computers

MDPI

Review

A Review of AI-Powered Controls in the Field of Magnetic Resonance Imaging

Mads Sloth Vinding * and Torben Ellegaard Lund

Center of Functionally Integrative Neuroscience (CFIN), Aarhus University, 8000 Aarhus, Denmark; torbenelund@cfin.au.dk

* Correspondence: msv@cfin.au.dk

Abstract

Artificial intelligence (AI) is increasingly reshaping the control mechanisms that govern magnetic resonance imaging (MRI), enabling faster, safer, and more adaptive operation of the scanner's physical subsystems. This review provides a comprehensive survey of recent AI-driven advances in core control domains: radio frequency (RF) pulse design and specific absorption rate (SAR) prediction, motion-dependent modeling of B_1^+ and B_0 fields, and gradient system characterization and correction. Across these domains, deep learning models—convolutional, recurrent, generative, and temporal convolutional networks—have emerged as powerful computational surrogates for numerical electromagnetic simulations, Bloch simulations, motion tracking, and gradient impulse response modeling. These networks achieve subject-specific field or SAR predictions within seconds or milliseconds, mitigating long-standing limitations associated with inter-subject variability, non-linear system behavior, and the need for extensive calibration. We highlight methodological themes such as physics-guided training, reinforcement learning for RF pulse design, subject-specific fine-tuning, uncertainty considerations, and the integration of learned models into real-time MRI workflows. Open challenges and future directions include unified multi-physics frameworks, deep learning approaches for generalizing across anatomies and coil configurations, robust validation across vendors and field strengths, and safety-aware AI design. Overall, AI-powered control strategies are poised to become foundational components of next-generation, high-performance, and personalized MRI systems.

Keywords: artificial intelligence; control; magnetic resonance imaging; deep learning; neural networks; optimization

Check for updates

Academic Editor: Paolo Bellavista

Received: 30 November 2025
Revised: 22 December 2025
Accepted: 22 December 2025
Published: 5 January 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

This review examines the role of artificial intelligence (AI) in control for magnetic resonance imaging (MRI). Comprehensive surveys of AI-driven control already exist in other domains, such as control theory [1], robotics [2], industrial automation [3], and aeronautics [4]. Conceptual advances and milestones from these fields can inspire and inform control strategies specific to MRI. Here, we focus on tasks related to operating and controlling the MRI system. As an imaging modality, MRI is the most versatile method compared to ultrasound (US), X-ray, computed tomography (CT), positron emission tomography (PET), etc. There are numerous MRI contrasts to choose from, e.g., longitudinal relaxation-time weighting (T_1w), transverse relaxation-time weighting (T_2w, T_2^*w), and proton-density weighting (ρ_0w) [5]; diffusion-weighted imaging (DWI) [6]; diffusion tensor imaging (DTI) [7]; perfusion-weighted imaging (PWI) [8]; blood-oxygenation-

level-dependent (BOLD) functional MRI (fMRI) [9]; MR elastography [10]; time-of-flight and phase-contrast MR angiography (TOF-/PC-MRA) [11]; MR spectroscopy (MRS) [12]; susceptibility-weighted imaging (SWI) [13]; temperature mapping [14], among others. This diversity enables MRI to support an exceptionally wide clinical range, including the following:

- Neurological imaging (neurodegenerative disease, stroke, aneurysms, brain tumors);
- Cardiovascular and thoracic imaging (cardiomyopathies, ischemia, aortic disease);
- Spinal and musculoskeletal imaging (intervertebral disc pathology, spinal cord compression, bone and soft-tissue tumors);
- Abdominal imaging (liver and kidney disease, malignancies, tumors).

Together, these contrasts and application domains define a large space of technical tasks on both the controlling and post-processing sides of MRI. In this review, we define post-processing as all operations that occur once raw imaging k-space (spatial frequency) data have been stored and the acquisition is complete. This includes image reconstruction (e.g., Fourier transforms and iterative reconstruction), noise reduction, artifact mitigation, quantitative parameter-map generation, image interpretation, and downstream diagnostic tasks. Numerous studies on AI applied to post-processing tasks are left for the reader elsewhere.

The acquisition-control side, which this review primarily focuses on, concerns steering the MRI system to acquire the desired contrast in the area of interest while optimizing precision and specificity. An MRI acquisition is executed by a pulse sequence, a microsecond-resolved schedule governing the available controls and their amplitudes. Some controlling and post-processing tasks can be treated as independent modules (or cascades of modules), but there can also be dependencies between controlling and post-processing tasks. Many modern acquisitions adapt sequence parameters based on quantitative maps or tissue estimates obtained from preparatory scans, which again may be supported by AI methods (in a post-processing fashion). Examples of these will also be covered herein.

2. Fundamentals of MRI Controls

At its core, MRI relies on the interaction between nuclear spins and externally applied magnetic fields. When a sample is placed in a (preferably strong) static magnetic field (commonly denoted B_0), certain nuclei with spin, e.g., ^{1}H nuclei, partially align with the field, creating a net magnetization. This magnetization can be perturbed in a desirable fashion by applying carefully timed and shaped radio frequency (RF) pulses with a carrier frequency matching the Larmor frequency specific to the nuclei and B_0. The subsequent evolution of the magnetization—governed by the Bloch equations—produces an RF signal whose characteristics reflect both the microscopic chemical environment and macroscopic tissue properties.

2.1. Gradients and RF

In MRI, spatial encoding in x, y, and z is achieved using magnetic field gradient pulses (G_x, G_y, G_z), which introduce approximately linear and time-varying shifts in the Larmor frequency across the object. The wide variety of MRI contrasts arises from manipulating RF and gradient waveforms according to specific pulse-sequence timing schemes, thereby encoding frequency, phase, and relaxation dynamics into the acquired k-space data. In this sense, gradient and RF waveforms constitute the primary controllable degrees of freedom in MRI.

RF transmission (Tx) is performed either using the whole-body coil integrated into the magnet bore or via local transmit coils. RF reception is typically performed using multi-channel local coils positioned near the anatomy of interest. Commercial scanners commonly use the body coil for transmission and a multi-channel head, spine, or torso coil

for reception. Parallel receive arrays—often comprising 32, 64, or more channels—enable accelerated imaging through parallel imaging techniques such as SENSE [15] or GRAPPA [16]. A similar concept for transmission has become increasingly popular in recent years, and parallel transmit systems are RF transmit arrays with multiple independently driven channels (e.g., 8, 16, or 24) [17,18]. The extra degrees of freedom can be used to mitigate RF inhomogeneity and improve excitation fidelity, and there is similarly room for RF pulse acceleration [19]. However, each extra RF channel expands the dimensionality of the RF control problem and significantly increases the computational calibration burden [20,21]. Hence, this is one challenge that receives attention with AI as it seems to facilitate orders of magnitude computation time speed-up [22,23]. An acceleration not yet achievable with conventional, non-AI computation methods.

The gradient coils are permanently installed in the magnet housing. They can be upgraded or replaced, but otherwise they largely remain fixed. Gradient-insert systems have been proposed to provide substantially higher gradient amplitudes and slew rates within a reduced field of view, enabling improved spatial and temporal resolution and potentially reducing whole-body peripheral nerve stimulation (PNS) or acoustic noise [24,25]. Inserts require dedicated installation and are not modular in the same sense as RF coils. From a control perspective, however, they are worked similarly to standard gradient channels.

2.2. The B_0 Superconductor and Shims

The main magnetic field B_0, generated by a superconducting solenoid, is static aside from slow temporal drift. The homogeneity of the B_0 field is optimized during installation with passive shim inserts in the magnet housing. Although the B_0 field is not typically treated as a controllable parameter during routine scanning, it is further optimized by active, subject-dependent shimming at the beginning of each session using spherical-harmonic shim coils up to third order. These shim coils are typically located in the gradient coil housing, which is again integrated out of sight within the magnet housing. Shim coils are, in principle, also a set of controls, and can be tailored to shim a voxel instead of an entire field-of-view. It is common, however, to optimize the active shims and leave them fixed for the entire session. More recently, local multi-channel shim arrays placed near the subject have been introduced to achieve higher-fidelity static field correction [26,27]. The currents in these shim channels can be regarded as additional controls, analogous to RF and gradient channels [28]. Although challenging the control optimization stage, advantages of this extra set of controls include additional degrees of freedom when perturbing the magnetization, in turn yielding an improved signal response [28]. But also that local multi-channel shim arrays are very cost-effective supplements, e.g., due to more budget-friendly current amplifiers.

2.3. Motion

Subject motion (subject movement, respiration, heartbeat, pulsating arteries, swallowing, eye movement, eye blinking) can be a source of image artifacts. There are various means to monitor motion to correct the acquisition during the scan or in post-processing, and AI may also play a role in achieving a successful scan. Motion can be monitored using external devices [29] or estimated from navigators and other intra-scan measurements [30]. However, motion is commonly minimized by oral instruction to remain still as much as possible and by physical stabilization, e.g., with pads around the head. Also straightened tape connecting nose and/ or forehead with the coil helmet is a practical haptic feedback for the subject [31,32]. A comfortable resting position is also assured to minimize the urge to move. In various scans, the subjects are instructed to hold their breath for a short while. Motion and the control of motion are only briefly touched upon in this review, but Nayak

et al. [33] have reviewed the concept of real-time MRI (RT-MRI), where the scans adjust to irregular motion.

2.4. Subject Interaction

In fMRI, subjects typically interact with external stimuli and may provide behavioral responses via button presses or eye-tracking. These responses can modify stimulus timing, introduce pauses, or trigger the onset of acquisitions [34]. To the best of the authors' knowledge, subject inputs are limited (for safety reasons) to control the course of events during the fMRI session, but not low-level scanner settings. However, RT-MRI [33] can be considered as subject-dependent control influence, and indeed, it can be imagined that AI would be useful in such a case.

3. Why We Need to Optimize Controls

Despite their central role in MRI, the available system controls—and the environment in which they operate—are far from ideal.

3.1. Gradient Controls

Gradient fields (G_x, G_y, and G_z), which define transmit and imaging k-space trajectories, exhibit spatial nonlinearities, particularly toward the periphery of the field of view, leading to geometric distortions and warping. Rapid gradient switching induces eddy currents in surrounding conductive structures, producing secondary magnetic fields that oppose the intended gradients. These manifest as temporal delays, residual offsets, and spatially varying phase errors. Hardware limitations, such as finite amplifier bandwidth and coil inductance, further cause actual gradient waveforms to lag behind or deviate from their commanded shapes. Given a nominal gradient waveform, there are pulse sequences designed to measure the actual gradient waveform or the trajectory it takes in the k-space [35,36]. This allows the user to apply RF control design to the measured gradient waveform, but the tedious preparation procedure must be performed for every nominal gradient waveform one intends to use. Alternatively, the gradient waveform or k-space trajectory can be measured with a set of field probes or "field cameras" [37]. They further enable immediate image reconstruction using the measured gradient waveform, so that at least the post-processing side benefits—RF pulse design based on a measured gradient waveform still requires a preparatory step in advance. A third option is to characterize the gradient system, e.g., with a field camera, and model the gradient system as a linear time-invariant (LTI) system [38]. The resulting gradient impulse response function (GIRF) can then be directly employed in RF pulse design without further preparation [39]. It is also common to use the characteristic profile of the system for pre-emphasis to prospectively account for gradient waveform distortions. Pre-emphasis deliberately modifies the commanded gradient current so that, after the physical system (coils, amplifiers, eddy currents) acts on it, the actual gradient waveform closely matches the desired one. Gradient inserts have a smaller region with linear encoding—requiring correction somewhere in the pipeline—and may need attention to concomitant fields too [24].

3.2. Gradient Effects

In addition, spatial and temporal imperfections, which to some extent can be accounted for by the (non-exhaustive) examples mentioned above or other means, are also the primary source of acoustic noise in MRI. Specific mechanical eigenmodes of the gradient assembly must be avoided to prevent excessive vibration, image artifacts, or automatic safety shutdowns. Fast gradient switching can also induce PNS, motivating strict constraints on gradient amplitudes and slew rates. Although aggressive gradient performance enables more rapid k-space traversal and thus higher spatial and temporal resolution,

international safety standards impose limits on dB/dt to prevent PNS [40]. High-duty-cycle gradient activity also causes heating of gradient coils and amplifiers, which can exacerbate waveform distortions and reduce fidelity over the course of a long scan. These limitations may become an active part of the control design to avoid certain system checks that block the pulse sequence, or, e.g., if subjects complain about PNS.

3.3. RF Controls and the B_1^+ Field

RF controls are likewise imperfect. Coil coupling, transmit chain nonlinearities, timing jitter, and hardware drift can degrade RF pulse performance. However, one dominant challenge today concerns the spatial distribution the B_1^+ field, i.e., the magnetic-field component of the RF electromagnetic transmission. The signal-to-noise ratio (SNR) scales roughly linearly with B_0, and because MRI is inherently a low-SNR imaging modality, there is strong motivation to pursue higher magnetic field strengths. Higher SNR can be traded for higher spatial and temporal resolution, and sensitivity to detect subtle contrasts. Higher fields also improve spectral resolution, which is particularly beneficial in MRS. Yet, as the field strength increases, the RF wavelength in tissue becomes shorter relative to the dimensions of the head or torso, producing substantial B_1^+ field inhomogeneity. Uncorrected field variation results in spatially non-uniform magnetization flip angles (FAs), which appear as bright and dark regions in images and ultimately compromise contrast. Importantly, B_1^+ field inhomogeneity is highly subject-dependent. This means a B_1^+ field sensitivity map that would hold the information needed to compensate spatially non-uniform FA ideally should be measured as a prerequisite for RF control optimization for each subject. However, as this prerequisite is a high demand—especially for parallel transmit systems—different methods for B_1^+ field inhomogeneity compensation exist, namely:

1. Subject-tailored pulses, which offer the highest performance but require B_1^+ field mapping and pulse optimization for each individual [41].
2. Robust universal pulses (UPs), pre-optimized across large ensembles of measured B_1^+ field maps to provide good average performance without per-subject optimization [42].
3. Hybrid strategies combining UPs with subject-tailoring [43].
4. Pulse sequences with inherent dedication to B_1^+ inhomogeneity mitigation in their parameter setup [44,45].
5. Pulse sequences applying B_1^+ inhomogeneity correction in the post-processing stage [45,46].

For severe B_1^+ field inhomogeneity, especially at ultrahigh field, the extra degrees of freedom of parallel transmit systems come in handy. But parallel transmit places additional strain on subject-tailored pulse preparation due to time-consuming multi-channel field mapping and increased online computation time. This favors general-purpose UPs. However, the optimal solution is likely a combination of the methods in the list above [47].

3.4. RF Energy Effects

RF transmission also deposits energy into the tissue, causing heating. This adheres to the electric (E) field component of the RF electromagnetic transmission. Currently, the specific absorption rate (SAR), which measures induced power per unit mass of tissue, is a proxy for temperature elevation constraints. Regulatory limits on SAR constrain the average and local RF power that may be delivered [40]. Especially with parallel transmit, local SAR constraints are challenging to tackle [21,48]. To make RF pulse optimization tractable, SAR is typically estimated using numerical whole-body models, e.g., the Virtual Family (VF) [49], and the resulting millions of individual local SAR constraints are compressed into manageable sets via virtual observation points (VOPs) [50]. SAR constraints on top of amplitude and average power constraints, etc., remain a computational burden that

would benefit from new and improved tools. However, because SAR is only a proxy for heat, there is a gap between the relevant biological parameter, tissue temperature, and our controls, leading to heating. Furthermore, there is also a gap to fill with respect to SAR being modeled versus having subject-dependent measurements to support more precise SAR or temperature estimates.

3.5. B_0 Field

In addition to B_1^+ field inhomogeneity, the main magnetic field B_0 also exhibits subject-dependent inhomogeneity. B_0 field maps are straightforward to measure and are commonly incorporated into RF pulse optimization. Off-resonance effects from B_0 field inhomogeneity can produce FA errors, slice misregistration, and blurred excitation profiles. Long RF pulses are particularly sensitive because off-resonance spins accumulate additional phase during the pulse, thereby reducing excitation fidelity.

3.6. Summary

These many control challenges in MRI have led to a rich and growing body of literature that addresses one or several issues through control optimization. It is beyond the scope of this article to review the methods herein, but we will focus on some of the latest novel AI-based methods.

4. Early Demonstrations of AI in Control in Magnetic Resonance

To the best of the authors' knowledge, the earliest attempt to apply neural networks (NNs) in magnetic resonance for control design was by Gezelter and Freeman in 1990 [51]. Their work was demonstrated for a nuclear magnetic resonance (NMR) spectroscopy application, which, in broad terms, is not MRI but merits mentioning here because of the frequent methodological overlap between the fields and because the authors explicitly discuss potential MRI-related extensions, including concepts resembling those described later in Section 5.4. Gezelter and Freeman considered an RF-pulse design problem targeting a $+90^\circ$ FA in a finite frequency band on the negative frequency axis and a -90° FA in the mirrored positive-frequency band. This target frequency-selective magnetization pattern can be produced experimentally on an NMR spectrometer using the shaped JANUS RF pulse. The authors constructed a NN to predict 16 Fourier coefficients that parametrize the JANUS pulse. The input ("retina") layer consisted of 100 neurons divided into two 50-neuron segments that received the absorption-mode and dispersion-mode excitation profiles in the frequency domain, respectively. Hence, the target magnetization was represented by 50 frequency bins across the spectral width for each of the transverse magnetization components M_x and M_y. The second (hidden) layer contained 200 fully connected (FC) neurons with sigmoid activation functions. The output layer comprised 16 neurons corresponding to the 16 Fourier coefficients of the JANUS pulse representation. Gezelter and Freeman trained the NN using an iterative loop in which each cycle began by generating a random set of 16 Fourier coefficients within prescribed limits. The Bloch equations were then solved to compute the frequency-domain excitation pattern produced by this randomly generated pulse. This excitation pattern served as the NN input. The NN's forward pass provided predicted Fourier coefficients, and these coefficients were used in the training procedure through standard back-propagation. The authors employed an exponentially decaying learning rate, and convergence was typically achieved after approximately 28,000 training samples. During inference, the NN was able to predict pulse shapes "*closely related to those obtained in earlier simulated annealing searches*" [51], demonstrating that neural networks could emulate a computationally expensive optimal-control-like search for RF pulse design.

5. MRI Control Designs Powered by AI

The following sections present an overview of AI-powered control designs, categorized by application and type, although some overlap occurs. Table 1 below gives a chronological overview of the presented AI-powered control designs.

Table 1. Chronological order (by year) of AI-powered control designs for MRI. Abbreviations are described in the text and summarized in the Abbreviations section.

Year	Authors	Employment	AI	RF [num. chan.]	$G_{x,y,z}$	B_0[T]	Experiments
1990	Gezelter & Freeman [51]	NMR dual band exc.	SvL, FC-NN	1	-	9.4	NMR, in vitro
2018	Mirfin et al. [52]	RF shim.	SvL, FC-NN	8	5-spoke	7	Head, in silico
	Ianni et al. [53]	RF shim.	ML, PIPRR	24	-	7	Head, in silico
2019	Vinding et al. [22,23]	2D spat.-select. exc./inv.	SvL, FC-NN	1	No *	3	Head, in vivo
	Tomi-Tricot et al. [54]	k_T-points	ML classifier	2 **	$5k_T$-points **	3	Abdom., in vivo
2020	Shin et al. [55]	SLR root-flip.	RL, FC-NN	1	-	3	Phantom
2021	Zhang et al. [56]	2D spat.- select. exc.	SvL, CNN	1	Yes, $G_{x,y}$	3	Phantom
	Vinding et al. [57]	2D spat.-select. exc.	SvL, CNN	1	No *	7	Head, in vivo
	Vinding et al. [58]	Dyn. RF shim.	SvL, CNN	1	No *	7	Head, in silico
	Shin et al. [59]	SLR exc.; SLR inv.; B_1^+-ins. adia. select. inv.; B_1^+-ins. adia. unselect. inv.	RL, RNN	1	No	3	Phantom
2022	Herrler et al. [60]	FOCUS	SvL, U-Net	8 **	2-spoke **	7	Head, in silico
	Eberhardt et al. [61]	RF shim.	SvL, ResNet, VAE	16	2-spoke	9.4	Head, in vivo
2023	Vinding et al. [62]	2D spat.-select. exc.	SvL, CNN	1	No *	7	Head, In silico
2024	Jang et al. [63]	SLR exc; B_1^+-ins. adia. select. inv.; spect.-spat. exc.; 2D spat.-select. exc.	PhG SSvL, FC-NN	1	No *	3	Knee, head, in vivo
	Kilic et al. [64]	Stat. RF shim.	PhG UsvL, CNN	8	-	7	Head, in silico
2025	Nagelstrasser et al. [65]	Dyn. RF shim.	SvL, CNN	8	Yes, $G_{x,y,z}$	7	Head, in vivo
	Lu et al. [66]	RF shim.	ResNet, Adam, NFD	8	-	7	Head, in silico

* Gradient waveforms $G_{x,y}$ pre-computed. ** The AI predicts which UP to use, not the controls themselves.

5.1. RF Shimming and Spoke Pulses

Mirfin et al. [52] presented an FC NN capable of predicting 8-channel parallel-transmit, 5-spoke RF shimming weights—including the transmit k-space locations—for slice inversion. The NN took as input full 2D, 8-channel B_1^+ field maps and, optionally, a B_0 field map. The output comprised 5×8 RF amplitudes, 5×8 RF phases, and 5 k-space spoke locations (k_x, k_y). The NN was trained in a supervised learning (SvL) manner using 60,000 samples (70% training, 30% validation) and evaluated using Bloch-simulation-based numerical experiments. The reported results showed "*inferior inversion performance as compared to joint-optimization of spokes pulses using a full optimal control*" [52], which the authors attributed to insufficient diversity in the training data. The work of Mirfin et al. was albeit a demonstration of "labeling" occurring for the SvL with a type of control optimizer [67]. Although this can be a lengthy procedure similar to manual annotation in image segmentation [68], the control optimization can be performed offline.

Ianni et al. [53] presented a machine learning (ML) method, RF Shim Prediction by Iteratively Projected Ridge Regression (PIPRR), designed to minimize dependence on subject-specific B_1^+ field maps and reduce computation time for slice-by-slice subject-tailored RF shimming. This is particularly valuable for many-channel systems, such as

24-channel parallel transmit at 7 T, where full B_1^+ field mapping and tailored RF shimming become time-consuming. Notably, instead of using complete 24-channel B_1^+ field maps as inputs, the PIPRR method predicts RF shim weights from compact subject features such as mask centroids, DC Fourier coefficients of the B_1^+ field maps, Fourier shape descriptors, and slice-position information. The PIPRR predictions exhibited slightly higher Coefficient of Variation (CoV) in the resulting B_1^+ patterns than the conventional tailored optimization, but SAR values were comparable or lower. Because PIPRR predictions were computed within a few milliseconds—and because full B_1^+ field mapping could be bypassed—significant practical acceleration was achieved. The PIPRR method was trained on 100 numerical head simulations and demonstrated with numerical simulations.

Building on an AI-based B_1^+ field mapping framework (see Section 6.2), Eberhardt et al. [61] proposed AI-driven prediction of 16-channel, 2-spoke RF shimming weights (including transmit k-space locations of the second spoke). The 16-channel B_1^+ field maps constituting the input were organized with the real, imaginary, and magnitude parts as the three color channels of portable network graphics images. They were fed into a Residual Network (ResNet) convolutional neural network (CNN) [69]. Instead of letting the CNN predict 32 real and 32 imaginary RF shimming weights (16 per spoke), plus (k_x, k_y) for the second spoke, in total 66 parameters, they appended a variational autoencoder (VAE) [70,71]. The VAE decoded a 16-dimensional latent representation into the 66-parameter spoke-pulse description. The CNN was trained by SvL to output this latent vector. The hybrid CNN–VAE architecture slightly outperformed a plain ResNet that directly regressed all 66 parameters, but otherwise predicted pulses with performance comparable to conventional methods.

Kilic et al. [64] introduced a CNN to predict 2D, 8-channel parallel transmit small-FA RF shimming weights based on B_1^+ field maps, similar in spirit to Mirfin et al. [52]. However, Kilic et al. concatenated the eight channels of the 7 T B_1^+ field maps (each of size 39×44) along a spatial dimension, using real and imaginary parts as two channels, effectively forming a $312 \times 44 \times 2$ tensor. This allowed a fully shift-equivariant CNN architecture. To avoid the need for large SvL libraries, they pursued unsupervised learning (UsvL) using a physics-guided (PhG) loss function in which each candidate shim was evaluated by a Bloch-simulation module. After the input layer, the CNN consisted of two convolutional layers, each followed by a max-pooling layer, and three FC layers, producing eight real and eight imaginary RF shim weights. No explicit amplitude, power, or SAR constraints were included in the network or loss function; thus, its performance was compared to the unregularized magnitude-least-squares (MLS) approach [72]. Despite removing the dependence on MLS reference samples in the training library, the authors discuss the dependence on many input samples. Reasonable results were achieved with 86 B_1^+ field maps (three subjects), while performance comparable to the full training database (3059 B_1^+ field maps, 115 subjects) required 514 B_1^+ field maps (21 subjects). Notably, the CNN occasionally outperformed the MLS baseline: for example, it avoided excitation-profile "holes" that sometimes appeared in MLS solutions. The method was demonstrated in numerical simulations.

Lu et al. [66] recently proposed Fast-RF-Shimming, an approach reported to be roughly 5000 times faster than conventional MLS optimization [72]. Training data were generated from numerically simulated 7 T, 8-channel B_1^+ field maps using a standard body model in Ansys HFSS (Canonsburg, PA, USA). Data augmentation, e.g., via rotation, yielded 24,576 2D B_1^+ field maps. Reference shim weights were computed using the Adam optimizer [73] with 300 random initializations per target to ensure high-quality solutions. Based on these inputs and reference data, they trained a ResNet [69] to predict RF shimming weights. Furthermore, they developed a Non-uniformity Field Detector (NFD) that predicts

whether the ResNet result yields a uniform shimming result. This optional NFD test was introduced to provide an additional quality metric beyond the root-mean-square error. The method was demonstrated numerically, with inference times of ∼0.14 s per subject versus ∼13 min for MLS.

5.2. Universal Pulses

Tomi-Tricot et al. [54] introduced SmartPulse, a method that merges the advantages of offline-computed, calibration-free UPs with elements of subject specificity. Rather than refining a UP through online optimization (as discussed in the next paragraph), SmartPulse predicts which UP from a precomputed library should be used for a given subject. This is achieved using an ML classifier that operates on 10 subject features. Three features are derived from the subject's anatomy obtained from the routinely acquired localizer scan: abdominal height, abdominal width, and the abdominal height-to-width ratio. The remaining seven features—age, sex, body weight, body height, body mass index, reference voltage, and a global SAR metric—are readily available at no additional cost in scan time. SmartPulse was evaluated in vivo at 3 T for a 2-channel transmit, 5-subpulses k_T-points pulse designed for an 11° FA excitation. SmartPulse outperformed subject-tailored static RF shimming, vendor-provided pulses, and conventional universal k_T-points pulses. However, fully subject-tailored k_T-points pulses still yielded the best FA homogeneity, although the visual differences in the resulting images were often subtle. An additional practical advantage was that subject-tailored static RF-shimming occasionally failed to converge for some subjects, whereas SmartPulse performed robustly across all cases.

Fast online-customized (FOCUS) parallel transmit pulse design was proposed earlier by Herrler et al. [43]. In essence, it uses a UP as an educated initial seed for a fast online, subject-tailored optimization—one that would be faster and more robust than starting from an uneducated initial seed. Herrler et al. [60] then proposed an NN-supported FOCUS method for 2-spoke, slice selective, large-FA excitation. Multiple samples of B_0 and B_1^+ field maps were used to form three clusters, including cluster-specific UPs. Clusters-specific U-Nets [74] were trained and used to predict the FA response to a sample of B_0 and B_1^+ field maps (input). The FA response was acquired from 2-spoke FOCUS pulses that used the cluster-specific UPs as an initial seed. The individual U-Net training sessions involved B_0 and B_1^+ field maps from 132 subjects. In operation, a subject-specific set of B_0 and B_1^+ field maps would be input to the U-Nets from the three clusters. The best FA response would determine which cluster-specific UP to use in the FOCUS workflow. The NN-supported FOCUS pulses provided significantly lower FA errors than coordinate-based UP initializations and circular polarization mode pulses.

Nagelstrasser et al. [65] recently presented 3D dynamic 8-channel transmit pulses for the turbo spin echo sequence. The 3D CNN loads 3D B_0 and complex B_1^+ field maps cast as input of size $17 \times 43 \times 36 \times 27$ (channels, spatial, spatial, spatial) and outputs for 120 time bins, three gradient waveforms plus eight RF waveforms (real + imaginary), totaling 19×120 individual controls. A total of 5500 input data samples were obtained from augmenting data from 132 scanned subjects [43]. For SvL, subject-tailored target pulses were computed for each input [75]. UPs of low and high SAR levels were computed for comparison. On FA CoV, the CNN-predicted pulses fell between the low- and high-SAR UPs, with the latter achieving the best performance. However, the second place was also obtained considering the specific energy dose (SED). In turn, the CNN-predicted pulses in turn yielded a higher CoV (∼14%) than the best high-SAR UP (∼12%), but the high-SAR UP on the other hand had a ∼40% higher SED than the CNN-predicted pulses.

5.3. Slice Selective, Non-Selective, and Adiabatic RF

Shin et al. [55] introduced DeepRF$_{SLR}$, an AI-framework aimed at reducing pulse duration or peak RF amplitude for multiband refocusing pulses. They reformulated the Shinnar–Le Roux (SLR) [76] root-flipping problem [77] as a reinforcement learning (RL) task, combining a deep FC-NN agent with a greedy tree search to efficiently find optimal root patterns. At large, DeepRF$_{SLR}$ substantially outperformed conventional Monte Carlo and exhaustive search approaches in both computational speed and optimization quality. Simulation studies and phantom MRI experiments showed that DeepRF$_{SLR}$ produced slice profiles comparable to standard minimum-phase SLR pulses while achieving approximately a two-fold reduction in pulse duration.

Later, Shin et al. [59] introduced DeepRF, a multi-purpose, single-channel transmit RL framework for autonomous RF pulse generation and refinement. In this setup, a recurrent neural network (RNN) agent interacted with a virtual MRI environment and generated millions (38,400,000) of candidate RF pulses, receiving reward signals based on metrics such as slice-profile fidelity and energy efficiency. The best-performing 256 candidate RF pulses were subsequently refined via a gradient-ascent module, thereby combining RL-based exploration with numerical optimization. Using this approach, DeepRF successfully designed four key types of MRI RF pulses: (1) slice-selective excitation, (2) slice-selective inversion, (3) B_1^+-insensitive volume inversion, and (4) B_1^+-insensitive selective inversion. DeepRF slice-selective excitation and inversion pulses required 17% and 11% less energy than the SLR counterparts, respectively. DeepRF B_1^+-insensitive volume and selective inversion pulses requested 9% and 2%, and less energy than the adiabatic hyperbolic secant counterparts. Both simulations and phantom experiments at 3 T demonstrated that DeepRF not only achieved high reproducibility but also uncovered new, non-adiabatic magnetization mechanisms beyond conventional methods.

Jang et al. [63] also presented a multi-purpose, single-channel transmit method: generalized RF pulse design using physics-guided self-supervised learning (GPS). GPS includes both offline and online stages. In the offline stage, a NN (four FC layers and two leaky-ReLU activation layers) is trained in a self-supervised learning (SSvL) manner using Bloch simulations (neglecting T_1 and T_2 relaxation) for a PhG loss estimation. The network takes a vectorized target magnetization profile and outputs an RF pulse with up to 2560 time-domain samples. Because training is target-specific, it requires no large library of examples and is therefore highly data-efficient. For online adaptation, the pre-trained network is compressed using low-rank adaptation techniques [78], enabling rapid refinement to compensate for measured or estimated subject-specific B_0 and B_1^+ inhomogeneities. The authors demonstrated GPS for: (1) SLR-type slice-selective excitation; (2) Adiabatic hyperbolic secant type B_1^+-insensitive selective inversion; (3) 2D spectral-spatial excitation; (4) Spiral-type 2D spatial-selective excitation. Experiments included both phantom and in vivo studies (knee and brain). Offline training required from a few minutes to a few hours, while online refinement took under five minutes. Consistent with findings from Shin et al. [59], the GPS framework also revealed magnetization dynamics distinct from those produced by classical RF pulse design approaches.

Both GPS [63] and DeepRF [59] demonstrated multi-purpose AI-based RF design capabilities, with GPS additionally covering 2D spatial-selective excitation. The next subsection considers AI-based 2D spatial-selective RF pulse design in more detail.

5.4. 2D Spatial-Selective RF Pulses

Vinding et al. [22,23] presented DeepControl, whose first version used an FC NN that took as input a 2D target magnetization profile (64×64) and predicted a single-channel, 2D spatial-selective RF pulse designed to reproduce the target. The NN comprised three

FC layers with ReLU activations and a regression output layer of size 1278, corresponding to 639 complex-valued time samples (real and imaginary components). DeepControl was trained in SvL fashion using optimal control-based pulses [21,79], where hard amplitude constraints had been enforced during pulse computation. A fixed, precomputed gradient waveform performing an inward spiral in transmit k-space was used for all designs. Training libraries were constructed from arbitrary target patterns (binary masks and grayscale images processed from ImageNet photographs [80]). Up to 20,000 optimal control pulses were generated for each FA class—30° and 90° excitation, and 180° inversion. As few as 1000 examples were sufficient for the network to learn to mimic the optimal control solutions. For a limited FA range, reliable FA scaling was achieved simply by scaling the input, consistent with the well-known proportionality between FA and RF amplitude for small-FA RF pulses. Preliminary work also showed that adding a B_0 map as a second input channel was feasible. The first version of DeepControl was demonstrated in phantom and in vivo experiments at 3 T. Inference times were approximately 7 ms, providing speedup of three orders of magnitude compared with the optimal control pulse design method.

Zhang et al. [56] proposed a multi-task CNN that jointly designs RF and gradient waveforms for echo-planar and SLR-type 2D spatial-selective excitation. A training library of 9000 samples with varying pulse configurations was generated. Inputs were random rectangular excitation targets sampled on $101 \times 101 \times 2$ grids (real and imaginary target components). Several kernel configurations and output strategies were examined for simultaneously predicting RF amplitudes and the "fast" and "slow" echo-planar gradient waveforms. Although the multi-task method did not employ subject-dependent field maps, the CNN reliably predicted fast- and slow-echo planar gradient waveforms, as well as SLR-type RF pulses matched to those trajectories. The performance was comparable to that of the conventional method, as demonstrated in phantom experiments at 3 T. Because SLR-based 2D pulses with echo-planar gradients can be computed comparatively quickly, the standard design method was faster than the proposed CNN in this case.

Around the same time as Ref. [56] appeared, Vinding et al. [57] proposed the second version of DeepControl. This time also adopting CNNs as in Ref. [56], but to accommodate B_0 and B_1^+ field maps. The FC-NN version in earlier work had difficulty handling B_0 variations [22,23], whereas CNNs proved far more stable with an expanded input representation of size $64 \times 64 \times 3$ (target, B_0, and B_1^+ field maps in the third dimension). For the training library, the input samples were constructed from ImageNet photographs [80]. The images were reconfigured to numerically reflect realistic MRI conditions, but the image contents were otherwise intentionally completely arbitrary and unrealistic from an MRI perspective. This strategy was chosen to enforce random variation into the data as opposed to using realistic, true MRI data as seen in other studies. It further served the purpose of being a rapid, inexpensive, and abundant source of data, which realistic, true MRI data is quite the opposite of. To eliminate risks of insufficient training data, a library of 100,000 samples were computed with a faster optimal control method [81] than that used earlier. Phantom and in-vivo experiments at 7 T demonstrated effective compensation for B_0 and B_1^+ inhomogeneities. The inference time was around 10 ms, still three orders of magnitude faster than the conventional optimal control method [81].

Two-dimensional spatial-selective RF pulses are often used to control magnetization in a region smaller than the object, e.g., for reduced field of view imaging [82], and RF shimming is typically aimed at a region resembling the whole subject (either in slices, e.g., with spokes pulses, or the entire volume, e.g., with k_T-points pulses). Nevertheless, spiral-based 2D spatial-selective RF pulses, including those produced by DeepControl, can in principle be used for RF shimming by setting the target pattern equal to the subject outline. Vinding et al. [58] tested the second version of DeepControl [57] with target

shapes resembling the whole brain. The goal was to determine whether the CNN trained exclusively on small, randomly shaped, single-region targets could generalize to whole-brain targets. This was indeed possible, and the network also successfully accommodated targets consisting of multiple disjoint regions (e.g., structures in the lower brain).

A recurring challenge in DeepControl was enforcing RF amplitude limits. The optimal control pulses in the training sets obeyed hard amplitude constraints, but the networks were not guaranteed to preserve them, occasionally producing pulses with amplitudes that exceeded specification. For example, in the second version of DeepControl [57] (target FA 30°), overshoots occurred in about 2% of the test cases. The problem became more pronounced in the RF-shimming application [58] because larger target regions increased pulse energy requirements. To address this issue, Vinding et al. [62] introduced a clipping layer immediately before the regression output to strictly enforce amplitude limits. In the study for the third version of DeepControl, the target FA was 90°, which increased the overshoot risk to 18% without clipping. However, the clipping layer prevented overshoots entirely, without significant or visible changes in performance.

6. AI-Powered Support for MRI Control Design

In the following sections, we provide an overview of AI-powered methods that support MRI control design, either by supplying input data (e.g., field maps or anatomical priors) or by performing quality or safety assessment of the designed controls. Table 2 below gives a chronological overview of the presented AI-powered support methods.

Table 2. Chronological order (by year) of AI-powered support for control design in MRI. Abbreviations are described in the text and summarized in the Abbreviations section.

Year	Authors	Employment	AI	B_0 [T]	Experiments
2002	Hwang et al. [83]	G_x, G_y, G_z, pre-emphasis	back-prop. NN	3	Unknown
2020	Meliadò et al. [84]	SAR	U-Net, PatchGAN	7	Pelvis, in vivo
	Haskell et al. [85]	B_0	ResNet	3	Head, in vivo
2021	Abbasi-Rad et al. [86]	1 Tx B_1^+, SAR, adiabatic RF	U-Net	7	Head, in vivo
	Gokyar et al. [87]	SAR	U-Net	3 + 7	Body + head, in silico
	Plumley et al. [88]	B_1^+, motion	U-Net, cGAN	7	Head, in silico
2022	Eberhardt et al. [61]	16 Tx B_1^+, 2-spoke RF	ResNet, VAE	9.4	Head, in vivo
	Brink et al. [89]	SAR, tissue segmentation	ForkNet	7	Head, in vivo
	Brink et al. [90]	SAR, electromagn. params.	ForkNet	7	Head, in vivo
	Liu et al. [91]	G_x, G_y, G_z	LSTM	3	Head, in vivo
2023	Krüger et al. [92]	8 Tx B_1^+	U-Net	7	Thorax, in vivo
2024	Motyka et al. [93]	B_0, motion	U-Net	7	Head, in vivo
2025	Krüger et al. [94]	8 Tx B_1^+	Complex CNN	7	Head, in vivo
	Dan et al. [95]	G_x, G_y, G_z, pre-emphasis	LSTM	3	Head, in vivo
2026	Martin et al. [96]	G_x, G_y, G_z	TCN	7	Ferret, ex vivo

6.1. B_0 Field Mapping

Acquiring a B_0 field map—typically performed at the beginning of an MRI session—is generally fast, reliable, and not considered burdensome. The standard approach uses a dual-gradient-echo acquisition in which the phase difference between echoes encodes the B_0 off-resonance. While B_0 field maps can be used as inputs to MRI control design in

the same way as B_1^+ maps, they are commonly also used to correct some distortions in other scans, e.g., as presented in Refs. [85,97–99]. Such distortion corrections fall on the post-processing side, which this review does not cover, but given the work in this area, actual B_0 field maps were retrieved from distorted images by NNs in Refs. [85,98,99].

For example, Haskell et al. [85] consider the challenge of distortion correction of longer fMRI sessions with high temporal resolution, where repeated acquisitions of standard B_0 field maps interleaved with the fMRI scans is inconvenient. Building on the work of Zeng et al. [97], they use a residual CNN to estimate temporally resolved B_0 field maps from single-echo spiral-based fMRI scans. This principle could be leveraged in future applications that require frequent updates to MRI control parameters.

Another example along this line will be discussed in Section 6.4 on AI-based assessment of motion-induced changes to B_0 field maps. Beyond such use cases—and aside from plausible post-processing topics such as denoising or super-resolution—the authors are not aware of broader investigations into AI-based prediction of B_0 field maps from localizer or scout scans. In contrast, estimating B_1^+ field maps is considerably more challenging, and AI methods have received far more attention.

6.2. B_1^+ Field Mapping

For convenience, conventional (non-AI based) B_1^+ field mapping sequences include (but are not limited to) Refs. [100–104]. Although acquisition speed, precision, and dynamic range continue to improve, there is growing interest in AI-supported B_1^+ field mapping techniques.

The T_2-FLAIR sequence is SAR-intensive due to the required adiabatic inversion pulses. Abbasi-Rad et al. [86] used 1-channel B_1^+ field maps to scale the RF pulses slice by slice while preserving adiabaticity. This in turn saved up to ~ 90 s of SAR-induced delay for the same image T_2-FLAIR quality, i.e., mitigating B_1^+ inhomogeneity with less SAR load. The magnitude B_1^+ field maps were obtained through a modified U-Net fed with 3D localizer scans and showed robustness towards oblique transverse, sagittal, and coronal orientations. The application was brain scans at 7 T (in vivo).

Eberhardt et al. [61] sought to accelerate multi-channel B_1^+ field mapping for parallel transmit coils by acquiring only a subset of the channel-specific maps and synthesizing the missing channels using generative adversarial networks (GANs). They compared the pix2pixHD [105] and spatially-adaptive denormalization (SPADE) [106] GANs, and found that, under their given computational resources, which limited the SPADE model size, pix2pixHD performed favorably. SPADE can predict all missing channels in a single pass, whereas pix2pixHD predicts them recursively. SPADE, when trained as a VAE, would also be simpler to integrate into the full 2-spoke RF-shimming pipeline (see Section 5.1). The authors noted that with stronger computational resources, SPADE might have matched or surpassed pix2pixHD.

For single-slice, 8-channel B_1^+ field map prediction, Krüger et al. [92] developed a U-Net that ingested real and imaginary 2D localizer data from all 32 receive channels plus the magnitude (input size: $96 \times 96 \times 65$). After a shared U-Net encoder, the signal splits into eight parallel decoders, each outputting the channel-specific B_1^+ field maps (overall output size: $96 \times 96 \times 16$). The challenging application was the human heart and thorax at 7 T, and the CNN was able to predict B_1^+ field maps successfully compared to the ground truth B_1^+ field maps in only 0.2 s per slice. The in vivo demonstration included phase-only RF shimming based on the predicted B_1^+ field maps.

Krüger et al. [94] subsequently extended the single-slice 2D B_1^+ field map prediction method with complex-valued CNNs. The application this time was brain scans at 7 T with an 8-channel transmit/receive coil. Cost function, activation functions, CNN, input, and output were complex-valued (input size: $96 \times 96 \times 9$, output size: $96 \times 96 \times 8$). The

authors furthermore successfully demonstrated B_1^+ field map prediction for multiple slice orientations with the complex-valued CNNs. On performance, which was evaluated in comparison to CNNs that handle real and imaginary parts side-by-side, the benefit of complex-valued CNNs over their counterparts was less pronounced. The indications observed (although without statistical significance), e.g., regarding structural similarity, could, with further investigation, potentially yield an advantage to enhanced phase prediction [94]. Whole-head B_1^+ field map estimation was performed (slice-by-slice) within a second, and results were demonstrated in vivo using 4-subpulses k_T-points pulses.

The B_1^+ field maps are important in RF pulse design for reaching target FAs. However, the corresponding E fields emanating along with the B_1^+ fields yield a concern with regard to SAR.

6.3. SAR Assessment

Because local SAR cannot be measured directly in vivo, SAR safety assessment typically relies on electromagnetic simulations that compute E and B_1^+ field maps, and SAR maps for a given RF coil setup and a digitalized anatomical model. Such tissue models are usually generic but anatomically realistic numerical phantoms, like VF [49]. However, because these phantoms are not subject-specific, safety margins must account for inter-subject variability, leading to conservative limits in RF pulse design.

Brink et al. [89] proposed a 2.5D CNN based on the ForkNET topology [107] to segment T_1-weighted images into internal air, bone, muscle, fat, white matter, gray matter, cerebrospinal fluid, and eye compartments. This allows the construction of subject-specific anatomical models in which each voxel is assigned relative permittivity, conductivity, and density from standard reference tables. Segmentation, which typically requires hours, was reduced to 14 s. Multi-contrast scans from 10 healthy volunteers at 7 T were used for data generation (semi-automated reference labels), but the network was trained only with a T_1-weighted scan as input (3D MPRAGE, 2 min 54 s acquisition time). For a single-channel birdcage coil, peak 10-g SAR errors remained below 3%, while a 6.3% safety margin sufficed for conservative SAR estimates in 95% of cases in the multi-channel setting. The complete pipeline—segmentation plus SAR computation—required 6 min and 30 min for single- and multi-channel setups, respectively.

Brink et al. [90] also proposed to use the ForkNET framework [107] for directly predicting the three tissue parameters that effectively go into the electromagnetic simulation: relative permittivity, conductivity, and density. But this time, the input would be the localizer scan (9 s acquisition time). Data was obtained from 20 healthy subjects at 7 T. The output that would also include masks for background and internal air cavities would be fed to a fast electromagnetic solver, outputting parallel transmit B_1^+ field maps and 10-g local SAR maps. Network inference took 1 s, and total simulation times were 30 s (birdcage) and 45 s (parallel transmit). Their method on peak 10-g SAR ultimately indicated a conservative margin of 8.5% would suffice for 95% coverage.

Because full electromagnetic simulations remain computationally demanding, Meliadò et al. [84] proposed predicting local SAR maps directly from measured complex B_1^+ maps using a CNN. To meet the 10-g averaging requirement—which corresponds to a tissue cube of approximately 2.2 cm side length—the network input consisted of five consecutive 2D slices of the B_1^+ field map (each 0.5 cm thick), and the output was the 10-g SAR map of the center slice. The input was decomposed into five real B_1^+ images, five imaginary B_1^+ images, and five binary tissue-air mask images. The network contained a U-Net generator and a PatchGAN classification discriminator, yielding a Pix2Pix-style [108] conditional generative adversarial network (cGAN). Data from 23 volunteers [109] and 250 random phase settings yielded 5750 training samples that included E- and B_1^+-fields and 10-g local

SAR maps from finite-difference time-domain electromagnetic simulations (ground truth). The network produced realistic SAR maps within a few milliseconds. A reduced safety factor enabled an estimated 25% reduction in scan time for SAR-limited protocols. The method was validated in vivo at 7 T for pelvis imaging.

Gokyar et al. [87] demonstrated by an in silico proof-of-principle that SAR maps via a 2D U-Net can be predicted from anatomical maps. Training data consisted of single-channel transmit, 3 T, and 7 T body model simulations. A four-stage U-Net was found optimal with a mean square error <1% for both 7 T head and 3 T body cases, and structural similarity indices >90% and >80% for 7 T head and 3 T body cases, respectively.

6.4. Motion Compensation

Motion is a major challenge in MRI. Subject-tailored parallel-transmit RF pulses are sensitive to subject motion, since the B_1^+ fields strongly depend on anatomy-to-coil positioning. Even small rotations of approximately 5° can induce 12–19% FA errors at 7 T [110].

Plumley et al. [88] proposed a cGAN-based framework to predict motion-induced changes in B_1^+ field maps at 7 T using an 8-channel transmit coil. Eight-layer encoder-decoder U-Net generators were trained separately for B_1^+ magnitude and phase, processing volumes of size $256 \times 256 \times 8$. Training data consisted of B_1^+ field maps simulated in Sim4Life (ZMT Zurich MedTech AG, Zurich, Switzerland) for grids of translations and yaw rotations applied to numerical head models (Ella and Duke for training; Billie for validation/testing; Dizzy for testing) [111]. Five B_1^+ magnitude- and five B_1^+ phase-predicting networks were trained for a subset of the simulated grid points. In operation—after a given head movement to a new, arbitrary position—one would cascade the relevant networks in a manner to propagate the baseline B_1^+ field map towards one that represents the new head pose. Compared with ground-truth simulations, predicted B_1^+ field maps exhibited smaller errors than those introduced by motion itself. When used to design 5-spoke parallel transmit pulses, the predicted B_1^+ field maps yielded reduced FA errors relative to using uncorrected B_1^+ field maps.

Subject motion also perturbs the B_0 field, degrading image and spectroscopy quality. Traditional dynamic B_0 correction methods (e.g., navigators or field probes) may increase sequence complexity or require additional hardware. Motyka et al. [93] instead propose a deep learning approach to predict motion-induced B_0 field changes directly from head pose information, enabling motion correction without additional scan time or sequence modification. The 3D U-Net input consisted of (i) an anatomical reference at the initial pose, (ii) the corresponding baseline B_0 field map, and (iii) the anatomy transformed into a new pose via a six-degree-of-freedom rigid transformation. The input size was $128 \times 128 \times 80 \times 3$. The output was the B_0 field map at the new pose (size $128 \times 128 \times 80$). The authors trained a 4-stage 3D U-Net on 11 subjects with 30 arbitrary head positions (and corresponding B_0 field maps) each, and started to see B_0 field map predictions approaching ground truth in the test set—testing involved four extra subjects. Additional improvement was then achieved using a fast, subject-specific, fine-training stage. Data from six head positions would be gathered at the beginning of a session and used to fine-tune a subject-specific U-Net, using the common U-Net as an initial guess. This fine-training stage required about 1 min of training, and it was shown that as few as three head positions (instead of six) would potentially suffice for fine-training. Overall, the U-Net was able to predict B_0 field map changes, agreeing well with the echo-planar-imaging (EPI) based B_0 field maps.

6.5. Gradient Discrepancy Compensation

Compensating for discrepancies between nominal and effective gradient waveforms benefits both image reconstruction and sequence performance. Early work by Hwang et al. [83] demonstrated a NN for optimizing pre-emphasis settings for the EPI sequence. Dan et al. [95] recently revisited this problem using a bidirectional long short-term memory (LSTM) network [112] that iteratively learned pre-emphasis parameters, reporting that the corrected gradient waveforms achieved "*near-perfect alignment with the ideal waveform*."

As much as GIRF-based gradient correction methods have proven useful, their limitation of modeling the gradient system as being LTI has a certain impact when non-linear effects are more severe [113]. To develop and evaluate a nonlinear, data-driven gradient system model, Martin et al. [96] presented a temporal convolutional network (TCN) that accurately predicts readout gradient deviations and improves image reconstruction. Using a pulse-sequence-based trajectory-mapping technique [114], they acquired 18 gradient waveforms (spirals, chirps, multisines, trapezoids, triangles) across all three gradient axes and with amplitudes from 70 to 700 mT/m. The TCN consisted of a cascade of five residual blocks that each twice performed causal convolution followed by a weight normalization, a Gaussian Error Linear Unit activation function, and a dropout layer. A 1D convolution (kernel-size 1) further bypassed the two causal convolution steps. Hyperparameters were fine-tuned with Optuna [115]. The TCN method was compared to measured data (test set); a GIRF-based correction method (established via the measured data set rather than from a field camera); and the related method of Liu et al. [91] who proposed an LSTM method for gradient waveform prediction. Experiments on phantoms and ex vivo ferret brains at a pre-clinical 7 T system demonstrated that the TCN achieved the best predictive accuracy among all methods. Improved diffusion parameter maps were also observed. The authors noted some areas for further development, e.g., the TCNs were trained for each gradient channel separately, so cross-axis coupling was not captured; training used phantom data only, omitting potential subject-specific dynamic effects; and temperature effects were not explicitly modeled but are known to influence gradient behavior [113]. However, regarding temperature, it was noted that imaging experiments were performed with an initially "cold" gradient system, but the extended acquisition time for training data acquisition likely reflected a "warmed-up" system. Finally, the TCN was trained with the scanner's built-in eddy-current compensation enabled, demonstrating that even with such compensation, residual discrepancies between nominal and effective gradients remain.

7. Conclusions and Future Directions

Deep learning methods have rapidly matured into practical tools for addressing long-standing physical and engineering limitations in MRI. Figure 1 summarizes the tasks for which AI has been employed, and Figure 2 summarizes the AI paradigms and network types exploited to address the tasks. Across RF safety and parallel transmit calibration, motion compensation, and gradient system characterization, data-driven models consistently demonstrate the ability to approximate complex electromagnetic or system-dynamical behavior with unprecedented speed. These approaches reduce reliance on time-consuming numerical simulations, enable subject-specific predictions during the scan, and open the door to more adaptive, personalized, and efficient imaging workflows.

In the context of SAR assessment, convolutional and adversarial architectures have shown that subject-specific anatomical modeling, electromagnetic parameter estimation, and even direct SAR prediction can be performed orders of magnitude faster than traditional simulation. Importantly, the emerging consensus is not that deep learning necessarily replaces full-wave solvers, but that it augments them: neural networks provide fast surrogates whose inaccuracies are well-characterized and can be mitigated with modest safety

margins. As such, data-driven SAR frameworks offer a path toward safer and more flexible parallel transmit pulse design, particularly at ultra-high field.

In motion handling, deep learning models demonstrate the ability to predict changes in both B_1^+ and B_0 fields arising from rigid head motion, thereby complementing or reducing the need for complex hardware, navigators, or real-time field probes. Notably, subject-specific fine-tuning strategies indicate that modest amounts of individualized data can substantially improve performance, hinting at a future in which deep learning models dynamically track the subject throughout the exam.

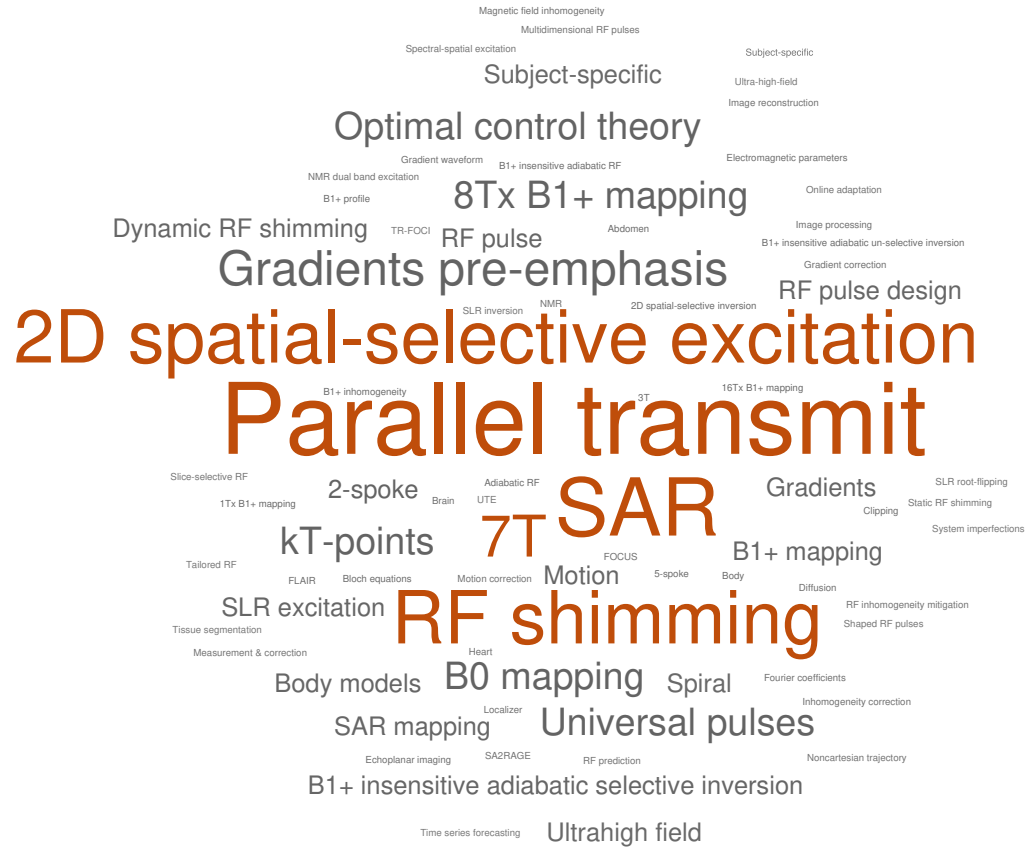


Figure 1. Digested word cloud on article keywords and employments of studies mentioned in Tables 1 and 2. The processed list of words has excluded "MRI", and anything related to AI. Articles lacking keywords were given representative keywords at the authors' discretion. Similar keywords were made coherent, e.g., for "Parallel transmit" and "Parallel transmission", the latter was changed to the first form.

For gradient system characterization, temporal convolutional and recurrent models now achieve predictive accuracy beyond classical LTI approaches, especially when non-linearities become significant. These results underscore a broader trend: MRI hardware, historically modeled with simplifying assumptions for tractability, can often be better described by data-driven models that capture dynamics not explicitly represented in traditional frameworks.

Despite these advances, several challenges remain before widespread clinical integration. First, many proposed methods rely on limited datasets, often restricted to a few numerical models, a small subject cohort, or a narrow range of coil and sequence configurations. Broader, harmonized datasets—potentially enabled by multi-centre data sharing, privacy-preserving learning strategies, or synthetic data generation—will be essential for robust generalization. Second, prospective safety validation remains non-trivial: deterministic worst-case analyses must be adapted to probabilistic models, and uncertainty

quantification must become standard practice. Safety first. Third, integration into real-time MRI pipelines requires rigorous software engineering and regulatory-ready frameworks that ensure reproducibility, traceability, and quality assurance.

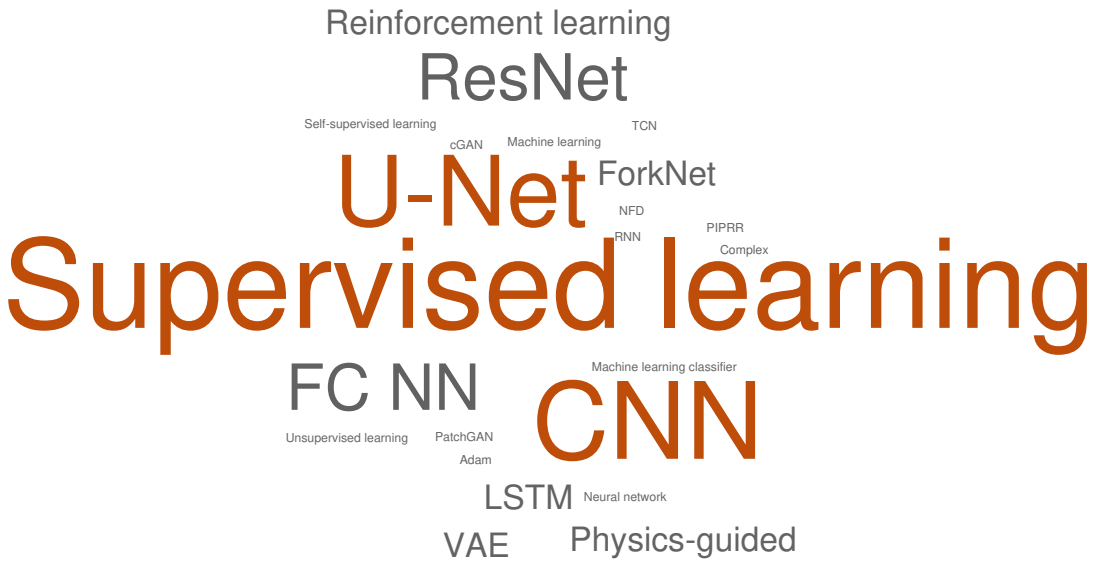


Figure 2. Digested word cloud on AI paradigm and network types of studies mentioned in Tables 1 and 2.

Looking forward, the most promising direction is perhaps a shift from task-specific networks toward holistic, physics-informed system models that jointly learn electromagnetic behaviour, motion, gradient dynamics, and even thermal or mechanical coupling. Such unified models could underpin fully adaptive MRI systems where parallel transmit pulses, shim fields, trajectories, and reconstructions are co-optimized in real time. Advances in neural operators, differentiable physics, and on-scanner GPU acceleration will likely be central to realizing this vision. In parallel, establishing standardized benchmarks and reporting practices will allow the community to more reliably compare methods and accelerate translation.

In summary, deep learning is not merely simplifying existing workflows; it could be reshaping how the MRI community conceptualizes system modeling. By efficiently capturing complex interactions between hardware, anatomy, and motion, data-driven approaches enable more agile, subject-specific, and high-performance imaging. Continued progress will depend on methodological rigor, broader datasets, uncertainty-aware safety frameworks, and close collaboration between MRI physicists, engineers, and the ML community. With these elements in place, AI has the potential to significantly expand what is achievable in MRI—particularly in the demanding regimes of ultra-high field, parallel transmit, and high-speed imaging.

Author Contributions: Writing—original draft preparation, M.S.V.; writing—review and editing, M.S.V. and T.E.L.; visualization, M.S.V. and T.E.L. All authors have read and agreed to the published version of the manuscript.

Funding: M.S.V. received financial support from Familien Andresens selskab til fremme af medicinsk forskning during the preparation of this manuscript.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Acknowledgments: During the preparation of this manuscript, the author(s) used ChatGPT (GPT-5) for the purposes of literature search support and refinement of the manuscript text. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
B_0	Static, main magnetic field
B_1^+	Magnetic component of electromagnetic field
BOLD	Blood Oxygenation Level Dependent
(c)GAN	(conditional) Generative Adversarial Network
CNN	Convolutional Neural Network
CT	Computed Tomography
DC	Direct Current
DL	Deep Learning
DTI	Diffusion Tensor Imaging
DWI	Diffusion Weighted Imaging
E	Electric component of electromagnetic field
EPI	Echo Planar Imaging
FA	Flip Angle
FC	Fully Connected
FLAIR	Fluid Attenuated Inversion Recovery
fMRI	Functional MRI
FOCUS	Fast online-customized
GAN	Generative Adversarial Network
GIRF	Gradient Impulse Response Function
GPS	Physics-Guided Self-supervised learning
GPU	Graphical Processing Unit
GRAPPA	Generalized Autocalibrating Partially Parallel Acquisitions
$G_{x,y,z}$	Magnetic field gradients along x, y, and z
JANUS	God of doors and gates in Roman mythology
LTI	Linear Time-Invariant
MLS	Magnitude Least Squares
MPRAGE	Magnetization Prepared Rapid Gradient Echo
MRA	Magnetic Resonance Angiography
MRI	Magnetic Resonance Imaging
MRS	Magnetic Resonance Spectroscopy
NA	Not applicable
NN	Neural Network
PatchGAN	Patch Generative Adversarial Network
PC	Phase Contrast
PET	Positron Emission Tomography
PhG	Physics-Guided
PIPRR	Prediction by Iteratively Projected Ridge Regression
PNS	Peripheral Nerve Stimulation
PWI	Perfusion Weighted Imaging
ReLU	Rectified Linear Unit
ResNet	Residual Network
RF	Radio Frequency
RL	Reinforcement Learning
RNN	Recurrent Neural Network
RT-MRI	Realtime MRI

SAR	Specific Absorption Rate
SENSE	SENSitivity Encoding
SLR	Shinnar-Le Roux
SNR	Signal to Noise Ratio
SPADE	Spatially-Adaptive Denormalization
SED	Specific Energy Dose
SSvL	Self-Supervised Learning
SvL	Supervised Learning
SWI	Susceptibility Weighted Imaging
T	Tesla
T_1	Longitudinal relaxation time
T_2	Transverse relaxation
TCN	Temporal Convolutional Network
TOF	Time of Flight
Tx	Transmit
UsvL	Unsupervised Learning
UP	Universal Pulse
US	Ultrasound
VAE	Variational AutoEncoder
VF	Virtual Family
VOP	Virtual Observation Points

References

1. Zhou, H. A Comprehensive Review of Artificial Intelligence and Machine Learning in Control Theory. *Appl. Comput. Eng.* **2024**, *116*, 43–48. [CrossRef]
2. Soori, M.; Arezoo, B.; Dastres, R. Artificial intelligence, machine learning and deep learning in advanced robotics, a review. *Cogn. Robot.* **2023**, *3*, 54–70. [CrossRef]
3. Rahman, M.A.; Shahrior, M.F.; Iqbal, K.; Abushaiba, A.A. Enabling Intelligent Industrial Automation: A Review of Machine Learning Applications with Digital Twin and Edge AI Integration. *Automation* **2025**, *6*, 37. [CrossRef]
4. Zaoui, A.; Tchuente, D.; Wamba, S.F.; Kamsu-Foguem, B. Impact of artificial intelligence on aeronautics: An industry-wide review. *J. Eng. Technol. Manag.* **2024**, *71*, 101800. [CrossRef]
5. Haacke, E.M.; Brown, R.W.; Thompson, M.R.; Venkatesan, R. *Magnetic Resonance Imaging: Physical Principles and Sequence Design*; Wiley-Liss: New York, NY, USA, 1999.
6. Bihan, D.L.; Breton, E.; Lallemand, D.; Grenier, P.; Cabanis, E.; Laval-Jeantet, M. MR imaging of intravoxel incoherent motions: Application to diffusion and perfusion in neurologic disorders. *Radiology* **1986**, *161*, 401–407. [CrossRef] [PubMed]
7. Basser, P.; Mattiello, J.; LeBihan, D. MR Diffusion Tensor Spectroscopy and Imaging. *Biophys. J.* **1994**, *66*, 259–267. [CrossRef] [PubMed]
8. Rosen, B.R.; Belliveau, J.W.; Vevea, J.M.; Brady, T.J. Perfusion imaging with NMR contrast agents. *Magn. Reson. Med.* **1990**, *14*, 249–265. [CrossRef]
9. Ogawa, S.; Lee, T.M.; Kay, A.R.; Tank, D.W. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc. Natl. Acad. Sci. USA* **1990**, *87*, 9868–9872. [CrossRef]
10. Muthupillai, R.; Lomas, D.J.; Rossman, P.J.; Greenleaf, J.F.; Manduca, A.; Ehman, R.L. Magnetic Resonance Elastography by Direct Visualization of Propagating Acoustic Strain Waves. *Science* **1995**, *269*, 1854–1857. [CrossRef]
11. Dumoulin, C.L.; Hart, H.R. Magnetic resonance angiography. *Radiology* **1986**, *161*, 717–720. [CrossRef]
12. Bottomley, R.R. Spatial Localization in NMR Spectroscopy in Vivo. *Ann. N. Y. Acad. Sci.* **1987**, *508*, 333–348. [CrossRef]
13. Haacke, E.M.; Xu, Y.; Cheng, Y.C.N.; Reichenbach, J.R. Susceptibility weighted imaging (SWI). *Magn. Reson. Med.* **2004**, *52*, 612–618. [CrossRef]
14. Poorter, J.D.; Wagter, C.D.; Deene, Y.D.; Thomsen, C.; Ståhlberg, F.; Achten, E. Noninvasive MRI Thermometry with the Proton Resonance Frequency (PRF) Method: In Vivo Results in Human Muscle. *Magn. Reson. Med.* **1995**, *33*, 74–81. [CrossRef]
15. Pruessmann, K.P.; Weiger, M.; Scheidegger, M.B.; Boesiger, P. SENSE: Sensitivity encoding for fast MRI. *Magn. Reson. Med.* **1999**, *42*, 952–962. [CrossRef]
16. Griswold, M.A.; Jakob, P.M.; Heidemann, R.M.; Nittka, M.; Jellus, V.; Wang, J.; Kiefer, B.; Haase, A. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn. Reson. Med.* **2002**, *47*, 1202–1210. [CrossRef] [PubMed]
17. Zhu, Y. Parallel excitation with an array of transmit coils. *Magn. Reson. Med.* **2004**, *51*, 775–784. [CrossRef] [PubMed]

18. Katscher, U.; Börnert, P.; Leussler, C.; van den Brink, J.S. Transmit SENSE. *Magn. Reson. Med.* **2003**, *49*, 144–150. [CrossRef] [PubMed]
19. Vinding, M.S.; Brenner, D.; Tse, D.H.Y.; Vellmer, S.; Vosegaard, T.; Suter, D.; Stöcker, T.; Maximov, I.I. Application of the limited-memory quasi-Newton algorithm for multi-dimensional, large flip-angle RF pulses at 7T. *Magn. Reson. Mater. Phys. Biol. Med.* **2017**, *30*, 29–39. [CrossRef]
20. Goodwin, D.L.; Vinding, M.S. Accelerated Newton-Raphson GRAPE methods for optimal control. *Phys. Rev. Res.* **2023**, *5*, L012042. [CrossRef]
21. Vinding, M.S.; Guérin, B.; Vosegaard, T.; Nielsen, N.C. Local SAR, global SAR, and power-constrained large-flip-angle pulses with optimal control and virtual observation points. *Magn. Reson. Med.* **2017**, *77*, 374–384. [CrossRef]
22. Vinding, M.S.; Skyum, B.; Sangill, R.; Lund, T.E. Ultrafast (milliseconds), multidimensional RF pulse design with deep learning. *Magn. Reson. Med.* **2019**, *82*, 586–599. [CrossRef]
23. Vinding, M.S.; Skyum, B.; Sangill, R.; Lund, T.E. Rapid multi-dimensional RF pulse design with deep learning. In Proceedings of the ISMRM 27th Annual Meeting & Exhibition, Montreal, QC, Canada, 11–16 May 2019; p. 4643.
24. Foo, T.K.F.; Tan, E.T.; Vermilyea, M.E.; Hua, Y.; Fiveland, E.W.; Piel, J.E.; Park, K.; Ricci, J.; Thompson, P.S.; Graziani, D.; et al. Highly efficient head-only magnetic field insert gradient coil for achieving simultaneous high gradient amplitude and slew rate at 3.0T (MAGNUS) for brain microstructure imaging. *Magn. Reson. Med.* **2020**, *83*, 2356–2369. [CrossRef]
25. Versteeg, E.; Klomp, D.W.J.; Siero, J.C.W. A silent gradient axis for soundless spatial encoding to enable fast and quiet brain imaging. *Magn. Reson. Med.* **2022**, *87*, 1062–1073. [CrossRef] [PubMed]
26. Juchem, C.; Nixon, T.W.; McIntyre, S.; Boer, V.O.; Rothman, D.L.; de Graaf, R.A. Dynamic multi-coil shimming of the human brain at 7T. *J. Magn. Reson.* **2011**, *212*, 280–288. [CrossRef]
27. Stockmann, J.P.; Witzel, T.; Keil, B.; Polimeni, J.R.; Mareyam, A.; LaPierre, C.; Setsompop, K.; Wald, L.L. A 32-channel combined RF and B_0 shim array for 3T brain imaging. *Magn. Reson. Med.* **2016**, *75*, 441–451. [CrossRef]
28. Vinding, M.S.; Lund, T.E.; Stockmann, J.P.; Guerin, B. INSTANT (INtegrated Shimming and Tip-Angle NormalizaTion): 3D flip-angle mitigation using joint optimization of RF and shim array currents. In Proceedings of the ISMRM & SMRT Virtual Conference & Exhibition, Virtual, 8–14 August 2020; p. 0612.
29. 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. [CrossRef] [PubMed]
30. Gallichan, D.; Marques, J.P.; 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. [CrossRef]
31. Krause, F.; Benjamins, C.; Eck, J.; Lührs, M.; van Hoof, R.; Goebel, R. Active head motion reduction in magnetic resonance imaging using tactile feedback. *Hum. Brain Mapp.* **2019**, *40*, 4026–4037. [CrossRef]
32. Knudsen, L.; Bailey, C.J.; Blicher, J.U.; Yang, Y.; Zhang, P.; Lund, T.E. Improved sensitivity and microvascular weighting of 3T laminar fMRI with GE-BOLD using NORDIC and phase regression. *NeuroImage* **2023**, *271*, 120011. [CrossRef]
33. Nayak, K.S.; Lim, Y.; Campbell-Washburn, A.E.; Steeden, J. Real-Time Magnetic Resonance Imaging. *J. Magn. Reson. Imaging* **2022**, *55*, 81–99. [CrossRef] [PubMed]
34. Voyvodic, J.T. Real-Time fMRI Paradigm Control, Physiology, and Behavior Combined with Near Real-Time Statistical Analysis. *NeuroImage* **1999**, *10*, 91–106. [CrossRef]
35. Onodera, T.; Matsui, S.; Sekihara, K.; Kohno, H. A method of measuring field-gradient modulation shapes. Application to high-speed NMR spectroscopic imaging. *J. Phys. E Sci. Instruments* **1987**, *20*, 416. [CrossRef]
36. Takahashi, A.; Peters, T. Compensation of multi-dimensional selective excitation pulses using measured *k*-space trajectories. *Magn. Reson. Med.* **1995**, *34*, 446–456. [CrossRef] [PubMed]
37. Dietrich, B.E.; Brunner, D.O.; Wilm, B.J.; Barmet, C.; Gross, S.; Kasper, L.; Haeberlin, M.; Schmid, T.; Vannesjo, S.J.; Pruessmann, K.P. A field camera for MR sequence monitoring and system analysis. *Magn. Reson. Med.* **2016**, *75*, 1831–1840. [CrossRef]
38. Vannesjo, S.J.; Haeberlin, M.; Kasper, L.; Pavan, M.; Wilm, B.J.; Barmet, C.; Pruessmann, K.P. Gradient system characterization by impulse response measurements with a dynamic field camera. *Magn. Reson. Med.* **2013**, *69*, 583–593. [CrossRef]
39. Aigner, C.S.; Rund, A.; Abo Seada, S.; Price, A.N.; Hajnal, J.V.; Malik, S.J.; Kunisch, K.; Stollberger, R. Time optimal control-based RF pulse design under gradient imperfections. *Magn. Reson. Med.* **2020**, *83*, 561–574. [CrossRef]
40. *IEC 60601-2-33:2022*; Medical Electrical Equipment—Part 2-33: Particular Requirements for the Basic Safety and Essential Performance of Magnetic Resonance Equipment for Medical Diagnosis. IEC: Geneva, Switzerland, 2022.
41. Malik, S.J.; Keihaninejad, S.; Hammers, A.; Hajnal, J.V. Tailored excitation in 3D with spiral nonselective (SPINS) RF pulses. *Magn. Reson. Med.* **2012**, *67*, 1303–1315. [CrossRef]
42. 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. [CrossRef] [PubMed]

43. Herrler, J.; Liebig, P.; Gumbrecht, R.; Ritter, D.; Schmitter, S.; Maier, A.; Schmidt, M.; Uder, M.; Doerfler, A.; Nagel, A.M. Fast online-customized (FOCUS) parallel transmission pulses: A combination of universal pulses and individual optimization. *Magn. Reson. Med.* **2021**, *85*, 3140–3153. [CrossRef] [PubMed]
44. Marques, J.P.; Kober, T.; Krueger, G.; van der Zwaag, W.; Van de Moortele, P.F.; Gruetter, R. MP2RAGE, a self bias-field corrected sequence for improved segmentation and T_1-mapping at high field. *NeuroImage* **2010**, *49*, 1271–1281. [CrossRef]
45. Ruijters, L.; Lund, T.E.; Vinding, M.S. Improving B_1^+-inhomogeneity tolerance by resolving non-bijection in MP2RAGE R_1 mapping: A 2D look-up table approach demonstrated at 3 T. *Magn. Reson. Med.* **2025**, *93*, 1712–1722. [CrossRef]
46. Delgado, P.R.; Kuehne, A.; Periquito, J.S.; Millward, J.M.; Pohlmann, A.; Waiczies, S.; Niendorf, T. B_1 inhomogeneity correction of RARE MRI with transceive surface radiofrequency probes. *Magn. Reson. Med.* **2020**, *84*, 2684–2701. [CrossRef] [PubMed]
47. Mauconduit, F.; Massire, A.; Gras, V.; Amadon, A.; Vignaud, A.; Boulant, N. PASTeUR package extension with MP2RAGE for robust T1 mapping technique in parallel transmit at 7T. In Proceedings of the ISMRM & SMRT Virtual Conference & Exhibition, Virtual, 8–14 August 2020; p. 3709.
48. Guérin, B.; Gebhardt, M.; Cauley, S.; Adalsteinsson, E.; Wald, L.L. Local specific absorption rate (SAR), global SAR, transmitter power, and excitation accuracy trade-offs in low flip-angle parallel transmit pulse design. *Magn. Reson. Med.* **2014**, *71*, 1446–1457. [CrossRef]
49. Christ, A.; Kainz, W.; Hahn, E.G.; Honegger, K.; Zefferer, M.; Neufeld, E.; Rascher, W.; Janka, R.; Bautz, W.; Chen, J.; et al. The Virtual Family—Development of surface-based anatomical models of two adults and two children for dosimetric simulations. *Phys. Med. Biol.* **2009**, *55*, N23. [CrossRef]
50. Eichfelder, G.; Gebhardt, M. Local specific absorption rate control for parallel transmission by virtual observation points. *Magn. Reson. Med.* **2011**, *66*, 1468–1476. [CrossRef] [PubMed]
51. Gezelter, J.D.; Freeman, R. Use of neural networks to design shaped radiofrequency pulses. *J. Magn. Reson. 1969* **1990**, *90*, 397–404. [CrossRef]
52. Mirfin, C.; Glover, P.; Bowtell, R. Optimisation of parallel transmission radiofrequency pulses using neural networks. In Proceedings of the Joint Annual Meeting ISMRM-ESMRMB, Paris, France, 16–21 June 2018; p. 3388.
53. Ianni, J.D.; Cao, Z.; Grissom, W.A. Machine learning RF shimming: Prediction by iteratively projected ridge regression. *Magn. Reson. Med.* **2018**, *80*, 1871–1881. [CrossRef]
54. Tomi-Tricot, R.; Gras, V.; Thirion, B.; Mauconduit, F.; Boulant, N.; Cherkaoui, H.; Zerbib, P.; Vignaud, A.; Luciani, A.; Amadon, A. SmartPulse, a machine learning approach for calibration-free dynamic RF shimming: Preliminary study in a clinical environment. *Magn. Reson. Med.* **2019**, *82*, 2016–2031. [CrossRef]
55. Shin, D.; Ji, S.; Lee, D.; Lee, J.; Oh, S.H.; Lee, J. Deep Reinforcement Learning Designed Shinnar-Le Roux RF Pulse Using Root-Flipping: $\text{DeepRF}_{\text{SLR}}$. *IEEE Trans. Med. Imaging* **2020**, *39*, 4391–4400. [CrossRef]
56. Zhang, Y.; Jiang, K.; Jiang, W.; Wang, N.; Wright, A.J.; 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*, e4443. [CrossRef]
57. Vinding, M.S.; Aigner, C.S.; Schmitter, S.; Lund, T.E. DeepControl: 2DRF pulses facilitating B_1^+ inhomogeneity and B_0 off-resonance compensation in vivo at 7 T. *Magn. Reson. Med.* **2021**, *85*, 3308–3317. [CrossRef]
58. Vinding, M.S.; Aigner, C.S.; Stockmann, J.; Guerin, B.; Schmitter, S.; Lund, T.E. DeepControl: AI-powered slice flip-angle homogenization by 2D RF pulses. In Proceedings of the ISMRM & SMRT Annual Meeting & Exhibition, Online, 15–20 May 2021; p. 0785.
59. 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*, 985–994. [CrossRef]
60. Herrler, J.; Liebig, P.; Majewski, K.; Gumbrecht, R.; Ritter, D.; Meixner, C.R.; Maier, A.; Dörfler, A.; Nagel, A.M. Neural Network-supported Fast Online-Customized (FOCUS) parallel transmit (pTx) pulses for slice-selective, large tip angle excitation. In Proceedings of the Joint Annual Meeting ISMRM-ESMRMB, London, UK, 7–12 May 2022; p. 3312.
61. Eberhardt, B.; Poser, B.A.; Shah, N.J.; Felder, J. B1 field map synthesis with generative deep learning used in the design of parallel-transmit RF pulses for ultra-high field MRI. *Z. Med. Phys.* **2022**, *32*, 334–345. [CrossRef]
62. Vinding, M.S.; Lund, T.E. Clipped DeepControl: Deep neural network two-dimensional pulse design with an amplitude constraint layer. *Artif. Intell. Med.* **2023**, *135*, 102460. [CrossRef]
63. Jang, A.; He, X.; Liu, F. Physics-guided self-supervised learning: Demonstration for generalized RF pulse design. *Magn. Reson. Med.* **2025**, *93*, 657–672. [CrossRef] [PubMed]
64. Kilic, T.; Liebig, P.; Demirel, O.B.; Herrler, J.; Nagel, A.M.; Ugurbil, K.; Akçakaya, M. Unsupervised deep learning with convolutional neural networks for static parallel transmit design: A retrospective study. *Magn. Reson. Med.* **2024**, *91*, 2498–2507. [CrossRef]
65. Nagelstrasser, S.; Herrler, J.; Vinding, M.S.; Egger, N.; Liebig, P.; Uder, M.; Zaiss, M.; Nagel; Armin, A.M. Comparison of deep learning predicted and universal pTx pulses for a 3D turbo spin echo sequence at 7T. In Proceedings of the European Society for Magnetic Resonance in Medicine and Biology, Marseille, France, 8–11 October 2025.

66. Lu, Z.; Liang, H.; Lu, M.; Wang, X.; Yan, X.; Huo, Y. Fast-RF-Shimming: Accelerate RF shimming in 7T MRI using deep learning. *Meta-Radiology* **2025**, *3*, 100166. [CrossRef] [PubMed]
67. Cao, Z.; Donahue, M.J.; Ma, J.; Grissom, W.A. Joint Design of Large-Tip-Angle Parallel RF Pulses and Blipped Gradient Trajectories. *Magn. Reson. Med.* **2016**, *75*, 1198–1208. [CrossRef]
68. Nielsen, P.S.; Georgsen, J.B.; Vinding, M.S.; Østergaard, L.R.; Steiniche, T. Computer-Assisted Annotation of Digital H&E/SOX10 Dual Stains Generates High-Performing Convolutional Neural Network for Calculating Tumor Burden in H&E-Stained Cutaneous Melanoma. *Int. J. Environ. Res. Public Health* **2022**, *19*, 14327. [CrossRef]
69. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep Residual Learning for Image Recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 27–30 June 2016.
70. Kingma, D.P.; Welling, M. Auto-Encoding Variational Bayes. *arXiv* **2022**, arXiv:1312.6114.
71. Zhao, S.; Song, J.; Ermon, S. InfoVAE: Balancing Learning and Inference in Variational Autoencoders. *Proc. AAAI Conf. Artif. Intell.* **2019**, *33*, 5885–5892. [CrossRef]
72. Setsompop, K.; Wald, L.L.; Alagappan, V.; Gagoski, B.A.; 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. [CrossRef]
73. Kingma, D.P.; Ba, J. Adam: A Method for Stochastic Optimization. *arXiv* **2017**, arXiv:1412.6980. [CrossRef]
74. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. *arXiv* **2015**, arXiv:1505.04597. [CrossRef]
75. Majewski, K. Simultaneous optimization of radio frequency and gradient waveforms with exact Hessians and slew rate constraints applied to k_T-points excitation. *J. Magn. Reson.* **2021**, *326*, 106941. [CrossRef]
76. Pauly, J.; Le Roux, P.; Nishimura, D.; Macovski, A. Parameter relations for the Shinnar-Le Roux selective excitation pulse design algorithm. *IEEE Trans. Med. Imaging* **1991**, *10*, 53–65. [CrossRef]
77. Sharma, A.; Lustig, M.; Grissom, W.A. Root-flipped multiband refocusing pulses. *Magn. Reson. Med.* **2016**, *75*, 227–237. [CrossRef] [PubMed]
78. Hu, E.J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W. LoRA: Low-Rank Adaptation of Large Language Models. *arXiv* **2021**, arXiv:2106.09685.
79. Vinding, M.S.; Maximov, I.I.; Tošner, Z.; Nielsen, N.C. Fast numerical design of spatial-selective rf pulses in MRI using Krotov and quasi-Newton based optimal control methods. *J. Chem. Phys.* **2012**, *137*, 054203. [CrossRef]
80. Deng, J.; Dong, W.; Socher, R.; Li, L.J.; Li, K.; Li; Fei-Fei, L. ImageNet: A large-scale hierarchical image database. In Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition, Miami, FL, USA, 20–25 June 2009; pp. 248–255. [CrossRef]
81. Vinding, M.S.; Goodwin, D.L.; Kuprov, I.; Lund, T.E. Optimal control gradient precision trade-offs: Application to fast generation of DeepControl libraries for MRI. *J. Magn. Reson.* **2021**, *333*, 107094. [CrossRef]
82. Wink, C.; Ferrazzi, G.; Bassenge, J.P.; Flassbeck, S.; Schmidt, S.; Schaeffter, T.; Schmitter, S. 4D flow imaging with 2D-selective excitation. *Magn. Reson. Med.* **2019**, *82*, 886–900. [CrossRef]
83. Hwang, S.C.; Yao, C.; Chen, J.H. EPI Gradient Preemphasis Optimization: Using Neural Networks. In Proceedings of the ISMRM 10th Scientific Meeting & Exhibition, Honolulu, HI, USA, 18–24 May 2002; p. 0829.
84. Meliadò, E.F.; Raaijmakers, A.J.E.; Sbrizzi, A.; Steensma, B.R.; Maspero, M.; Savenije, M.H.F.; Luijten, P.R.; van den Berg, C.A.T. A deep learning method for image-based subject-specific local SAR assessment. *Magn. Reson. Med.* **2020**, *83*, 695–711. [CrossRef]
85. Haskell, M.; Cao, A.A.; Noll, D.; Fessler, J. Deep learning field map estimation with model-based image reconstruction for off resonance correction of brain images using a spiral acquisition. In Proceedings of the ISMRM Workshop on Data Sampling & Image Reconstruction, Sedona, AZ, USA, 26–29 January 2020.
86. Abbasi-Rad, S.; O'Brien, K.; Kelly, S.; Vegh, V.; Rodell, A.; Tesiram, Y.; Jin, J.; Barth, M.; Bollmann, S. Improving FLAIR SAR efficiency at 7T by adaptive tailoring of adiabatic pulse power through deep learning B_1^+ estimation. *Magn. Reson. Med.* **2021**, *85*, 2462–2476. [CrossRef] [PubMed]
87. Gokyar, S.; Robb, F.J.L.; Kainz, W.; Chaudhari, A.; Winkler, S.A. MRSaiFE: An AI-Based Approach Towards the Real-Time Prediction of Specific Absorption Rate. *IEEE Access* **2021**, *9*, 140824–140834. [CrossRef] [PubMed]
88. Plumley, A.; Watkins, L.; Treder, M.; Liebig, P.; Murphy, K.; Kopanoglu, E. Rigid motion-resolved B_1^+ prediction using deep learning for real-time parallel-transmission pulse design. *Magn. Reson. Med.* **2022**, *87*, 2254–2270. [CrossRef]
89. Brink, W.M.; Yousefi, S.; Bhatnagar, P.; Remis, R.F.; Staring, M.; Webb, A.G. Personalized local SAR prediction for parallel transmit neuroimaging at 7T from a single T1-weighted dataset. *Magn. Reson. Med.* **2022**, *88*, 464–475. [CrossRef]
90. Brink, W.; Staring, M.; Remis, R.F.; Webb, A. Fast Subject-Specific Local SAR and B1+ Prediction for PTx at 7T using only an Initial Localizer Scan. In Proceedings of the Joint Annual Meeting ISMRM-ESMRMB & ISMRT 31st Annual Meeting, London, UK, 7–12 May 2022; p. 0387.
91. Liu, Q.; Nguyen, H.; Zhu, X.; Zhai, X.; Li, B.; Xu, J.; Zhang, W. Gradient Waveform Prediction Using Deep Neural Network. In Proceedings of the Joint Annual Meeting ISMRM-ESMRMB, London, UK, 7–12 May 2022; p. 0639.

92. Krueger, F.; Aigner, C.S.; Hammernik, K.; Dietrich, S.; Lutz, M.; Schulz-Menger, J.; Schaeffter, T.; Schmitter, S. Rapid estimation of 2D relative B_1^+-maps from localizers in the human heart at 7T using deep learning. *Magn. Reson. Med.* **2023**, *89*, 1002–1015. [CrossRef]
93. Motyka, S.; Weiser, P.; Bachrata, B.; Hingerl, L.; Strasser, B.; Hangel, G.; Niess, E.; Niess, F.; Zaitsev, M.; Robinson, S.D.; et al. Predicting dynamic, motion-related changes in B_0 field in the brain at a 7T MRI using a subject-specific fine-trained U-net. *Magn. Reson. Med.* **2024**, *91*, 2044–2056. [CrossRef] [PubMed]
94. Krueger, F.; Aigner, C.S.; Lutz, M.; Riemann, L.T.; Degenhardt, K.; Hadjikiriakos, K.; Zimmermann, F.F.; Hammernik, K.; Schulz-Menger, J.; Schaeffter, T.; et al. Deep learning-based whole-brain B_1^+-mapping at 7T. *Magn. Reson. Med.* **2025**, *93*, 1700–1711. [CrossRef]
95. Dan, G.; Zhong, Z.; Liu, Q.; Xu, J. Gradient Waveform Pre-emphasis using Deep Neural Network. In Proceedings of the ISMRM & ISMRT Annual Meeting & Exhibition, Honolulu, HI, USA, 10–15 May 2025; p. 3809.
96. Martin, J.B.; Alderson, H.E.; Gore, J.C.; Does, M.D.; Harkins, K.D. Modeling the MRI gradient system with a temporal convolutional network: Improved reconstruction by prediction of readout gradient errors. *Magn. Reson. Med.* **2026**, *95*, 286–298. [CrossRef]
97. Zeng, D.Y.; Shaikh, J.; Holmes, S.; Brunsing, R.L.; Pauly, J.M.; Nishimura, D.G.; Vasanawala, S.S.; Cheng, J.Y. Deep residual network for off-resonance artifact correction with application to pediatric body MRA with 3D cones. *Magn. Reson. Med.* **2019**, *82*, 1398–1411. [CrossRef]
98. Schote, D.; Winter, L.; Kolbitsch, C.; Zimmermann, F.; O'Reilly, T.; Webb, A.; Seifert, F.; Kofler, A. Physics-informed Deep Learning for Image Distortion Correction from B0-inhomogeneities in Low field MRI. In Proceedings of the Joint Annual Meeting ISMRM-ESMRMB, London, UK, 7–12 May 2022; p. 1816. [CrossRef]
99. Schote, D.; Winter, L.; Kolbitsch, C.; Rose, G.; Speck, O.; Kofler, A. Joint B_0 and Image Reconstruction in Low-Field MRI by Physics-Informed Deep-Learning. *IEEE Trans. Biomed. Eng.* **2024**, *71*, 2842–2853. [CrossRef] [PubMed]
100. Yarnykh, V.L. Actual flip-angle imaging in the pulsed steady state: A method for rapid three-dimensional mapping of the transmitted radiofrequency field. *Magn. Reson. Med.* **2007**, *57*, 192–200. [CrossRef]
101. Moortele, V.D.; Snyder, C.; DelaBarre, L.; Adriany, G.; Vaughan, T.; Ugurbil, K. Calibration Tools for RF Shim at Very High Field with Multiple Element RF Coils: From Ultra Fast Local Relative Phase to Absolute Magnitude B1+ Mapping. In Proceedings of the Joint Annual Meeting ISMRM-ESMRMB, Berlin, Germany, 19–25 May 2007; p. 1676.
102. Ehses, P.; Brenner, D.; Stirnberg, R.; Pracht, E.D.; Stöcker, T. Whole-brain B_1-mapping using three-dimensional DREAM. *Magn. Reson. Med.* **2019**, *82*, 924–934. [CrossRef]
103. Kent, J.L.; Dragonu, I.; Valkovič, L.; Hess, A.T. Rapid 3D absolute B_1^+ mapping using a sandwiched train presaturated TurboFLASH sequence at 7 T for the brain and heart. *Magn. Reson. Med.* **2023**, *89*, 964–976. [CrossRef]
104. Lutz, M.; Aigner, C.S.; Flassbeck, S.; Krueger, F.; Gatefait, C.G.F.; Kolbitsch, C.; Silemek, B.; Seifert, F.; Schaeffter, T.; Schmitter, S. B1-MRF: Large dynamic range MRF-based absolute B_1^+ mapping in the human body at 7T. *Magn. Reson. Med.* **2024**, *92*, 2473–2490. [CrossRef] [PubMed]
105. Wang, T.C.; Liu, M.Y.; Zhu, J.Y.; Tao, A.; Kautz, J.; Catanzaro, B. High-Resolution Image Synthesis and Semantic Manipulation With Conditional GANs. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Salt Lake City, UT, USA, 19–21 June 2018.
106. Park, T.; Liu, M.Y.; Wang, T.C.; Zhu, J.Y. Semantic Image Synthesis With Spatially-Adaptive Normalization. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 15–20 June 2019.
107. Rashed, E.A.; Diao, Y.; Hirata, A. Learning-based estimation of dielectric properties and tissue density in head models for personalized radio-frequency dosimetry. *Phys. Med. Biol.* **2020**, *65*, 065001. [CrossRef] [PubMed]
108. Isola, P.; Zhu, J.Y.; Zhou, T.; Efros, A.A. Image-to-Image Translation with Conditional Adversarial Networks. *arXiv* **2018**, arXiv:1611.07004.
109. Meliadò, E.F.; van den Berg, C.A.T.; Luijten, P.R.; Raaijmakers, A.J.E. Intersubject specific absorption rate variability analysis through construction of 23 realistic body models for prostate imaging at 7T. *Magn. Reson. Med.* **2019**, *81*, 2106–2119. [CrossRef]
110. Kopanoglu, E.; Plumley, A.J.D.; Erturk, M.A.; Deniz, C.M.; Wise, R.G. Implications of within-scan patient head motion on B1+ homogeneity and Specific Absorption Rate at 7T. In Proceedings of the ISMRM 27th Annual Meeting & Exhibition, Montreal, QC, Canada, 11–16 May 2019; p. 4686.
111. Gosselin, M.C.; Neufeld, E.; Moser, H.; Huber, E.; Farcito, S.; Gerber, L.; Jedensjö, M.; Hilber, I.; Gennaro, F.D.; Lloyd, B.; 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. [CrossRef]
112. Hochreiter, S.; Schmidhuber, J. Long Short-Term Memory. *Neural Comput.* **1997**, *9*, 1735–1780. [CrossRef] [PubMed]
113. Nussbaum, J.; Dietrich, B.E.; Wilm, B.J.; Pruessmann, K.P. Thermal variation in gradient response: Measurement and modeling. *Magn. Reson. Med.* **2022**, *87*, 2224–2238. [CrossRef] [PubMed]

114. Harkins, K.D.; Does, M.D. Efficient gradient waveform measurements with variable-prephasing. *J. Magn. Reson.* **2021**, *327*, 106945. [CrossRef] [PubMed]
115. Akiba, T.; Sano, S.; Yanase, T.; Ohta, T.; Koyama, M. Optuna: A Next-generation Hyperparameter Optimization Framework. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '19, New York, NY, USA, 4–8 August 2019; pp. 2623–2631. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.