training were fully constrained. Here, we have extended the convolutional neural network with a custom-made clipping layer that completely eliminates the risk of pulse amplitude overshoots, while preserving the ability to compensate for the inhomogeneous field conditions.

1. Introduction

In magnetic resonance imaging (MRI), the subject to be imaged is placed in a (main) magnetic field, denoted B_0 , nominally and typically of 1.5, 3, or 7 Tesla (T) formed by a large, superconducting solenoid. Hereby, certain nuclear-magnetic-resonance active atomic nuclei (e.g. of Hydrogen) form a small magnetization, which can be probed with the corresponding (Larmor) resonance frequencies of 64, 128, and 300 MHz, respectively. By a coil nearby, a second perturbing electromagnetic field, denoted B_1^+ , transmits radio-frequency (RF) pulses into the subject, e.g., to invert, excite or refocus the magnetization, as a means to efficiently exploit the magnetization's ability to induce a measurable electrical current in a receive coil nearby, which has an associated receive field denoted B_1^- . Of essence particularly for MRI, pulsed magnetic field gradients are used to encode spatial information into the signal. Magnetic field gradients stem from pulsed currents running typically in two Golay saddle coils and one circular Maxwell coil, covering the x -, y - and z -directions, respectively.

MRI has a plethora of different RF and gradient pulse schemes for different imaging applications. Simple fixed-shape RF pulses are typically one-dimensional (e.g., a Hamming filtered sinc-shape) and therefore just select the image slice (a slice of the imaged body) [1]. Neither the B_0 , B_1^+ , or B_1^- fields are perfectly homogeneous, and they depend partly on the subject being scanned. Variations are expected and may be seen in images as distortions, blurring, shading, etc. This communication concerns RF pulse design, where B_0 and B_1^+ fields

are part of the underlying math. The B_1^- field inhomogeneity and compensation thereof is another matter that will not be discussed further, but one account of this is for example the work of Caan et al. [2]. Simple RF pulses do not compensate spatial B_0 and B_1^+ field inhomogeneities, or at least not to the same extent as more advanced pulse designs. Advanced pulses may be selective in all three spatial dimensions as well as the spectral dimension if needed. For instance, 2D RF pulses are not slice-selective, but selective in the 2D image plane itself, and are therefore able to encode and compensate B_0 and B_1^+ field inhomogeneities in the image plane [3]. They may, e.g., be used to zoom the selection to a small region of interest, which the following image acquisition can focus on, resulting, e.g., in faster scanning with less energy deposition for the same resolution [4]. Yet, to compensate for subject-dependent field inhomogeneities, these tailored pulses need maps of these B_0 and B_1^+ fields as they need to plug into the math. While these maps can be measured rather swiftly nowadays [5], exploiting the field maps in tailored pulse designs typically require lengthy iterative numerical optimizations. In fact, there is an urge to achieve more difficult applications, e.g., multi-dimensional, multi-channel transmit at ultrahigh field (UHF). Meanwhile, we need to mitigate more and more demanding experimental challenges such as increasing B_1^+ inhomogeneities at UHF [3]. We must also comply with more and more demanding constraints, e.g., pulse power and local specific absorption rate (SAR). As a result, the pulse design methods often become too slow in practice and hamper clinical applicability [6].

* Corresponding author.

E-mail address: msv@cfin.au.dk (M.S. Vinding).




Fig. 1. The DeepControl method including the clipping layer. From the left, the image input layer loads a stack of three images of a given size into the network, mapping first of all the 2D RF pulse target area (here as a binary mask). Secondly, the B_0 field map shows the main magnetic field inhomogeneity. Thirdly, the RF coil's sensitivity through a B_1^+ field map. These maps are among three typical input parameters in 2D RF pulse design. The hidden layers perform 2D convolution followed by Swish layer activation functions. Before the final regression layer, there is a fully connected layer that matches the desired output array and the custom-made two-sided clipping layer that prevents the regression layer output from exceeding a user-defined limit by simple cutting. The regression layer contains the real and imaginary 2D RF pulse in a concatenated array. When using the network predicted pulse in a simulation or experiment, it must be accompanied by the given magnetic field gradient waveform shown in Fig. 6. This magnetic field gradient waveform was used, when building the supervised training library with the conventional 2D RF pulse optimization framework.

As in many fields and because of promising results, artificial intelligence in medicine is getting increased attention these years for solving various challenges. For example, image analysis [7], image segmentation [8], early neuro-degenerative disease detection [9], cancer management [10], automated capillary microscopy analysis [11], survival prediction in arrhythmia [12], stroke treatment outcome [13], and pandemic mitigation [14]. We recently introduced a deep learning neural network method to challenge long pulse design computations [15]. It used supervised learning of a vast library of 2D RF pulses made from different numerical optimization approaches, e.g., the optimal control methods, we have reported earlier [6,16,17]. Other MRI pulse types addressed with optimal control include, e.g., non-selective refocusing [18], high-resolution multi-slice excitation [19], and contrast-enhancement [20]. In our case, the neural network pulse prediction times were on average 7 ms, which was more than three orders of magnitude faster than the optimal control computations [15], and thus approached a real-time push-button solution. In the follow-up, we extended the neural network with convolution layers to better facilitate B_0 and B_1^+ field maps, and we demonstrated compensation of these field inhomogeneities in UHF experiments [21]. In both studies, however, the RF pulse amplitudes were constrained only in the optimal control training library. Meaning that we relied on a fully constrained training library with respect to pulse amplitude, and the network's ability to deep learn this limit. I.e., we had no handle on pulse amplitude in the network pulse predictions. Overall, the statistical risk of pulse amplitude overshoot was small. In the last study, for example, only 2% of the test cases were overshooting. The amount of overshoot was not alarming either, and we described methods to handle overshoots [21]. However, as we have later developed networks for RF pulses with stronger excitation flip angles of the magnetization (going from 30° in Ref. [21] to the 90° herein), we observed an unsatisfying greater risk of pulse amplitude overshoot. As we suggested in Ref. [21], this communication reports on a convolutional neural network that includes a custom-made clipping activation function layer. It completely prevents pulse amplitude overshoots without loss in performance and with the preserved ability to compensate for inhomogeneous field conditions.

2. Method

The method presented here is an updated version of that of Ref. [21]. In short, a vast supervision library is established, see Section 2.1,

where each entry has an input and an output. The input mimics a pulse design situation through three images: a 2D RF target pattern, a B_0 map, and a B_1^+ map. The output for the same library entry contains the 2D RF pulse that, based on the input, is computed by a conventional pulse design method, e.g., optimal control (Section 2.1.1). The convolutional neural network, Section 2.2, was designed first of all to handle the specific input and target array sizes, plus the novel clipping layer was included (Section 2.2.1). The hidden layers, activation functions, etc., were adjusted by monitoring the deep learning progress during several trial runs. All computations were performed in MATLAB R2021a (Mathworks, Natick, MA, USA) on a 28x2.2 GHz Intel Xeon Gold 5120 (Intel Corporation, Santa Clara, CA) workstation with 384 GB RAM.

2.1. Supervision library

The total library consisted of five hundred thousand individual entries. The input part of each entry consists of three maps of 64-by-64 pixels. An input entry is shown in Fig. 1 (left). Here we show a realistic case to illustrate the DeepControl method in operation. The first maps out (spatially) the target region of interest as a binary mask, i.e., what the 2D RF pulse is aimed to select for zoomed imaging. The other two are the B_0 and B_1^+ maps. In the library, however, the three images are not MRI related as such. The target patterns are made by placing ten points randomly on the grid as corners in a polygon. This polygon is then morphologically dilated and eroded twice with disk-shaped structure elements of radius 8, 12, 3, and 4 pixels, respectively. The randomness implemented here assures the user can select arbitrary target regions when employing the method [15,21,22].

As seeds for B_0 and B_1^+ field maps, we used photographs from ImageNet [23] cut to 64-by-64 pixels. This choice was made to induce a large variability in map patterns that for typical MR images would seem random. The B_0 photograph seeds were firstly subtracted by their median pixel value and then normalized. The normalized map was then multiplied by a random value selected from the ranges [-600 Hz, -50 Hz] and [50 Hz, 600 Hz]. These values correspond to typical maximum offset values observed in measured B_0 maps. The relation to frequency units is given by the Larmor precession frequency equation $f_0 = \frac{\gamma}{2\pi} B_0$, where γ is the natural gyromagnetic ratio constant for the given nucleus (here the ^1H nucleus). The excluded range from

(−50 Hz, 50 Hz) assured that the B_0 maps were not unrealistic and uncomplicated.

A B_1^+ field rotates the magnetization by an amount denoted a flip angle (FA). As such, the target map displays a region with a desired FA. In this study, a value of 0 or 1 in the target map corresponds to a desired, nominal FA of 0° or 90°, respectively. The actual, achieved FA (depending on B_1^+), though, is found by experiment or some form of RF pulse simulation, e.g., the generally applicable Bloch equations, or the Fourier transform for small FAs. An inhomogeneous B_1^+ may therefore not achieve the nominal FA everywhere, if not attended to in the RF pulse design. The B_1^+ photograph seeds were just normalized to the internal between [0, 1]. A value of 1 in such a sensitivity map corresponds to a point, where the B_1^+ field is capable of achieving the nominal FA.

A common superellipse mask was further applied to the maps to simulate potential object boundaries, with the outside area working as a “don’t-care” region that is not included in the conventional pulse optimizations either. Because the target pattern had values of 0 or 1, and B_1^+ maps had values also between 0 and 1, the B_0 maps were furthermore downscaled by 600 Hz to achieve values between −1 and 1 (depending on the random generator). This was done to avoid mixing different dynamic ranges in the library that could lead to numerical instability.

2.1.1. Optimal control

As a supervision library, the input maps described in the previous section need an output part as well. This is in our studies 2D RF pulses computed with optimal control. The optimal control applies, through the Bloch equations, all physically relevant parameters, which the convolutional neural network is blind to. The target maps are translated to a nominal FA of 90°, and the B_0 field maps are rescaled back by 600 Hz when drawn from the library.

The B_1^+ field maps are plugged in as sensitivity maps and stay unitless. The 2D RF pulses will then be computed in field strength or frequency units by the relation $f_1 = \frac{\omega_1}{2\pi} = \frac{\gamma B_1}{2\pi}$. The frequency unit refers to the *rate* by which the magnetization is rotated as a result of the applied B_1^+ field. In turn, we can impose limits on the 2D RF pulse amplitudes (real and imaginary channels) in frequency units. Our RF amplitude constraint was $c' = 1$ kHz, meaning the real and imaginary channels could vary freely in the range $[-c', +c']$. The choice of using frequency units etc. is one common convention. Among other conventions, it is also common to measure B_1^+ field maps in Tesla per Volt, which directly refers to the electric potentials applied to the RF coils by the RF amplifiers. RF pulses are then output and constrained in Volt [6].

The maps of 64-by-64 pixels are set up to span a 25-cm square. This is done by the underlying magnetic field gradient waveform, shown later, that encodes spatial information. The magnetic field gradient waveform dictates the minimum pulse duration, as it assures a full Nyquist sampling in the frequency space and conforms to amplitude and slew-rate limits. For the present setup, the total pulse duration was 7 ms. The time-step width was 10 μ s, i.e., resulting in 700 time steps. However, the total number of bins to compute was 2×700 for the real and imaginary channels.

Each 2D RF pulse optimization was allowed 50 iterations and our workstation computed 25 pulses in parallel. We used the *midpoint* optimization gradient of Ref. [16], where gradient precision tradeoffs were reported considering generation of DeepControl libraries. It also includes investigations on the time-step width, which influence the number of time steps for a given pulse duration.

2.2. Convolutional neural network

The presented deep neural network shown in Fig. 1 was inspired by the deep neural network of Ref. [21]. The first image input layer imports the three maps of 64-by-64 pixels. This is handed over to three

2D convolution layers with filter sizes 5, 3, and 2, and with 80, 600, and 60 filters in each, respectively. Each 2D convolution has a stride length of 2 in both directions, and each 2D convolution layer is followed by a Swish activation function layer. After the three convolution layers, there is a fully connected layer, which must have 1400 neurons because it must match the output of the final regression layer, which is the 2D RF pulse (real and imaginary channels). In between the fully connected layer and the regression layer, we have inserted the clipping layer, described in the next section, and marked in Fig. 1 by the red dashed lines.

2.2.1. Clipping layer

The clipping activation function layer serving as the main contribution in this work was made by modifying the function files of MATLAB’s clipped Rectified Linear Unit (ReLU) layer for ease of implementation. For comparison, the ReLU layer function is

$$f^{\text{ReLU}}(x) = \max(0, x) = \begin{cases} 0, & x < 0 \\ x, & x \geq 0 \end{cases} \quad (1)$$

and the Swish layer function is

$$f^{\text{Swish}}(x) = \frac{x}{1 + e^{-x}} \quad (2)$$

The clipped ReLU layer is:

$$f^{\text{clippedReLU}}(x, c) = \min(\max(0, x), c) = \begin{cases} 0, & x < 0 \\ x, & 0 \leq x < c \\ c, & x \geq c \end{cases} \quad (3)$$

where c is a clipping value. And the modified, two-sided clipping version used herein to constrain RF amplitudes is

$$f^{\text{Clipping}}(x, c) = \min(\max(-c, x), c) = \begin{cases} -c, & x < -c \\ x, & -c \leq x \leq c \\ c, & x > c \end{cases} \quad (4)$$

This type of clipping (or sometimes clamping) function was also used for clipping, e.g., deep neural network weights [24], and Adam [25] optimization gradients [26].

When computing the pulse library with optimal control, we imposed the hard constraint, c' (Section 2.1.1), on the RF amplitude (real and imaginary channels). This was done through MATLAB’s minimization function *fmincon*, where upper and lower bounds are enforced strictly. An alternative to this is a weighted spill-out norm square penalty in the objective function, with a corresponding penalty gradient and Hessian [27].

When the pulse library is imported into the deep learning framework, we normalize the RF waveforms by c' such that they vary between −1 and +1 for numerical stability (similar to what we did for the image inputs). Accordingly, the clipping value $c = 1$.

In the end, pulses predicted from the neural network are converted back to physical units through the c' factor.

2.3. Deep learning

We trained the deep neural network in Fig. 1 with the stochastic-gradient-descent-with-momentum algorithm [28]. The momentum factor was 0.95. The learning rate was constant at 0.006 and the L_2 -regularization factor was 1×10^{-4} . The total library of 500k cases was divided into fractions of 0.8 : 0.1 : 0.1, which were used for training, validation, and testing, respectively. The minibatch size was 1280 and we trained for 1000 epochs. We used a Tesla P100 GPU (Nvidia, Santa Clara, USA) on the workstation for the deep learning.

For comparison, we also performed the same training on the neural network, where the clipping layer was deactivated. Here, we call our network DeepControl, and we use the terms “clipped” or “non-clipped” to distinguish between the situations, where the clipping layer was activated or de-activated during training, respectively.

We stored the intermediate non-clipped networks for every epoch as checkpoints to monitor trends of pulse amplitude overshoots as deep learning improves the network.

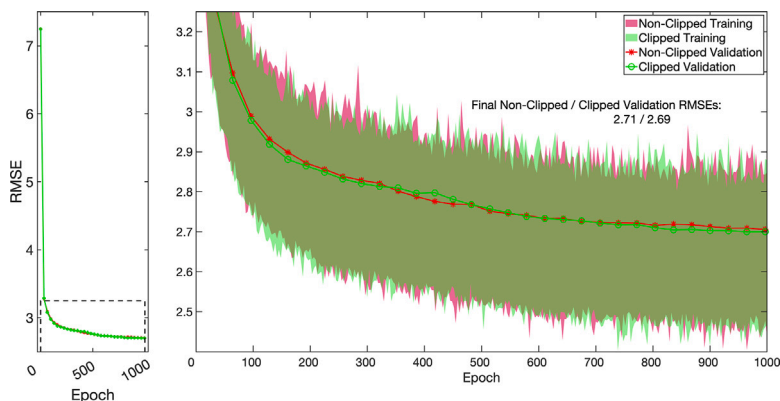


Fig. 2. Deep learning root-mean-square error (RMSE) as a function of epoch number for both the clipped and non-clipped DeepControl networks. Left, the RMSE of the validation subset. The dashed box signifies the zoomed view of the right plot, which also contains the RMSE of the training batches.




Fig. 3. Histograms of the maximum absolute real and imaginary RF amplitudes. Left, the non-clipped DeepControl network against the clipped DeepControl network. Center, the clipped DeepControl network against the optimal control cases. Right, the non-clipped DeepControl network against the optimal control cases. The green dashed lines are the desired trends. Data are from the library test subset.

2.4. Realistic field maps

While we perform our statistical analysis with the library test subset of 50k cases, we also tested each network (clipped and non-clipped) with a set of realistic field maps adopted from Ref. [29]. They are shown by example in Fig. 1. Here, the target pattern was a 15-by-16 pixels hand-drawn rectangle on a 64-by-64 pixels canvas in Paintbrush (Apple Inc., Cupertino, CA, USA).

3. Results

Generating the library takes around one week. The subsequent deep learning progress of the clipped and non-clipped networks, shown in Fig. 2, lasts approximately one and a half to three days for one thousand epochs. This conservative estimate is based on the fact that we ran multiple deep learning processes on the GPU simultaneously including various other demanding tasks occasionally. The two training sessions (clipped and non-clipped) ended with near identical root-mean-square errors (RMSE) of the library validation subset.

The test subset of the supervision library was used to probe the network performances with unseen data. With each network (clipped and non-clipped), we predicted 2D RF pulses from each test input and contrasted these against each other and the corresponding optimal control pulses.

The maximum absolute RF values are contrasted in Fig. 3. For these plots, we pooled the real and imaginary RF channels together. For example, ordinate counts above 1 kHz in the left and right plots correspond to either the real or imaginary or both RF channels of the non-clipped network pulses having values outside the range $[-c', +c']$. If more than one instance of the RF channels spilled out, we only counted and plotted the maximum spill-out value. In total, 18.3% of the non-clipped network pulses overshoot the $c' = 1$ kHz limit somehow. The amount of overshoot was 84 ± 78 Hz (mean \pm standard deviation in the test subset). From the non-clipped networks stored at each epoch, we have estimated the amount of overshoot to be $17.6 \pm 1.2\%$. Figure S1 in the Supplementary Material shows the test case overshoot percentage as a function of epochs together with a histogram. The right panel of Fig. 3 shows that the non-clipped network for amplitudes lower than 0.75 kHz follows optimal control cases quite well. Above 0.75 kHz and beyond 1 kHz the non-clipped network finds solutions with relatively higher peak amplitudes than optimal control. The clipped network (center panel of Fig. 3), as expected, piles up just below the 1 kHz limit. The clipped and non-clipped networks (left panel of Fig. 3) are quite similar, until around the 1 kHz limit. For reference, the test subset pulse mean powers are displayed in Figure S2 in the Supplementary Material.

The Bloch simulations of the pulses lasted a few hours. We measured the normalized root mean square errors (NRMSE) of the resulting excitation profiles with respect to the target excitation profiles, as shown in Fig. 4. In conventional pulse optimization, the NRMSE or




Fig. 4. Histograms of the normalized root-mean-square errors (NRMSE) of the magnetization components computed with Bloch simulations. Left, the non-clipped DeepControl network against the clipped DeepControl network. Center and right, the clipped and non-clipped DeepControl networks against the optimal control cases, respectively. Data are from the library test subset.

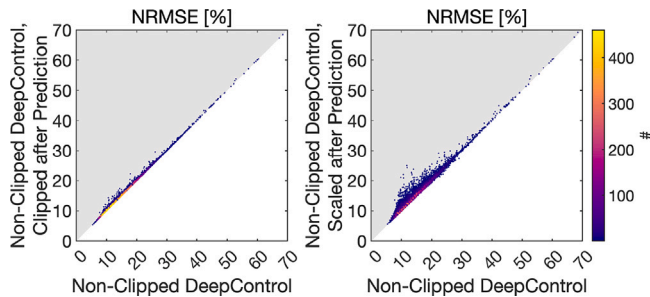


Fig. 5. NRMSE effects of regulating overshooting 2D RF pulses from the non-clipped DeepControl network, by clipping (left) or down-scaling (right) the 2D RF waveforms after pulse prediction. For clarity, the histograms only contain the overshooting pulses from the library test subset.

similar merit of the magnetization profile constitutes a direct measure of the optimization performance. In this framework, the NRMSE of the magnetization profile is indirect, because the magnetization profile is not part of the deep learning objective. The deep learning measures the RMSE of the network-predicted pulses with respect to the library pulses. However, Fig. 4 (center and right panels) show a good correspondence between the DeepControl and optimal control magnetization responses. The (left) panel shows a near identical performance between the two clipped and non-clipped networks. The mean and standard deviation NRMSE values were $10.59 \pm 4.47\%$, $11.39 \pm 4.70\%$, and $11.40 \pm 4.69\%$ for the optimal control, clipped DeepControl, and non-clipped DeepControl test cases, respectively.

Although the clipped network solves the problem of pulse overshoots, we investigated the effects of regulating pulses from the non-clipped network that overshoot the limit through two options as proposed in Ref. [15]. One option is to clip the pulse manually. If the internal checks of the MRI system do not block the pulse, the RF transmit chain will likely perform the clipping at some point. Another option is to down-scale the entire pulse waveform. The effects of performing clipping and scaling are shown in Fig. 5. Clipping and scaling after pulse prediction increased the NRMSE in the majority of the cases. Scaling had the greatest (negative) effect. The mean and standard deviation NRMSE values increased to $11.41 \pm 4.70\%$ and $11.49 \pm 4.76\%$ by clipping and scaling, respectively.

On the workstation, the pulse prediction times (mean \pm standard deviation in the test subset) for the clipped and non-clipped networks were 5.8 ± 1.3 ms and 3.9 ± 0.3 ms, respectively. The corresponding optimal control pulse computation times were around 20 to 30 s.

The video in the Supplementary Material displays a random selection of 10k simulations from the library test subset. It shows the photograph-based B_0 and B_1^+ field maps, and the randomly generated target areas, i.e., from the library input section. It also shows the 2D RF pulses predicted by the clipped and non-clipped networks,

and the corresponding optimal control pulses, i.e., the latter from the library output section. Finally, the simulated magnetization responses are shown.

3.1. Realistic field maps

The DeepControl pulses predicted for the realistic field maps are shown together with the optimal control pulse in Fig. 6. We observe a pulse overshoot of 255 Hz of the imaginary RF waveform for the non-clipped pulse. We found (not shown) that this pulse overshoot vanished if the target mask was trimmed to 13-by-14 pixels instead. The clipped network is as shown not overshooting the limit.

The bottom panel of Fig. 6 shows the magnetic field gradient waveform that in simulations and the optimal control framework is underlying all 2D RF pulses, i.e., it is played simultaneously to the RF pulses. While the magnetic field gradient waveform could have been further optimized in the optimal control and correspondingly been subject to a neural network output, we chose not to pursue this because this particular oscillating waveform is already time-optimized [30].

The pulses of Fig. 6 return the magnetization profiles shown in Fig. 7 also displaying the realistic field maps. We see a near-identical performance between the three contestants. However, regulating the pulse amplitude by clipping or scaling, as discussed in the previous section, influenced the magnetization patterns as shown in Fig. 8.

4. Discussion

A clipping layer has been introduced to the DeepControl neural network, that effectively prevents pulse amplitude overshoots. As compared to the same network except with the clipping layer deactivated, there is no significant difference in performance. Both networks (clipped and non-clipped) also yield almost the same performance as the conventional method of optimal control. The gain is a speed gain in pulse computation time by a factor of > 3000 .

This communication has mainly concerned the prevention of RF amplitude overshoots internally in the clipped neural network. We also investigated how regulating RF amplitude overshoots of a non-clipped network affected the magnetization responses. As proposed in Ref. [15], we applied clipping or scaling of the overshooting pulses. Our non-clipped network produced overshooting pulses in 18.3% of test cases. However, for the limit of 1 kHz, the amount of overshoot was on average below 0.1 kHz. For the type of magnetic field gradient waveform as shown in Fig. 6, the accompanying 2D RF pulse waveforms often appear spiky, with high transients. However, the highest amplitudes typically happen near the end, where the bulk magnetization is flipped. Therefore, manually clipping such 2D RF pulses does not, for the majority of our test cases, yield huge increases in NRMSE. Only a small fraction of the pulse waveform is clipped. Yet still, it may lead to changes in the magnetization profiles of around 10% as shown in one example. Down-scaling pulses to fit within the boundary, a common

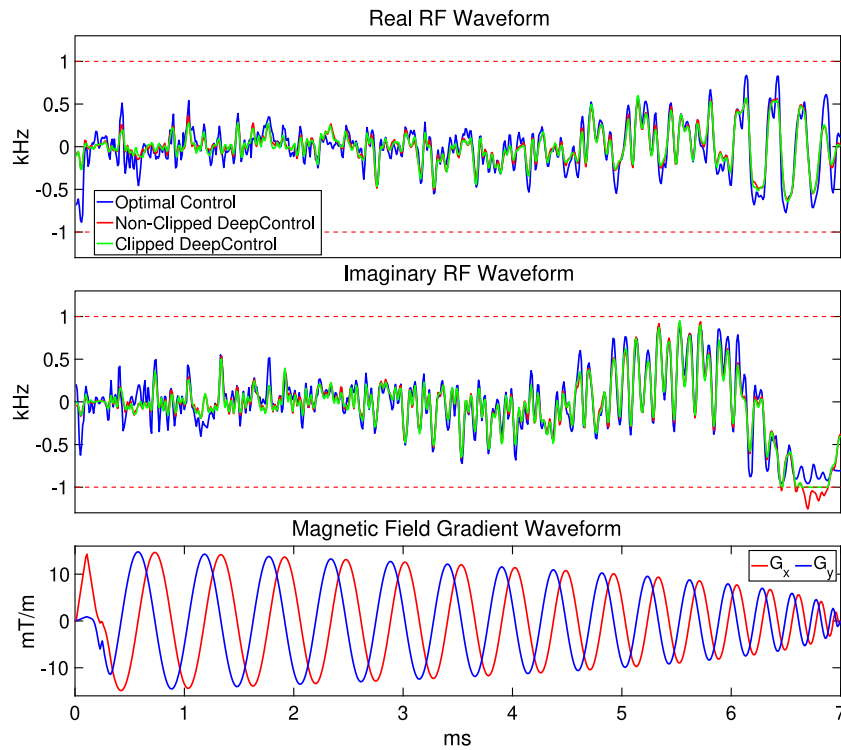


Fig. 6. The 2D RF pulses computed with realistic field maps, see Fig. 7. Top/central panel, the real/imaginary RF waveforms: optimal control (black), non-clipped DeepControl (red) and clipped DeepControl (green). The bottom panel shows the magnetic field gradient waveform accompanying the RF pulses in experiments for spatial encoding. The mean power of the optimal control, the non-clipped and the clipped DeepControl pulses were 0.23, 0.18, and 0.17 (kHz)², respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

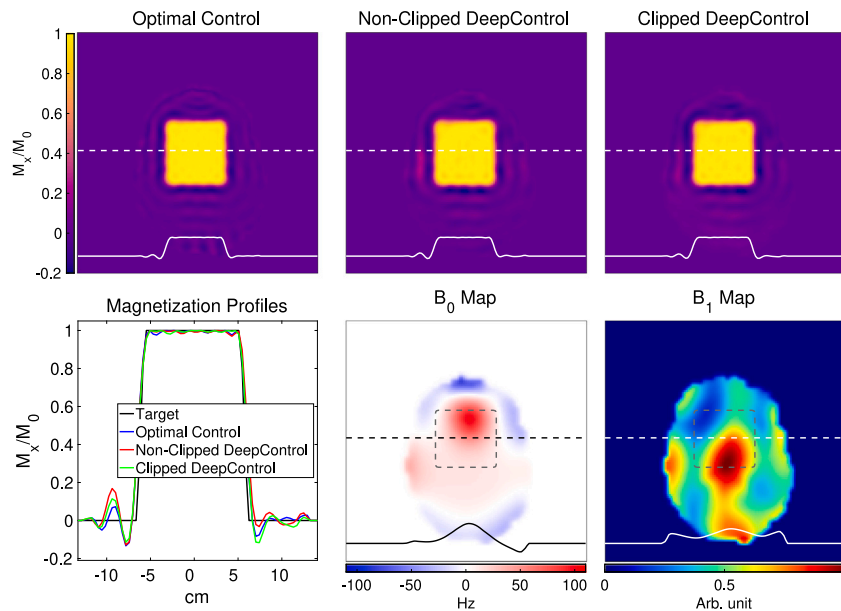


Fig. 7. Excitation maps (upper row) of the optimal control (left), non-clipped DeepControl (center) and clipped DeepControl (right) pulse responses. The dashed lines show, where the profiles (full lines) are from. They are overlaid in the lower left panel. The lower center and right panels show the B_0 and B_1^+ field maps, including the profiles and also the dashed frames signifying the target area.

way to change the desired FA, lead to higher NRMSE increases, mainly because we change the entire pulse waveform. As also suggested [15], we could also remove pulse overshoots, by trimming the target pattern area in our realistic case study. This approach can obviously be inconvenient, as there is no guarantee for success, or the final target may be unsatisfactory. To minimize the risk of pulse overshoot, one might further constrict the conventional pulses, in our case within

the optimal control, e.g., setting the limit to 0.9 kHz [15]. This will likely affect the overall performance of both the conventional optimal control and the DeepControl pulses, but it also alleviates the hardware. The clipped DeepControl network solves the overshooting problem and enables more efficient use of the system.

In Ref. [21], we used ReLU layers after convolution layers. In this study, we have replaced the ReLU layers with Swish layers. This




Fig. 8. Effects on the magnetization patterns of regulating the amplitude overshoot of the non-clipped DeepControl pulse from Fig. 6. The pulse has been clipped (left) or down-scaled (right) after the non-clipped DeepControl network predicted the pulse.

trades substantially longer training for lower RMSEs. Nevertheless, for faster progress, it is possible to obtain good results with ReLU layers (data not shown). We adjusted many of the other network and deep learning parameters before switching from ReLU to Swish layers. For example, the number of filters in the convolution layers, the mini-batch size, the learn rate and momentum factor. Sweeping the vast multi-dimensional parameter space is a considerable task given the long deep learning time (with either ReLU or Swish layers). Our success criteria were satisfactory NRMSE values and visual quality inspection of the magnetization profiles from simulations, which is in line with conventional pulse optimization approaches.

We did not make any effort to find and remove outlier cases from the library. Typical outlier cases happen by our way of generating the library when a target does not overlap with sufficient B_1^+ sensitivity. Those few co-incidents (< 1%) yielded unfinished, (sub-)optimal control pulses, which the trained networks also could not support. The deep learning we used, however, is robust (overall) to the presence of some amount of outlier cases in the library. We further propose to use pulse libraries such as those generated herein for investigations on automated solution acceptance/rejection, improved error metric standards, and potentially address artificial intelligence *explainability* [31,32].

Average power did not raise concern in this or the previous study [21], but pulse optimization pipelines sometimes include average power limits, and/or even local/global SAR constraints. We hope to address these constraints in future studies, especially local SAR for multi-channel RF systems [6]. Another limitation of DeepControl, which we wish to repeat [15], is that the network and the library output side are tied to one magnetic field gradient waveform (Fig. 6). This means to change, e.g., spatial resolution or span, one needs to start over with another magnetic field gradient waveform. This may also result in some changes to the network layer dimensions. Related to this topic, Zhang et al. [33] have proposed a multi-task solution including magnetic field gradient waveform design, and Shin et al. [26] can generate different pulse types by reinforcement learning. Further, we wish to extend DeepControl to 3D RF selection in the future. For that, 3D magnetic field gradient waveforms are needed, and this yearns for more network flexibility.

5. Conclusion

Utilizing deep learning in magnetic resonance imaging, to improve image reconstruction or diagnoses, or facilitate novel control designs receives increasing interest these years. Here we have shown that a custom-made clipping layer in a deep neural network actively limits predicted controls, such that they stay within user-prescribed boundaries or hardware restrictions. This sort of control limit was in the

previous networks only enforced in the supervision library. Hence, the deep learned neural network was not guaranteed to adopt these limits. Here we have shown that the clipping layer does not interfere with the training progress or the statistical performance of the controls.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We thank VILLUM Fonden, Eva og Henry Fraenkels Mindefond, Harboefonden, and Kong Christian Den Tiendes Fond. We also thank C.S. Aigner, I. Kuprov, D.F. Hansen, I.I. Maximov, and D.L. Goodwin for fruitful discussions.

Code availability

Our repository at <https://github.com/madssakre/DeepControl3> will contain the proposed method.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.artmed.2022.102460>.

References

- [1] Brown RW, Cheng Y-CN, Haacke EM, Thompson MR, Venkatesan R. *Magnetic resonance imaging: physical principles and sequence design*. 2nd ed.. Hoboken, New Jersey: John Wiley & Sons, Inc; 2014.
- [2] Caan MWA, Bazin P, Marques JP, Hollander G, Dumoulin SO, Zwaag W. MP2RAGEME: T_1 , T_2 *, and QSM mapping in one sequence at 7 tesla. *Human Brain Mapp* 2019;40:1786–98. <http://dx.doi.org/10.1002/hbm.24490>.
- [3] Dietrich S, Aigner CS, Kolbitsch C, Mayer J, Ludwig J, Schmidt S, Schaeffter T, Schmitter S. 3D Free-breathing multichannel absolute Mapping in the human body at 7T. *Magn Reson Med* 2021;85:2552–67. <http://dx.doi.org/10.1002/mrm.28602>.
- [4] Finsterbusch J. Fast-spin-echo imaging of inner fields-of-view with 2D-selective RF excitations. *J Magn Reson Imaging* 2010;31:1530–7. <http://dx.doi.org/10.1002/jmri.22196>.
- [5] Herrler J, Liebig P, Gumbrecht R, Ritter D, Schmitter S, Maier A, Schmidt M, Uder M, Doerfler A, Nagel AM. Fast online-customized (FOCUS) parallel transmission pulses: A combination of universal pulses and individual optimization. *Magn Reson Med* 2021;85:3140–53. <http://dx.doi.org/10.1002/mrm.28643>.
- [6] Vinding MS, Guérin B, Vosegaard T, Nielsen NC. Local, SAR and global SAR and power-constrained large-flip-angle pulses with optimal control and virtual observation points. *Magn Reson Med* 2017;77:374–84. <http://dx.doi.org/10.1002/mrm.26086>.
- [7] Bernal J, Kushibar K, Asfaw DS, Valverde S, Oliver A, Martí R, Lladó X. Deep convolutional neural networks for brain image analysis on magnetic resonance imaging: a review. *Artif Intell Med* 2019;95:64–81. <http://dx.doi.org/10.1016/j.artmed.2018.08.008>.
- [8] Boutillon A, Borotikar B, Burdin V, Conze P-H. Multi-structure bone segmentation in pediatric MR images with combined regularization from shape priors and adversarial network. *Artif Intell Med* 2022;132:102364. <http://dx.doi.org/10.1016/j.artmed.2022.102364>.
- [9] Li R, Wang X, Lawler K, Garg S, Bai Q, Alty J. Applications of artificial intelligence to aid early detection of dementia: A scoping review on current capabilities and future directions. *J Biomed Inform* 2022;127:104030. <http://dx.doi.org/10.1016/j.jbi.2022.104030>.
- [10] Parimbelli E, Wilk S, Cornet R, Sniatala P, Sniatala K, Glaser S, Fraterman I, Boekhout A, Ottaviano M, Peleg M. A review of AI and data science support for cancer management. *Artif Intell Med* 2021;117:102111. <http://dx.doi.org/10.1016/j.artmed.2021.102111>.
- [11] Helmy Abdou MA, Truong TT, Dykyy A, Ferreira P, Jul E. CapillaryNet: An automated system to quantify skin capillary density and red blood cell velocity from handheld vital microscopy. *Artif Intell Med* 2022;127:102287. <http://dx.doi.org/10.1016/j.artmed.2022.102287>.

- [12] Popescu DM, Shade JK, Lai C, Aronis KN, Ouyang D, Moorthy MV, Cook NR, Lee DC, Kadish A, Albert CM, Wu KC, Maggioni M, Trayanova NA. Arrhythmic sudden death survival prediction using deep learning analysis of scarring in the heart. *Nat Cardiovasc Res* 2022;1:334–43. <http://dx.doi.org/10.1038/s44161-022-00041-9>.
- [13] Nielsen A, Hansen MB, Tietze A, Mouridsen K. Prediction of tissue outcome and assessment of treatment effect in acute ischemic stroke using deep learning. *Stroke* 2018;49:1394–401. <http://dx.doi.org/10.1161/STROKEAHA.117.019740>.
- [14] Schü LC, Mello B, da Costa CA, Antunes RS, Rigo SJ, de Oliveira Ramos G, da Rosa Righi R, Scherer JN, Donida B. A rapid review of machine learning approaches for telemedicine in the scope of covid-19. *Artif Intell Med* 2022;129:102312. <http://dx.doi.org/10.1016/j.artmed.2022.102312>.
- [15] Vinding MS, Skyum B, Sangill R, Lund TE. Ultrafast (milliseconds), multidimensional RF pulse design with deep learning. *Magn Reson Med* 2019;82:586–99. <http://dx.doi.org/10.1002/mrm.27740>.
- [16] Vinding MS, Goodwin DL, Kuprov I, Lund TE. Optimal control gradient precision trade-offs: Application to fast generation of DeepControl libraries for MRI. *J Magn Reson* 2021;333:107094. <http://dx.doi.org/10.1016/j.jmr.2021.107094>.
- [17] Maximov II, Vinding MS, Tse DH, Nielsen NC, Shah NJ. Real-time 2D spatially selective MRI experiments: Comparative analysis of optimal control design methods. *J Magn Reson* 2015;254:110–20. <http://dx.doi.org/10.1016/j.jmr.2015.03.003>.
- [18] Massire A, Cloos MA, Vignaud A, Le Bihan D, Amadon A, Boulant N. Design of non-selective refocusing pulses with phase-free rotation axis by gradient ascent pulse engineering algorithm in parallel transmission at 7T. *J Magn Reson* 2013;230:76–83. <http://dx.doi.org/10.1016/j.jmr.2013.01.005>.
- [19] Aigner CS, Clason C, Rund A, Stollberger R. Efficient high-resolution RF pulse design applied to simultaneous multi-slice excitation. *J Magn Reson* 2016;263:33–44. <http://dx.doi.org/10.1016/j.jmr.2015.11.013>.
- [20] Reeth EV, Ratiney H, Tse Ve Koon K, Tesch M, Grenier D, Beuf O, Glaser SJ, Sugny D. A simplified framework to optimize MRI contrast preparation. *Magn Reson Med* 2019;81:424–38. <http://dx.doi.org/10.1002/mrm.27417>.
- [21] 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* 2021a;85:3308–17. <http://dx.doi.org/10.1002/mrm.28667>.
- [22] Vinding MS, Aigner CS, Stockmann J, Guerin B, Schmitter S, Lund TE. DeepControl: AI-powered slice flip-angle homogenization by 2D RF pulses. In: *International Society of Magnetic Resonance in Medicine*. 2021b, p. 0785.
- [23] Deng J, Dong W, Socher R, Li L-J, Li Kai, Fei-Fei Li. ImageNet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. 2009, p. 248–55. <http://dx.doi.org/10.1109/CVPR.2009.5206848>.
- [24] Merolla P, Appuswamy R, Arthur J, Esser SK, Modha D. Deep neural networks are robust to weight binarization and other non-linear distortions. 2016, arXiv preprint [arXiv:1606.01981](https://arxiv.org/abs/1606.01981).
- [25] Kingma DP, Ba J. Adam: A Method for Stochastic Optimization. 2017, URL: <http://arxiv.org/abs/1412.6980>.
- [26] Shin D, Kim Y, Oh C, An H, Park J, Kim J, Lee J. Deep reinforcement learning-designed radiofrequency waveform in MRI. *Nat Mach Intell* 2021;3:985–94. <http://dx.doi.org/10.1038/s42256-021-00411-1>.
- [27] Goodwin DL, Kuprov I. Modified Newton-Raphson GRAPE methods for optimal control of spin systems. *J Chem Phys* 2016;144:204107. <http://dx.doi.org/10.1063/1.4949534>.
- [28] Rumelhart DE, Hinton GE, Williams RJ. Learning representations by back-propagating errors. *Nature* 1986;323:533–6. <http://dx.doi.org/10.1038/323533a0>.
- [29] Grissom W. ISMRM RF Pulse Design Challenge. 2016, URL: <https://github.com/wgrissom/ISMRM-RF-Pulse-Design-Challenge>.
- [30] Lustig M, Kim S-J, Pauly JM. A fast method for designing time-optimal gradient waveforms for arbitrary k-space trajectories. *IEEE Trans Med Imaging* 2008;27:866–73. <http://dx.doi.org/10.1109/TMI.2008.922699>.
- [31] Combi C, Amico B, Bellazzi R, Holzinger A, Moore JH, Zitnik M, Holmes JH. A manifesto on explainability for artificial intelligence in medicine. *Artif Intell Med* 2022;133:102423. <http://dx.doi.org/10.1016/j.artmed.2022.102423>.
- [32] Amey JL, Keeley J, Choudhury T, Kuprov I. Neural network interpretation using descrambler groups. *Proc Natl Acad Sci USA* 2021;118. <http://dx.doi.org/10.1073/pnas.2016917118>.
- [33] Zhang Y, Jiang K, Jiang W, Wang N, Wright AJ, Liu A, Wang J. Multi-task convolutional neural network-based design of radio frequency pulse and the accompanying gradients for magnetic resonance imaging. *NMR Biomed* 2021;34. <http://dx.doi.org/10.1002/nbm.4443>.