

Gradient and shim pre-emphasis by inversion of a linear time-invariant system model

S. Johanna Vannesjo^{1*}, Yolanda Duerst¹, Laetitia Vionnet¹, Benjamin E. Dietrich¹, Matteo Pavan¹, Simon Gross¹, Christoph Barmet^{1,2}, Klaas P. Pruessmann¹

¹*Institute for Biomedical Engineering, University of Zurich and ETH Zurich, Zurich, Switzerland*

²*Skope Magnetic Resonance Technologies, 8044 Zurich, Switzerland*

** Current address: FMRIB centre, Nuffield Department of Clinical Neurosciences, University of Oxford, Oxford, UK*

Word count body text: 5655

Word count appendices: 475

Running title: Gradient and shim pre-emphasis by inversion of an LTI model

Correspondence address

Prof. Klaas P. Pruessmann

Institute for Biomedical Engineering

University and ETH Zurich

Gloriastrasse 35

CH-8092 Zurich, Switzerland

E-mail: pruessmann@biomed.ee.ethz.ch

Tel: +41 44 632 6696

Fax: +41 44 632 1193

Abstract

Purpose

The goal of this contribution is to enhance the fidelity and switching speed of gradient and shim fields by advancing pre-emphasis towards broadband and full cross-term correction.

Theory and Methods

The proposed approach is based on viewing gradient and shim chains as linear, time-invariant (LTI) systems. Pre-emphasis is accomplished by inversion of a broadband digital system model. In the multiple-channel case, it amounts to a matrix of broadband filters that perform concerted self- and cross-term correction. This approach is demonstrated with gradients and shims up to the third order in a 7T whole-body MR system.

Results

Pre-emphasis by LTI model inversion is first verified by studying settling speeds and response behavior without and with the correction. It is then demonstrated for rapid shim updating, achieving substantially enhanced fidelity of field dynamics and shim settling within about 1 ms. In single-shot EPI acquisitions in-vivo this benefit is shown to translate into enhanced geometric fidelity.

Conclusions

The fidelity of gradient and shim dynamics can be greatly enhanced by pre-emphasis based on inverting a general LTI system model. Permitting shim settling on the millisecond scale, broadband multiple-channel pre-emphasis promises to render higher-order shimming fully versatile at the level of MRI sequence design.

Keywords

Higher-order shimming

Dynamic shimming

Pre-emphasis

SIRF

Introduction

In MRI systems, spatial encoding relies mostly on gradient fields operated in the kHz range. For effective encoding the gradient fields must be highly accurate functions of space and time. However, the response characteristics of gradient systems are compromised by a variety of mechanisms such as eddy currents, gradient coil interactions, mechanical vibrations, and amplifier limitations. Resulting imperfections are not limited to the time domain but also include cross-term responses of unintended spatial structure. Eddy current effects can be significantly reduced, but not completely eliminated, by active shielding of the gradient coils (1,2). Likewise, coil vibrations, which give rise to oscillatory field responses at mechanical resonance frequencies, can be suppressed but not fully removed by appropriate hardware design. Overall, despite increasing sophistication of gradient hardware, residual imperfections remain and can cause artifacts in sensitive imaging and spectroscopy applications.

A second main type of field aberration in MRI is inhomogeneity of the static magnetic field, which arises mostly from imperfections of the main magnet and susceptibility differences among different tissue types and air. Static field inhomogeneity is typically counteracted by a combination of passive and active shimming. The latter, relying on variable currents in a set of shim coils (3), is usually adjusted for each subject and scan volume. For further refinement, shimming can be optimized for sub-volumes (4) or on a slice-by-slice basis (5–9). This option, however, requires switching of shim fields at the sequence level and thus on the millisecond scale. In present-day MR systems, higher-order shim coils are typically not designed for such dynamic use. In particular, they are rarely actively shielded, which would come at significant expense in terms of bore space and efficiency. Switching of unshielded shims gives rise to large eddy currents in nearby conducting structures, which impede dynamic operation by settling times up to the order of a second. In addition, similar to gradients the dynamic operation of shims is impaired by vibrations and mutual coupling.

For both gradients and shims, field response imperfections can be reduced by pre-emphasis, i.e., by pre-distorting input waveforms by suitable filters, under the assumption that the system is linear and time-invariant (LTI). Pre-emphasis filters have traditionally been based on modeling eddy current effects as a set of exponentially decaying response components (10–13). However, such models cover only part of the actual degrees of freedom of system responses. Eddy currents flow in a number of structures admitting a continuum of current patterns that can only be approximated by a discrete set of exponentials. Oscillatory responses of mechanical origin are not captured at all in this way and amplifier characteristics, too, do not necessarily lend themselves to compact exponential expansions. The system complexity is multiplied by cross-responses between spatial field components, which again reflect diverse physics and

whose number increases rapidly with that of the spatial terms of interest. Additionally, unlike per-channel pre-emphasis, cross-term correction is an entangled problem since any inputs added for cross-term compensation cause secondary cross-term perturbations of their own.

The purpose of the present work is to explore a more complete pre-emphasis approach for accurate dynamic field actuation. We propose to base pre-emphasis on a system perspective that encompasses more complex temporal response behavior as well as the multiple-input, multiple-output (MIMO) nature of gradient and shim systems. To this end we rely solely on the LTI assumption, according to which a system's behavior is fully described by a matrix of impulse response functions (13). Impulse response characterization of gradient chains has been accomplished using phantom measurements (14) as well as arrays of NMR field probes (15). Using the latter approach, LTI characterization has recently been expanded to combined gradient and higher-order shim systems and their cross-terms (16). Starting from such a system representation, here we report pre-emphasis by inversion of the full LTI model (17,18). The feasibility of this approach is demonstrated with a 3rd-order shim system, achieving bandwidths of 1-3 kHz and thus the ability to switch shim settings in 0.3-1 ms. As one prospective application we explore slice-wise shim updating for echo-planar imaging at 7T, both in a phantom and in the brain *in vivo*.

Theory

LTI system model

The basic LTI system model is described in Appendix A, in its single-channel and MIMO forms. In short, a gradient and/or shim system is modeled as a matrix of transfer functions, $\mathbf{H}(\omega)$, such that

$$\mathbf{O}(\omega) = \mathbf{H}(\omega)\mathbf{I}(\omega), \quad [1]$$

where ω denotes angular frequency and the matrix element $H_{k,l}(\omega)$ relates input in the l -th gradient or shim channel, $I_l(\omega)$, to output in the k -th spatial field component, $O_k(\omega)$. The field components should be chosen such as to span the diversity of fields that the system generates. Most commonly, gradient and shim fields are designed to approximate sets of spherical harmonics, which then form a natural choice of output expansion. Alternatively, the field patterns that gradients and shims actually produce, e.g., in stationary operation, can be determined by measurement and used to form a tailored output basis. Both approaches naturally link each input channel with one output field pattern that is commonly called the respective self-term. When considering only one channel at a time, Eq. [1] reduces to its scalar form [A.2].

Single-channel pre-emphasis

Pre-emphasis aims to counteract the system-induced waveform distortion by pre-compensating the demand spectrum, $I'(\omega)$, through a filter, $H_{PE}(\omega)$, to yield the input to the respective gradient/shim channel, $I(\omega)$.

$$I(\omega) = H_{PE}(\omega) I'(\omega) \quad [2]$$

Ideally, pre-emphasis should result in unit net response across the demand bandwidth, which would be achieved with $H_{PE}(\omega) = H(\omega)^{-1}$. However, such full compensation fails when exceedingly low values of $H(\omega)$ cause it to boost input waveforms beyond amplifier specifications. A trade-off between response behavior and amplifier requirements is struck by defining a target transfer function, $H_T(\omega)$, that draws only on the range of frequencies in which $H(\omega)$ is substantially non-zero. The corresponding pre-emphasis filter is then given by

$$H_{PE}(\omega) = \frac{H_T(\omega)}{H(\omega)}. \quad [3]$$

MIMO pre-emphasis

This approach can be readily expanded to describe pre-emphasis for multiple actuation channels and their cross-terms. Using the matrix-vector notation of Eq. [1], MIMO pre-emphasis amounts to passing a demand vector, $\mathbf{I}'(\omega)$, through a filter matrix, $\mathbf{H}_{PE}(\omega)$, yielding the pre-compensated input vector

$$\mathbf{I}(\omega) = \mathbf{H}_{PE}(\omega) \mathbf{I}'(\omega). \quad [4]$$

The individual entries of the pre-emphasis matrix, $H_{PE_{lm}}(\omega)$, are filters that translate the m -th component of the demand vector into input to the l -th channel. In analogy with the single-channel case, the choice of $\mathbf{H}_{PE}(\omega)$ can be guided by a target transfer matrix, $\mathbf{H}_T(\omega)$, and should then fulfill or approximate

$$\mathbf{H}(\omega) \mathbf{H}_{PE}(\omega) = \mathbf{H}_T(\omega). \quad [5]$$

If $\mathbf{H}(\omega)$ is square and has full rank an exact solution is given by

$$\mathbf{H}_{PE}(\omega) = \mathbf{H}^{-1}(\omega) \mathbf{H}_T(\omega). \quad [6]$$

With channel numbering such that self-term responses appear along the diagonal, one natural choice of $\mathbf{H}_T(\omega)$ is the identity matrix, \mathbf{Id} , multiplied by a scalar target filter $H_T(\omega)$:

$$\mathbf{H}_T(\omega) = H_T(\omega) \mathbf{Id}. \quad [7]$$

The choice of the scalar target filter is subject to essentially the same trade-offs as in the single-channel case. However, physical bandwidth limitations frequently differ between channels. Thus, alternative to

Eq. [7], it may be desirable to use different target filters for different channels. This is formally straightforward by different entries along the diagonal of $\mathbf{H}_T(\omega)$ but requires care not to exceed hardware specifications by cross-term pre-emphasis. A conservative choice is to constrain the target bandwidth of all channels to the narrowest native bandwidth, however at the potential expense of slowing down some channels more than necessary. The assumption of $\mathbf{H}(\omega)$ being square and invertible will frequently hold but not in all relevant scenarios. A discussion of MIMO pre-emphasis for more general $\mathbf{H}(\omega)$ is provided in Appendix B.

Methods

Shim control

The proposed pre-emphasis approach was implemented on a whole-body 7T Philips Achieva system (Philips Healthcare, Best, The Netherlands) equipped with full 3rd-order spherical harmonic shim coils, including a 0th-order B_0 coil. Control of the 0th-, 2nd-, and 3rd-order shim channels was realized via a set of 16-bit digital-to-analogue converters (DAC) with a sampling rate of 25 kS/s and a range of ± 10 V, (NI 9264, cDAQ-9188, National Instruments, Austin, USA), connected to analogue inputs of the shim amplifiers. First-order shim field actuation was performed via the MR system's gradient channels, using custom-built analogue voltage summation units that add shim input waveforms to sequence gradient waveforms from the console. To reserve sufficient amplitude for sequence gradients, 1st-order shim input waveforms were limited to ± 2 mT/m amplitude. They were generated with DACs of higher sampling rate (NI 9263: 100 kS/s, range ± 10 V, 16 bits) to fully cover the larger operational bandwidth of the gradient chains. The DACs were controlled from a PC using LabView (National Instruments).

This setup permitted dynamic operation of all gradient and shim channels independently of the scanner console. Shim input waveforms were hence not subject to the MR system's built-in digital eddy current compensation whereas imaging gradient waveforms generated by the scanner console still underwent default compensation. Where required, the shim waveforms were synchronized with the scanner sequence timing via a TTL trigger signal input to the DACs. System characterization was performed via the independent control chain for all gradient and shim channels.

Field measurements

Field measurements in the bore center were performed with field cameras (19,20) comprised of 16 ¹H NMR field probes (21) distributed on the surface of a sphere of 20 cm in diameter. Two field cameras were used interchangeably, one with probes doped to $T_1 \approx T_2 \approx 20$ ms ('long-lived') (Skopec Magnetic Resonance Technologies, Zurich, Switzerland) and the other with field probes doped to $T_1 \approx T_2 \approx 1.5$ ms

(‘short-lived’). With both cameras the probe signals were received via a custom-configured acquisition system (20), based on high-speed analog-to-digital converters and field-programmable gate arrays set up for sampling at a temporal resolution of 1 μ s. Field measurements were synchronized with shim actuation via a TTL trigger, and the internal clocks of the DACs and the acquisition system were locked onto a shared external source to establish a common time base.

Field probe calibration and processing of raw probe signals was performed as described in Ref. (19), yielding field time courses, $B_p(t)$, for each probe p . The measured fields were expanded in a spatial basis, $b_k(\mathbf{r})$, given by the stationary responses of the 16 gradient and shim channels, k , each related to a time-varying coefficient, $u_k(t)$ (22):

$$B(t, \mathbf{r}) = \sum_k u_k(t) b_k(\mathbf{r}). \quad [8]$$

The stationary responses were obtained by measuring the field changes at the probe positions caused by constant input to the respective gradient/shim channel. The spatial structure of each basis function within the shim volume was determined by field mapping in a large spherical phantom. For visualization purposes, the time-varying coefficients, $u_k(t)$, were scaled to the maximum field shift induced by the related basis function within a sphere of 10 cm radius centered at isocenter. This scaling facilitates comparison between different field components and will be indicated by the unit Hz_{max}.

To measure field responses extending longer than the signal life-time of the field probes, an interleaved measurement scheme was implemented (16). During the length of a shim pulse, the probes were excited multiple times at a TR allowing for sufficient relaxation between excitations. To fill the resulting gaps, each shim waveform was played out multiple times with an incremental shift of the recording intervals. Individual signal acquisitions had durations of 1 ms and 10 ms, respectively, with the short- and long-lived probes, at a TR of 10 ms / 100 ms. Thus in both cases 10 repetitions of a shim waveform were required for full coverage.

System characterization

Shim impulse response functions (SIRFs) were determined by measuring the shim field responses to frequency-modulated pulses input to one channel at a time, as detailed in Ref. (16). For the 0th-, 2nd- and 3rd-order channels, two different frequency-modulated pulses were used as inputs: a linear frequency sweep of 0-3 kHz, and a sweep with cubic frequency modulation of 0-30 Hz, each 10 seconds long. The latter pulse was designed to capture low-frequency responses with particularly high sensitivity. For the gradients, a linear frequency sweep covering 0-20 kHz in 10 seconds was used in addition to the 30 Hz and the 3 kHz sweeps. From the input waveforms and recorded outputs, the SIRFs were calculated using least-squares combination of data obtained with different sweep bandwidths, and were subsequently

smoothed with a kernel of frequency-dependent width, varying between <0.1 Hz for the lowest frequencies and 10 Hz for frequencies above 40 Hz. Using interleaved repetitions as described above, each input pulse required about 3 minutes of measurement, permitting full system characterization in approximately two hours.

Figure 1 shows the 30 Hz and the 3 kHz sweeps, in the time (Fig. 1a,b) and frequency domain (Fig. 1c,d) together with measured self-term field responses and the resulting self-term SIRF (Fig. 1e,f) for a sample channel (ZX).

Pre-emphasis implementation

Based on the hardware and field expansion described above, the gradient and shim system as a whole was modeled by a 16×16 matrix of measured transfer functions [Eq. 1].

Two different approaches were used to define target system responses. In the first, aiming to achieve unit frequency response within a specified bandwidth, a target transfer function was defined in the frequency domain as a raised cosine window:

$$H_T(\omega) = \begin{cases} 1, & |\omega| < 2\pi f_c(1 - \beta) \\ \frac{1}{2} + \frac{1}{2} \cos\left(\frac{|\omega| - 2\pi f_c(1 - \beta)}{4f_c\beta}\right), & 2\pi f_c(1 - \beta) < |\omega| < 2\pi f_c(1 + \beta) \\ 0, & |\omega| > 2\pi f_c(1 + \beta) \end{cases} \quad [9]$$

The full-width half-maximum bandwidth (FWHM) of the window is $2f_c$ and β is a factor defining the width of the transition band. In the present work, for all uses of this target function the FWHM was set to 4 kHz and β was set to 0.4, yielding a transition width of 1.6 kHz. As the purpose of this target filter was to achieve a plateau of unit response within a given bandwidth it will be referred to as plateau filter in the following.

In the second approach, a target step response was defined in the time domain as

$$h_T(t) = \frac{1}{2} + \frac{1}{2} \operatorname{erf}(\alpha t), \quad [10]$$

where erf denotes the Gauss error function. The parameter α determines the rise time of the step and was chosen between 2000 s^{-1} and 20000 s^{-1} , corresponding to FWHM bandwidths of 1-10 kHz. The target transfer function was obtained by differentiation of $h_T(t)$, to yield the target impulse response, and subsequent transformation to the frequency domain. This type of filter will hereafter be called a Gaussian filter.

In defining target transfer matrices, the self-term target transfer functions were chosen to be the same for all channels. Allowing for one demand waveform for each of the 16 output field terms resulted

in square target transfer matrices of the form given by Eq. [7]. The pre-emphasis matrix was then calculated with Eq. [6], involving a separate matrix inversion for each sampled frequency.

Pre-emphasis was carried out as specified by Eq. [4] for each given demand vector, $\mathbf{I}'(\omega)$. For convenient calculation, the pre-emphasis filters were resampled in the frequency domain to match the sampling pattern of the input. The filters were then transformed to the time domain and clipped to a length of 6 s, corresponding to complete settling of all impulse responses. Input waveforms with pre-emphasis, $i_l(t)$, were obtained by discrete time-domain convolution and summation:

$$i_l(t_u) = \sum_{m=1}^{N_m} \sum_v h_{PE_{l,m}}(t_u - \tau_v) i'_m(\tau_v). \quad [11]$$

where $h_{PE_{l,m}}(t)$ and $i'_m(t)$ are the inverse Fourier transforms of the corresponding elements of $\mathbf{H}_{PE}(\omega)$ and $\mathbf{I}'(\omega)$. The indices u and v count the discrete sampling points of the input and demand waveforms, respectively.

The effect of pre-emphasis was assessed by re-measuring the system transfer matrix both with full MIMO pre-emphasis and, for comparison, with self-term pre-emphasis only. To study effects on switching performance, field dynamics were recorded for shim steps without pre-emphasis, with self-term pre-emphasis and with MIMO pre-emphasis. Full 3rd-order MIMO pre-emphasis was limited to the Gaussian 1 kHz target due to critical lows in the native responses of two 3rd-order shims (Z^2X , Z^2Y).

To validate the proposed mode of pre-emphasis in a typical dynamic shim setting, a sequence of shim updates was carried out at an interval of 100 ms, using actual per-slice shim values taken from a preceding in-vivo brain scan. The shim updates were performed without pre-emphasis, with self-term pre-emphasis only and with full MIMO pre-emphasis at 1 kHz bandwidth. Based on concurrent field recordings, filtered to 1.8 kHz bandwidth for noise reduction, the root-mean-square error (RMSE) between measured and prescribed shim fields was determined. In the RMSE calculation, intervals of ± 1 ms around each shim step were excluded to allow for the nominal rise time.

Imaging

The proposed mode of pre-emphasis was also used for per-slice shim updating in imaging experiments with single-shot EPI readouts. Image data of a phantom and a healthy volunteer was acquired using a 32-channel head coil (Nova Medical, Wilmington, DE). In-vivo imaging was performed in compliance with local ethics guidelines.

Initially, to calculate slice-wise shim settings, B_0 maps were acquired as follows: phantom – 40 transverse slices, FOV 220x220 mm², resolution 3x3x3 mm³, TE 3/4 ms, flip angle 45°, TR 9.3 ms; in-vivo – 50 transverse slices, FOV 240x240 mm², resolution 3.5x3.5x2 mm³, TE 2.96/3.96 ms, flip angle 10°, TR 5.21 ms, 6 averages. For each slice a shim setting was determined by least-squares fitting

the gradient/shim spatial basis to the measured field. The fit encompassed the two adjacent slices on either side and included shim amplitude limits as inequality constraints. Only voxels within a mask of the object, determined by thresholding the magnitude image of the field map acquisition, were accounted for. The resulting series of shim steps was then tested for feasibility also with pre-emphasis (1 kHz). The latter occasionally boosted input waveforms beyond hardware limits. In these cases, the calculation of shim settings was iterated, gradually tightening the amplitude constraints for limiting channels, until pre-emphasis was feasible throughout.

Single-shot EPI data were collected from the same slices, with timing and resolution typical of BOLD fMRI: phantom – resolution $3 \times 3 \times 3 \text{ mm}^3$, FOV $220 \times 220 \text{ mm}^2$, TE 25 ms, TR 3 s, flip angle 70° , BW/pixel 32 Hz; in vivo – resolution $1.25 \times 1.25 \times 2 \text{ mm}^3$, FOV $190 \times 230 \text{ mm}^2$, TE 25 ms, TR 2.6 s, flip angle 70° , SENSE 4, BW/pixel 30.6 Hz. The EPI protocol was performed i) without shim, ii) with a global shim (in vivo only), iii) with per-slice shim updating without pre-emphasis, iv) with self-term pre-emphasis (phantom only) and v) with full MIMO pre-emphasis. Shim updating was performed 5 ms before slice excitation.

Results

Self-term pre-emphasis

Demonstrations of self-term pre-emphasis are shown in Figs. 2 and 3, using the XY and Z channels, respectively, as examples. The native transfer function of the XY channel exhibits a narrow peak at DC due to long-lived eddy currents and a more complex resonance pattern around 1 kHz (Fig. 2a-c). The pre-emphasis filter, obtained with the plateau target, compensates for reduced response at low frequencies and excess response in the high-frequency range and has suitable structure around 1 kHz to equalize the resonance. The phase of the filter is the negative of the channel's measured phase. The net transfer function measured with pre-emphasis exhibits the desired flat plateau up to the designed cut-off frequency at 1.2 kHz and zero phase as intended. Residual deviations from the target are up to 0.6% in magnitude and 0.02 rad in phase.

A rectangular shim pulse with and without pre-emphasis is shown in Figure 2d-f. Without pre-emphasis, the field settles slowly due to long-living eddy currents whereas with pre-emphasis it reaches 99% of the targeted value within less than 2 ms (Fig. 2e). After each shim step, oscillations are observed in the native response that correspond to the aforementioned resonance pattern around 1 kHz (Fig. 2f). Pre-emphasis cancels these oscillations along with the eddy current effects.

In the shielded Z gradient channel long-lived eddy currents are very small, making up only about 0.3% of the field response (Fig. 3a-c). Pre-emphasis, targeting a Gaussian transfer function at 10 kHz

bandwidth, further reduced their effect to about 0.06%. Pre-emphasis also eliminated a set of mechanical resonances in the interval of 1.3-2 kHz (Fig. 3c). Due to active shielding the native Z channel implements a rectangular pulse with much greater fidelity than the XY channel (Fig. 3d). Still, detailed views show that pre-emphasis accelerates settling and suppresses oscillations after switching (Fig. 3e-f).

Figure 4 shows self-term pre-emphasis of the XY channel using target system responses of four different bandwidths: a 4 kHz plateau and Gaussian targets of 1, 2, and 3 kHz in width (Fig. 4c,d). The calculated step response of the plateau function had an initial overshoot and slight oscillations after the step. In comparison, the error-function step responses rise more smoothly while settling within a similar time frame despite lower bandwidth. The four target responses resulted in the pre-emphasis filters and step input waveforms shown in Fig. 4e,f. SIRFs and step responses with pre-emphasis were then measured, yielding the results shown in Fig. 4g,h, which resemble the target responses closely.

MIMO pre-emphasis

The joint correction of self- and cross-term responses by MIMO pre-emphasis is illustrated in Fig. 5, using responses to Z2Y input as an example. The native self-term response of this channel is reduced to about 50% by eddy currents and exhibits severe distortion due to mechanical behavior around 1 kHz (Fig. 5a). Self-term pre-emphasis yields a response that closely follows the targeted Gaussian of 1 kHz bandwidth and has the desired flat phase response. In the system used in this work the strongest cross-term of Z2Y was found to be that producing Y field. The native Z2Y→Y response ranges up to 3.3 within the shown bandwidth up to 1.2 kHz (Fig. 5b). This means that upon driving the Z2Y channel at certain frequencies the Y cross-term response is more than three times stronger than the intended Z2Y field (H_{max} scaling). Using self-term pre-emphasis only, the cross-term response is further amplified in parts of the spectrum, to up to a factor of 5. With MIMO pre-emphasis the cross-term response is countered by suitable concurrent input to the Y channel as well as other channels required to suppress cross-terms of Z2Y. As a result the measured Z2Y→Y response with MIMO pre-emphasis is near-zero while the Z2Y self-term response does not differ significantly between the two types of pre-emphasis.

In the time domain, a rectangular Z2Y shim pulse is shown in terms of its self-term (Fig. 5c) and Z2Y→Y cross-term field dynamics (Fig. 5d). With self-term pre-emphasis the measured self-term field time course closely follows the desired waveform. However, self-term treatment alone boosts the Z2Y→Y cross-term response from a factor of up to 2.7 times the targeted self-term step to a maximum factor of 4.9. MIMO pre-emphasis eliminates the cross-term field to a large degree. The small residuum is of very slow dynamics, corresponding to a residual error close to DC that is also apparent in Fig. 5b.

Figure 6 illustrates MIMO pre-emphasis at the system level, showing the entire transfer matrix measured without pre-emphasis (Fig. 6a) and with full MIMO pre-emphasis, targeting Gaussian response

of 1 kHz bandwidth (Fig. 6b). Without pre-emphasis the shape and usable bandwidth of the self-term response varies widely between channels, as do the magnitude and spectral structure of cross-term responses. Virtually all significant off-diagonal entries reflect cross-term output of order lower than that of the respective input channel. This is partly due to the fact that the 0th- and 1st-order coils of the system used are actively shielded while the 2nd and 3rd orders are not. A notable exception is the Z2 response to driving the B0 channel, reaching up to about 1% of this input's self-term response (maximum field within 10 cm radius). The largest cross-terms by far are those of Z2, Z2Y, Z3, and Z2X to B0, Y, Z and X, respectively. For these, the cross-term response exceeds the desired self-term field by up to a factor of 7.6.

Figure 6b shows the result of re-measuring the transfer matrix with MIMO pre-emphasis enabled. Successful pre-emphasis is reflected by the self-term responses consistently showing the targeted spectral shape and all cross-terms being reduced to the noise floor of the measurements involved. Note that the residua reflect noise not only of the measurement with pre-emphasis but also of the initial measurement of the native transfer matrix. The latter propagates with variable scaling according to the targeted transfer gain and the conditioning of the native matrix, which both vary with frequency.

Shim updating

Figure 7 shows a sequence of shim updates without pre-emphasis, with self-term pre-emphasis, and with MIMO pre-emphasis. Without pre-emphasis, shim settling is slow compared to the update rate causing partly large errors in shim levels depending on the history of updates. Self-term pre-emphasis improves the time courses of most 2nd- and 3rd-order shim fields whereas it further deteriorates those of the 0th- and 1st-order fields due to cross-terms. With MIMO pre-emphasis all field time courses approximate the demanded staircase dynamics closely, with switching times of about 1 ms as confirmed in Fig. 4. Noticeable deviations remain in the 0th-order field, however considerably diminished relative to using no pre-emphasis or self-term pre-emphasis alone. For the shim system used here, the step from self-term to MIMO pre-emphasis has the largest impact on the 0th- and 1st-order field terms. Nonetheless several higher-order terms also show significant improvements with MIMO pre-emphasis. Besides removing exponential settling due to eddy currents MIMO pre-emphasis also eliminated oscillations after large shim steps in several field terms.

Table 1 summarizes these results quantitatively in terms of the root-mean-square error (RMSE) between measured and targeted shim field time courses. Without pre-emphasis the RMSE ranges between 2.6 and 139.2 Hz_{max} in different field terms. With self-term pre-emphasis the RMSE is reduced by about one order of magnitude in most 2nd- and 3rd-order terms (0.4-2.9 Hz_{max}) whereas it is increased in all 0th- and 1st-order terms (17.5-158.8 Hz_{max}). With MIMO pre-emphasis the RMSE is consistently low in all

field terms, with a range of 0.2-3.7 Hz_{max}. For all field terms but one the RMSE reduced upon going from self-term to MIMO pre-emphasis.

Imaging

The results of the EPI study in the phantom and in vivo are shown in Figs. 8 and 9, respectively. Single-shot images acquired without shimming exhibit typical distortion relative to the contours of reference spin-warp images (Figs. 8a, 9a). Static global shimming mitigates distortion only somewhat (Fig. 9b) and can even deteriorate field uniformity for individual slices. Without pre-emphasis, slice-wise shim updating is overall detrimental due to the shortcomings of the native dynamic system response (Figs. 8b, 9c). Self-term pre-emphasis alone is similarly ineffective and partly even exacerbates distortions by boosting cross-term output (Fig. 8c). MIMO pre-emphasis, finally, achieves substantially enhanced geometric fidelity by enabling slice-wise shimming with settling in about 1 ms (Figs. 8d, 9d).

Discussion

The pre-emphasis approach explored in this work is based on a full LTI perspective of gradient and shim systems. The LTI picture has long been the basis of pre-emphasis (10–13) and has previously been expanded to incorporate cross-terms (11–13). However, established pre-emphasis implementations rely on describing the system responses by a limited number of exponential terms. Cross-term compensation is typically performed for few selected terms only, which are usually determined in a second iteration after self-term compensation. Relative to traditional pre-emphasis the key advantages of the proposed approach are greater generality regarding temporal response characteristics and complete accounting for cross-term behavior. Measurement and inversion of the full system model permits addressing any system imperfections that are linear and time-invariant without assumptions about underlying mechanisms. Moreover, global system inversion yields self- and cross-term pre-emphasis filters in a single step after one-time system characterization. The report provided here is based on earlier conference proceedings focusing on the broadband LTI approach per se (17) and MIMO considerations (18). In the meantime, pre-emphasis based on a broadband LTI perspective has also, independently, been reported by Goora et al. (23), addressing the case of self-term correction for gradient channels.

MIMO pre-emphasis has here been implemented for a combined gradient and shim system up to 3rd order and has been demonstrated to enable per-slice shim updating without delays for shim settling. Per-slice shimming has the potential to reduce distortion and signal drop-out, especially for techniques that use long readouts and/or gradient echoes with long echo times such as most implementations of fMRI (5,8,24). Other uses of dynamic higher-order fields, such as real-time physiological field compensation (22,25) and higher-order image encoding (26,27), will likely also benefit from higher-fidelity field

dynamics. The arguably broadest area of application, however, is enhanced pre-emphasis of common 1st-order gradient systems. In addition to compensating for exponentially decaying eddy currents, the proposed approach has been found effective at eliminating oscillatory gradient response imperfections of mechanical origin, which have previously been reported to induce artifacts in spectroscopy (28) and phase-sensitive imaging applications (29). One challenge of pre-emphasis for gradient chains, however, is the computational demand of full waveform convolution, which increases with bandwidth.

The LTI assumption underlying pre-emphasis holds to a large degree for typical gradient and shim hardware and the physics involved. Nonetheless, real-world system responses will always exhibit limited degrees of non-linearity due, e.g., to amplifier imperfections or concomitant fields (30,31). Deviations from linear response on the order of a few per mil have been reported for gradient systems (15,30). In a higher-order shim system non-linearity of about one percent has been observed (16). Furthermore, violations of time-invariance occur, for instance, when the system response varies with the thermal state of hardware components (32). Addressing these effects will require extensions to the system model. For slow thermal changes it may be feasible to update pre-emphasis filters based on pre-calibration of the temperature dependence of the system response (33) or by harvesting the response during operation (34). Non-linear responses may be more difficult to address, given the diversity of possible origins and manifestations.

Alternatively to feed-forward pre-emphasis, it has been suggested to improve the gradient system response by real-time feed-back based on field sensors that are sensitive to specific spatial field components (35). This approach is not affected by LTI violations and could also correct for field perturbations originating from sources unrelated to the system (22). However it crucially relies on accurate field sensing and high bandwidth of the feed-back loop, posing considerable technological challenges.

One fundamental limitation of the proposed approach, shared by all prospective field correction strategies, is that cross-term compensation can only be performed for spatial field components that are represented by the available gradient and shim sets. Cross-terms of higher spatial order must be addressed either at the hardware level or retrospectively at the level of data processing if possible.

Pre-emphasis by full model inversion hinges on detailed and accurate knowledge of gradient/shim response behavior, including cross-terms. In this work, system characterization was performed with a field camera based on NMR probes (15,17). Instead, any measurement technique delivering sufficient spatial and temporal information can serve the same purpose. Field camera measurements offer the practical advantage of yielding simultaneous recordings of fields of different spatial distribution at high temporal resolution.

Calculation of the pre-emphasis matrix involves inversion of the transfer matrix $\mathbf{H}(\omega)$. The stability of this operation depends on the magnitude of the singular values of the matrix. Small singular values tend to occur where native transfer functions are small or as a consequence of strong cross-terms. They have two adverse effects. Firstly, they give rise to large entries in $\mathbf{H}^{-1}(\omega)$ and thus in $\mathbf{H}_{PE}(\omega)$, which increase amplifier requirements. Secondly, they cause amplification of noise and any other errors upon matrix inversion, which degrade the accuracy of eventual field dynamics. The generic way of containing these effects is to choose the target filters such as to attenuate or mask critical entries in $\mathbf{H}^{-1}(\omega)$ upon multiplication with $\mathbf{H}_T(\omega)$ according to Eq. [6]. In this work this was achieved by limiting the target bandwidth such as to avoid critical lows of certain transfer functions at higher frequencies. Alternatively, small singular values could be addressed by regularizing the matrix inversion. However, regularization generally introduces alternate systematic error, which ought to be traded off against the problems that it solves. For best results a regularization strategy should hence co-minimize amplifier requirements, propagation of error from the transfer matrix, and the error incurred by the regularization.

Pre-emphasis by a matrix of general filters offers full freedom of choice in the design of the target system response. No single target filter will be optimal for all applications, as demonstrated by comparing the plateau and Gaussian filters used in this work. The plateau target had a better frequency representation and may hence be preferable for waveforms with large higher-frequency content such as EPI readout gradients. Gaussian targets, on the other hand, delivered faster step response settling in the time-domain, despite smaller bandwidth. The choice of target filter thus depends on the intended application and can be determined using general strategies for filter design. While not demonstrated in this work, the target filter may also differ between channels. For instance, different targets could reflect differences in bandwidth requirements, such as between gradients and higher-order shims. A prerequisite for this approach, however, is that cross-term contamination from broader- to narrower-band channels does not exceed the native bandwidths of the latter.

Pre-emphasis often entails substantial amplification of inputs. For unshielded coils, dynamic inputs may need to be a multiple of static inputs for the same field amplitude. Each given system, however, is constrained by amplitude and slew rate limits of the amplifiers involved, resulting in a trade-off between switching speed and amplitude range. This trade-off is less pronounced but also present for shielded 1st-order gradients, especially at large bandwidths. Conventionally, amplifier reserve for pre-emphasis is ensured by permanently reducing amplitude limits for the demand. For unshielded shim channels such general reduction of amplitude range may be unnecessarily limiting, particularly for slow and static operation. Therefore in the present work we employed waveform-dependent iterative adjustment of constraints on demand amplitude.

Conclusion

The fidelity of gradient and shim dynamics can be substantially enhanced by pre-emphasis based on inverting a general LTI system model. This approach has been found effective for addressing both complex temporal response characteristics and multiple-channel cross-term behavior. Deployed with a combined gradient and shim system of full 3rd order it has enabled dynamic shim updating with settling times of about 1 ms. Offering accurate dynamics on this time scale, higher-order shimming promises to reach full versatility at the level of MRI sequence design.

Acknowledgements

Technical and financial support from Philips Healthcare, Best, The Netherlands, is gratefully acknowledged. This work was funded in part by the Swiss Commission for Technology and Innovation (CTI), Grant number 13756.1 PFFLM-NM.

Appendix A

A single-channel LTI system is characterized by its impulse response function, $h(t)$, which governs the relation between an input signal, $i(t)$, and the system output, $o(t)$, by acting as a convolution kernel:

$$o(t) = \int_{-\infty}^{\infty} h(t - \tau) i(\tau) d\tau . \quad [\text{A.1}]$$

In the frequency domain this operation amounts to a multiplication:

$$O(\omega) = H(\omega) I(\omega) , \quad [\text{A.2}]$$

where $O(\omega)$, $I(\omega)$ and $H(\omega)$ denote the Fourier transforms of $o(t)$, $i(t)$, and $h(t)$, respectively. The frequency-domain representation of the impulse response function is termed the transfer function.

For MIMO systems this formalism is expanded by introducing a separate response function for each pair of input and output channels:

$$\begin{aligned} o_{k,l}(t) &= \int_{-\infty}^{\infty} h_{k,l}(t - \tau) i_l(\tau) d\tau \\ &\quad \Downarrow FT \\ O_{k,l}(\omega) &= H_{k,l}(\omega) I_l(\omega) \end{aligned} \quad [\text{A.3}]$$

with the indices k and l counting the output and input channels, respectively. The total output in channel k is thus given by

$$O_k(\omega) = \sum_l H_{k,l}(\omega) I_l(\omega), \quad [\text{A.4}]$$

leading to the matrix-vector notation

$$\mathbf{O}(\omega) = \mathbf{H}(\omega)\mathbf{I}(\omega), \quad [\text{A.5}]$$

with

$$\begin{aligned} \mathbf{O}(\omega) &= \left(O_1(\omega), O_2(\omega), \dots, O_{N_k}(\omega) \right)^T \\ \mathbf{I}(\omega) &= \left(I_1(\omega), I_2(\omega), \dots, I_{N_l}(\omega) \right)^T \end{aligned}$$

and

$$\mathbf{H}(\omega) = \begin{pmatrix} H_{1,1}(\omega) & \cdots & H_{1,N_l}(\omega) \\ \vdots & \ddots & \vdots \\ H_{N_k,1}(\omega) & \cdots & H_{N_k,N_l}(\omega) \end{pmatrix}.$$

N_k and N_l denote the number of output and input channels, respectively. The number of transfer functions required to fully describe the system thus amounts to $N_k \times N_l$.

Appendix B

If the numbers of physical input and output channels do not match or $\mathbf{H}(\omega)$ is rank-deficient optimal pre-emphasis in the least-squares sense is obtained via the Moore-Penrose pseudoinverse:

$$\mathbf{H}_{\text{PE}}(\omega) = \mathbf{H}^+(\omega)\mathbf{H}_T(\omega). \quad [\text{B.1}]$$

When $N_l > N_k$ the system is redundant in that some or all of the output field patterns considered can be produced by multiple channel combinations. One example of this scenario is the presence of both an unshielded and a shielded coil for the same shim term to offer large amplitude for static shimming as well as faster, albeit weaker, response for dynamic use. Upon such redundancy the pseudo-inverse delivers that pre-emphasis which minimizes the l_2 -norm of $\mathbf{I}(\omega)$, thereby also minimizing the sum of the square integrals of the physical input waveforms. When, on the other hand, $N_l < N_k$, the space spanned by the output field terms cannot be fully accessed from the input level. For such a system, pre-emphasis calculated by the pseudo-inverse produces those feasible field dynamics that exhibit the least-squares deviation from the targeted behavior. With rank-deficient $\mathbf{H}(\omega)$ these two situations coincide. Pseudo-inverse pre-emphasis then produces the least-squares set of waveforms that minimize the violation of target behavior.

The pre-emphasis matrix has output dimension N_l , the number of physical input channels. Its input dimension equals the number of demand waveforms that one wishes to consider, N_m . A natural choice is to include one demand waveform for each output field pattern, thus $N_m=N_k$. However, the formalism equally accommodates any other number of demand waveforms whose meaning at the field level is declared through the target transfer matrix $\mathbf{H}_T(\omega)$. An example of practical interest is demand for only gradient fields, yet with cross-term shim pre-emphasis to suppress eddy current and vibration effects of other spatial orders. In this case $\mathbf{H}_{PE}(\omega)$ will be of size $N_l \times 3$, yet equally obtained by Eqs. [6,B.1].

References

1. Mansfield P, Chapman B. Active magnetic screening of gradient coils in NMR imaging. *J Magn Reson* 1986;66:573–576.
2. Carlson JW, Derby KA, Hawryszko KC, Weideman M. Design and evaluation of shielded gradient coils. *Magn Reson Med* 1992;26:191–206.
3. Golay MJE. Field Homogenizing Coils for Nuclear Spin Resonance Instrumentation. *Review of Scientific Instruments* 1958;29:313–315.
4. Poole M, Bowtell R. Volume parcellation for improved dynamic shimming. *Magn Reson Mater Phy* 2008;21:31–40.
5. Blamire A, Rothman D, Nixon T. Dynamic shim updating: A new approach towards optimized whole brain shimming. *Magn Reson Med* 1996;36:159–165.
6. de Graaf RA, Brown PB, McIntyre S, Rothman DL, Nixon TW. Dynamic shim updating (DSU) for multislice signal acquisition. *Magn Reson Med* 2003;49:409–416.
7. Sengupta S, Welch EB, Zhao Y, Foxall D, Starewicz P, Anderson AW, Gore JC, Avison MJ. Dynamic B0 shimming at 7 T. *Magn Reson Imaging* 2011;29:483–496.
8. Koch K, McIntyre S, Nixon T, Rothman D, De Graaf R. Dynamic shim updating on the human brain. *J Magn Reson* 2006;180:286–296.
9. Juchem C, Nixon TW, Diduch P, Rothman DL, Starewicz P, De Graaf RA. Dynamic shimming of the human brain at 7 T. *Concept Magn Reson B* 2010;37B:116–128.
10. Morich MA, Lampman DA, Dannels WR, Goldie FT. Exact temporal eddy current compensation in magnetic resonance imaging systems. *IEEE T Med Imaging* 1988;7:247–254.

11. Jehenson P, Westphal M, Schuff N. Analytical method for the compensation of eddy-current effects induced by pulsed magnetic field gradients in NMR systems. *J Magn Reson* 1990;90:264–278.
12. van Vaals JJ, Bergman AH. Optimization of eddy-current compensation. *J Magn Reson* 1990;90:52–70.
13. Eccles CD, Crozier S, Westphal M, Doddrell DM. Temporal spherical-harmonic expansion and compensation of eddy-current fields produced by gradient pulses. *J Magn Reson Ser A* 1993;103:135–141.
14. Addy NO, Wu HH, Nishimura DG. Simple method for MR gradient system characterization and k-space trajectory estimation. *Magn Reson Med* 2012;68:120–129.
15. Vannesjo SJ, Haeberlin M, Kasper L, Pavan M, Wilm BJ, Barmet C, Pruessmann KP. Gradient system characterization by impulse response measurements with a dynamic field camera. *Magn Reson Med* 2013;69:583–593.
16. Vannesjo SJ, Dietrich BE, Pavan M, Brunner DO, Wilm BJ, Barmet C, Pruessmann KP. Field camera measurements of gradient and shim impulse responses using frequency sweeps. *Magn Reson Med* 2014;72:570–583.
17. Vannesjo SJ, Dietrich BE, Barmet C, Wilm BJ, Brunner DO, Pruessmann KP. Measurement and pre-emphasis of shim responses using frequency sweeps. In: *Proceedings of the 20th Annual Meeting of ISMRM, Melbourne, Australia; 2012.* p 142.
18. Vannesjo SJ, Dietrich BE, Pavan M, Barmet C, Pruessmann KP. Digital Cross-Term Pre-Emphasis for Higher-Order Dynamic Shimming. In: *Proceedings of the 21th Annual Meeting of ISMRM, Salt Lake City, USA; 2013.* p 671.
19. Barmet C, De Zanche N, Pruessmann KP. Spatiotemporal magnetic field monitoring for MR. *Magn Reson Med* 2008;60:187–197.
20. Dietrich BE, Brunner DO, Wilm BJ, Barmet C, Gross S, Kasper L, Haeberlin M, Schmid T, Vannesjo SJ, Pruessmann KP. A field camera for MR sequence monitoring and system analysis. *Magn Reson Med* 2016;75:1831–1840.
21. De Zanche N, Barmet C, Nordmeyer-Massner JA, Pruessmann KP. NMR probes for measuring magnetic fields and field dynamics in MR systems. *Magn Reson Med* 2008;60:176–186.

22. Duerst Y, Wilm BJ, Dietrich BE, Vannesjo SJ, Barmet C, Schmid T, Brunner DO, Pruessmann KP. Real-time feedback for spatiotemporal field stabilization in MR systems. *Magn Reson Med* 2015;73:884–893.
23. Goora FG, Colpitts BG, Balcom BJ. Arbitrary magnetic field gradient waveform correction using an impulse response based pre-equalization technique. *J Magn Reson* 2014;238:70–76.
24. Fillmer A, Vannesjo SJ, Pavan M, Scheidegger M, Pruessmann KP, Henning A. Fast iterative pre-emphasis calibration method enabling third-order dynamic shim updated fMRI. *Magn Reson Med* 2016;75:1119–1131.
25. van Gelderen P, de Zwart JA, Starewicz P, Hinks RS, Duyn JH. Real - time shimming to compensate for respiration - induced B0 fluctuations. *Magn Reson Med* 2007;57:362–368.
26. Hennig J, Welz AM, Schultz G, Korvink J, Liu Z, Speck O, Zaitsev M. Parallel imaging in non-bijective, curvilinear magnetic field gradients: a concept study. *Magn Reson Mater Phy* 2008;21:5–14.
27. Stockmann JP, Ciris PA, Galiana G, Tam L, Constable RT. O-space imaging: Highly efficient parallel imaging using second-order nonlinear fields as encoding gradients with no phase encoding. *Magn Reson Med* 2010;64:447–456.
28. Clayton DB, Elliott MA, Leigh JS, Lenkinski RE. ¹H Spectroscopy without solvent suppression: Characterization of signal modulations at short echo times. *J Magn Reson* 2001;153:203–209.
29. Busch J, Vannesjoe SJ, Giese D, Barmet C, Pruessmann K, Kozerke S. Pre-emphasis Compensation of Oscillatory Phase Offsets in Phase-Contrast Flow Measurements. In: *Proceedings of the 20th Annual Meeting of ISMRM, Melbourne, Australia; 2012.* p 1172.
30. Brodsky EK, Samsonov AA, Block WF. Characterizing and correcting gradient errors in non-cartesian imaging: Are gradient errors linear time-invariant (LTI)? *Magn Reson Med* 2009;62:1466–1476.
31. Bernstein MA, Zhou XJ, Polzin JA, King KF, Ganin A, Pelc NJ, Glover GH. Concomitant gradient terms in phase contrast MR: Analysis and correction. *Magn Reson Med* 1998;39:300–308.

32. Busch J, Vannesjo SJ, Barmet C, Pruessmann KP, Kozerke S. Analysis of temperature dependence of background phase errors in phase-contrast cardiovascular magnetic resonance. *J Cardiovasc Magn Reson* 2014;16:97.
33. Dietrich BE, Reber J, Brunner DO, Wilm BJ, Pruessmann KP. Analysis and Prediction of Gradient Response Functions Under Thermal Load. In: *Proceedings of the 24th Annual Meeting of ISMRM*, Singapore; 2016. p 3551.
34. Wilm BJ, Dietrich BE, Reber J, Vannesjo SJ, Pruessmann KP. Gradient response harvesting for continuous system characterization during MR sequences. In: *Proceedings of the 24th Annual Meeting of ISMRM*, Singapore, Singapore; 2016. p 544.
35. Edler K, Hoult D. Spherical Harmonic Inductive Detection Coils for Dynamic Pre-emphasis. *Magnetic Resonance in Medicine* 2008;60:277–287.

Table 1

	B0	X	Y	Z	XY	ZY	Z2	ZX	X2-Y2
No PE	14.4	139.2	127.0	93.6	4.7	15.5	2.6	30.8	11.0
Self PE	17.5	158.8	150.1	116.8	2.0	2.4	2.9	2.5	1.9
MIMO PE	3.7	1.5	2.7	1.3	0.6	1.1	0.9	0.6	1.8
	Y3	XYZ	Z2Y	Z3	Z2X	Z(X2-Y2)	X3		
No PE	10.2	22.0	10.9	6.9	7.8	16.3	21.6		
Self PE	0.7	0.5	0.7	0.9	0.5	0.9	0.4		
MIMO PE	0.2	0.4	0.3	0.5	0.3	0.6	0.6		

Table 1: Root-mean-square-errors between measured and targeted shim fields without pre-emphasis (no PE), with self-term pre-emphasis (self PE) and with full MIMO pre-emphasis (MIMO PE), for a representative sequence of shim updates. All values are scaled to Hz_{\max} for the respective field term.

Figures

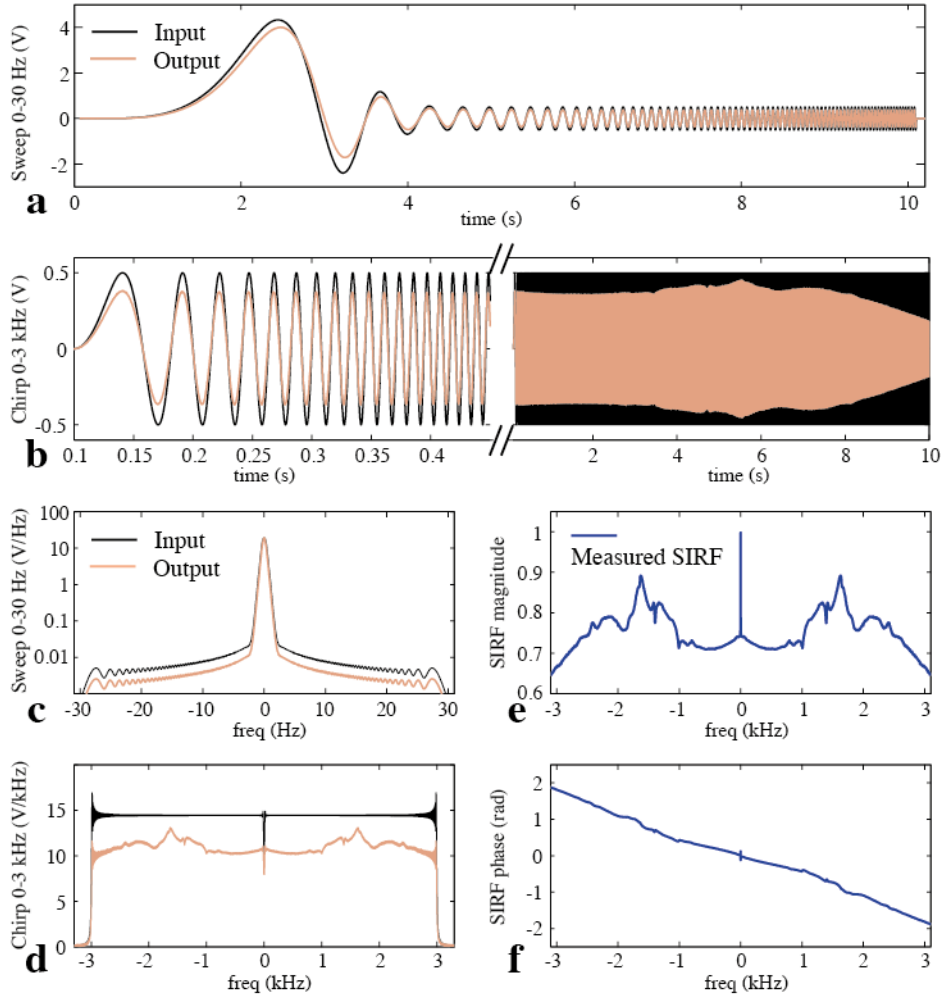


Fig. 1: Frequency-modulated pulses used to determine the transfer functions of each shim channel, shown **a,b**) in the time domain and **c,d**) in the frequency domain. The measured self-term response of the ZX channel is displayed together with the input waveforms. Note the logarithmic scale in **c**. Frequency-domain division yields the measured self-term SIRC for the ZX shim: **e**) magnitude, **f**) phase.

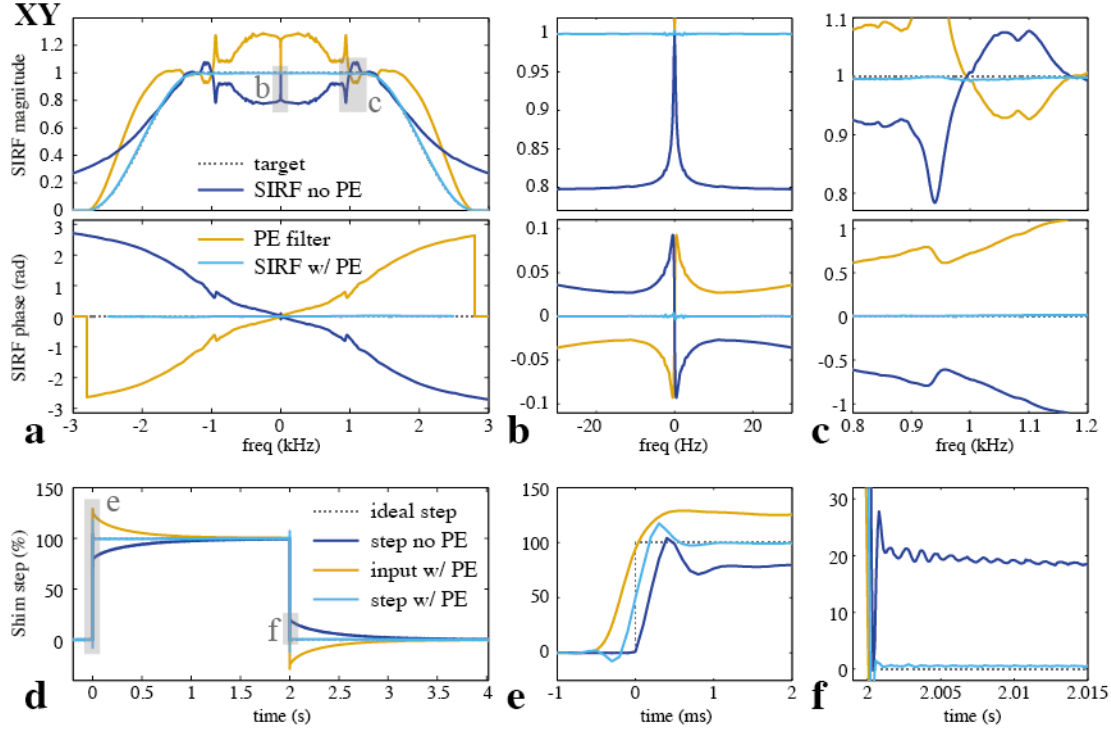


Fig. 2: Measured self-term SIRFs and shim steps for the XY shim channel. **a)** Magnitude and phase of the measured self-term SIRFs without and with pre-emphasis, shown together with the plateau target response and the calculated pre-emphasis filter. The SIRF phase is shown only where the magnitude response is sufficient to yield accurate phase estimates. **b-c)** Magnitude and phase plots showing details of the measured SIRFs around **b)** the eddy-current peak at the center, and **c)** mechanical resonances. **d)** Measured trapezoidal shim pulse with and without pre-emphasis, shown together with an ideal step function and the pre-emphasized input waveform. **e-f)** Details of the trapezoidal shim pulse showing **e)** the initial rising edge and **f)** settling after the end of the pulse. Note the scaling of the frequency axes in kHz (a,c) or Hz (b), and the time axes in seconds (d,f) or milliseconds (e).

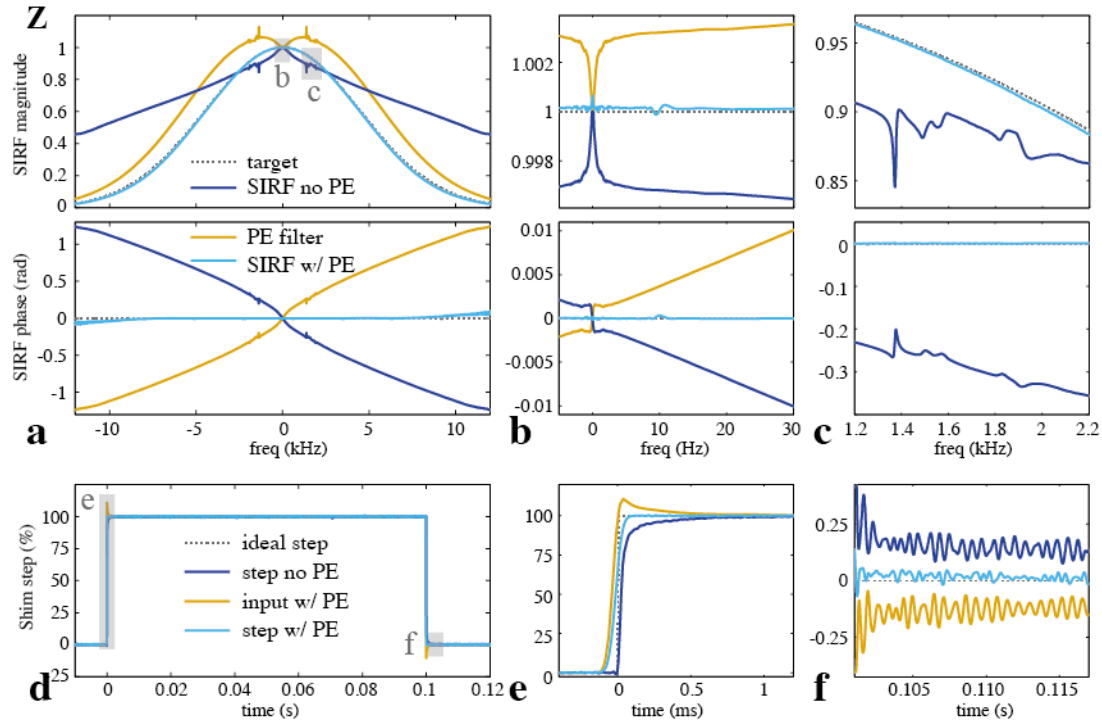


Fig. 3: Measured self-term SIRFs and shim steps for the Z shim channel. **a)** Magnitude and phase of the measured self-term SIRFs without and with pre-emphasis, shown together with the Gaussian target response and the calculated pre-emphasis filter. The SRF phase is shown only where the magnitude response is sufficient to yield accurate phase estimates. **b-c)** Magnitude and phase plots showing details of the measured SIRFs around **b)** the eddy-current peak at the center, and **c)** mechanical resonances. **d)** Measured trapezoidal shim pulse with and without pre-emphasis, shown together with an ideal step function and the pre-emphasized input waveform. **e-f)** Details of the trapezoidal shim pulse showing **e)** the initial rising edge and **f)** settling after the end of the pulse. Note the scaling of the frequency axes in kHz (a,c) or Hz (b), and the time axes in seconds (d,f) or milliseconds (e).

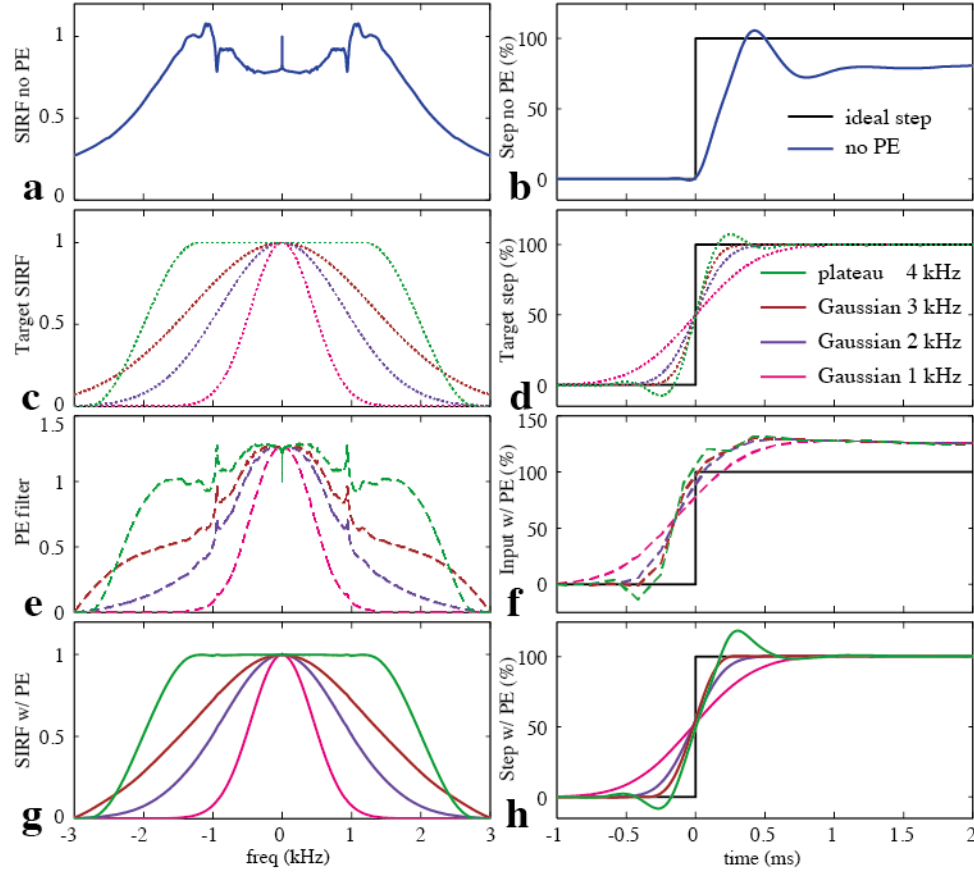


Fig. 4: Measured self-term SIRFs and shim steps for the XY shim channel. Top row: **a)** the native SIRF and **b)** the field response to a step input function. Second row: **c)** different target filters (Gaussian and plateau of bandwidths between 1 kHz and 4 kHz), and **d)** their corresponding step responses. Third row: **e)** calculated pre-emphasis filters and **f)** corresponding filtered step functions. Bottom row: **g)** measured SIRFs and **h)** time-domain shim steps with the respective pre-emphasis.

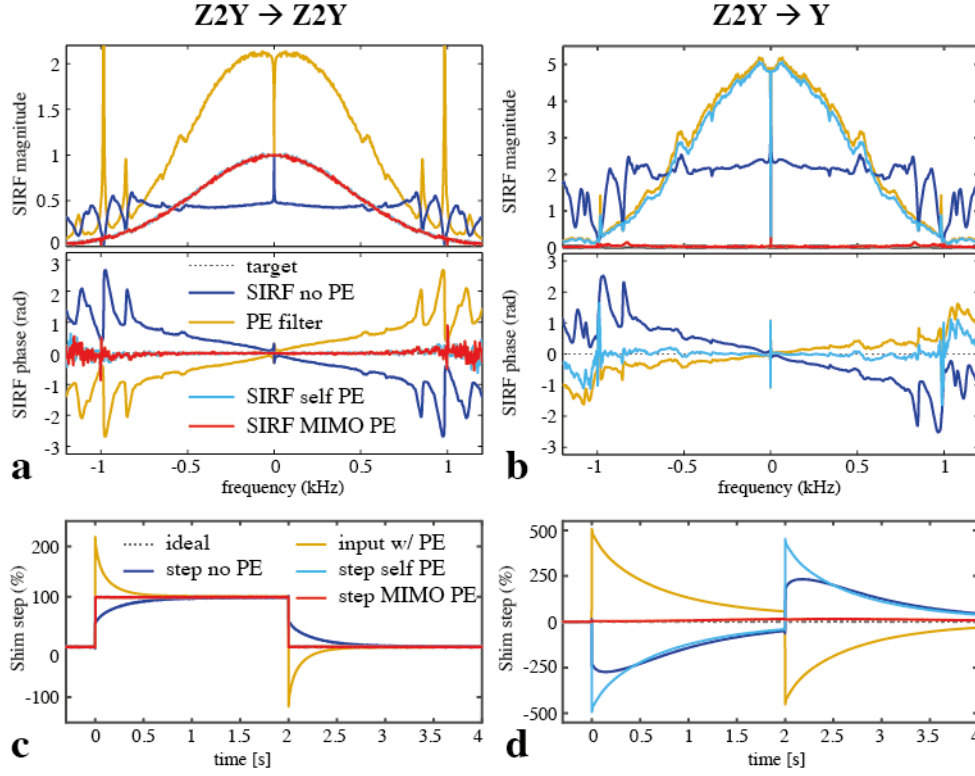


Fig. 5: Measured self- and cross-term SIRFs and shim steps for input to the Z2Y channel. **a)** Magnitude and phase of the self-term SIRF without pre-emphasis, with self-term pre-emphasis only and with MIMO pre-emphasis, shown together with the target response and the MIMO pre-emphasis filter for the self-term. **b)** Magnitude and phase of a cross-term SIRF ($Z2Y \rightarrow Y$) without pre-emphasis, with self-term pre-emphasis only and with MIMO pre-emphasis, shown together with the target response and the respective MIMO pre-emphasis filter. The phase of the cross-term SIRF with MIMO pre-emphasis is not displayed because the magnitude is too low to permit an accurate phase estimate. **c)** Measured self-term of a trapezoidal shim pulse without pre-emphasis, with self-term pre-emphasis and with MIMO pre-emphasis, shown together with an ideal step and the input after pre-emphasis. **d)** Measured cross-term ($Z2Y \rightarrow Y$) of a trapezoidal shim pulse without pre-emphasis, with self-term pre-emphasis and with MIMO pre-emphasis, shown together with the ideal response and the input after MIMO pre-emphasis.

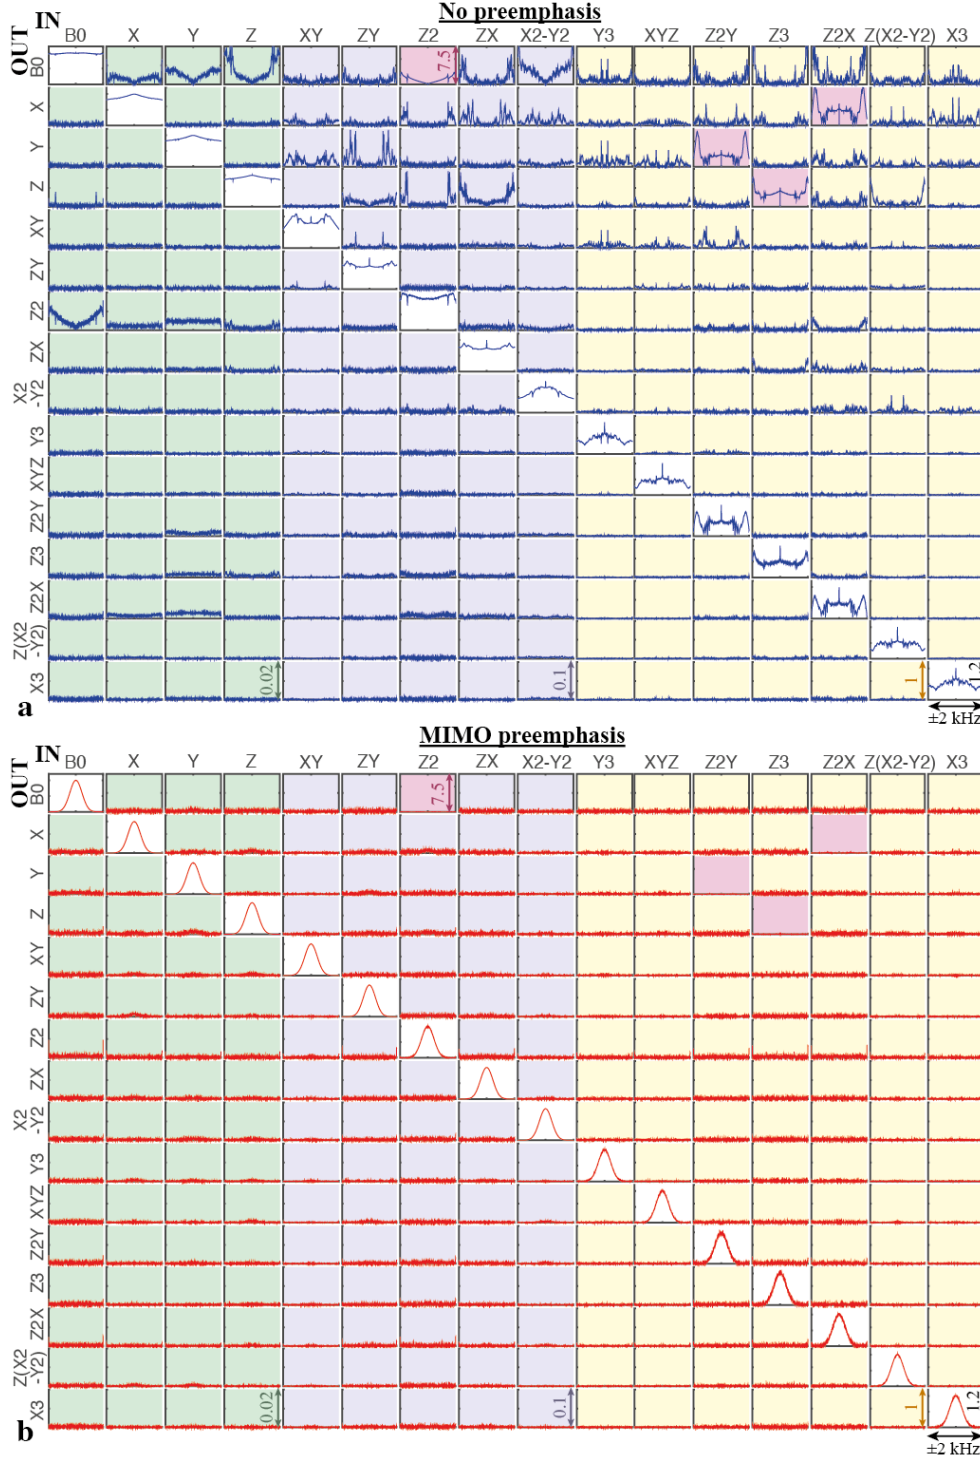


Fig. 6: Measured SIRF matrix, **a)** without pre-emphasis, and **b)** with full MIMO pre-emphasis. The columns represent the driving shim channels and the rows the measured field terms. The horizontal frequency axes cover ± 2 kHz. The vertical range is 0-1.2 for self-terms. For cross-terms it is: 0-0.02 for B0,X,Y,Z input (green); 0-0.1 for 2nd-order input (purple); 0-1 for 3rd-order input (yellow). Four exceptionally strong cross-terms are plotted with range 0-7.5 (magenta).

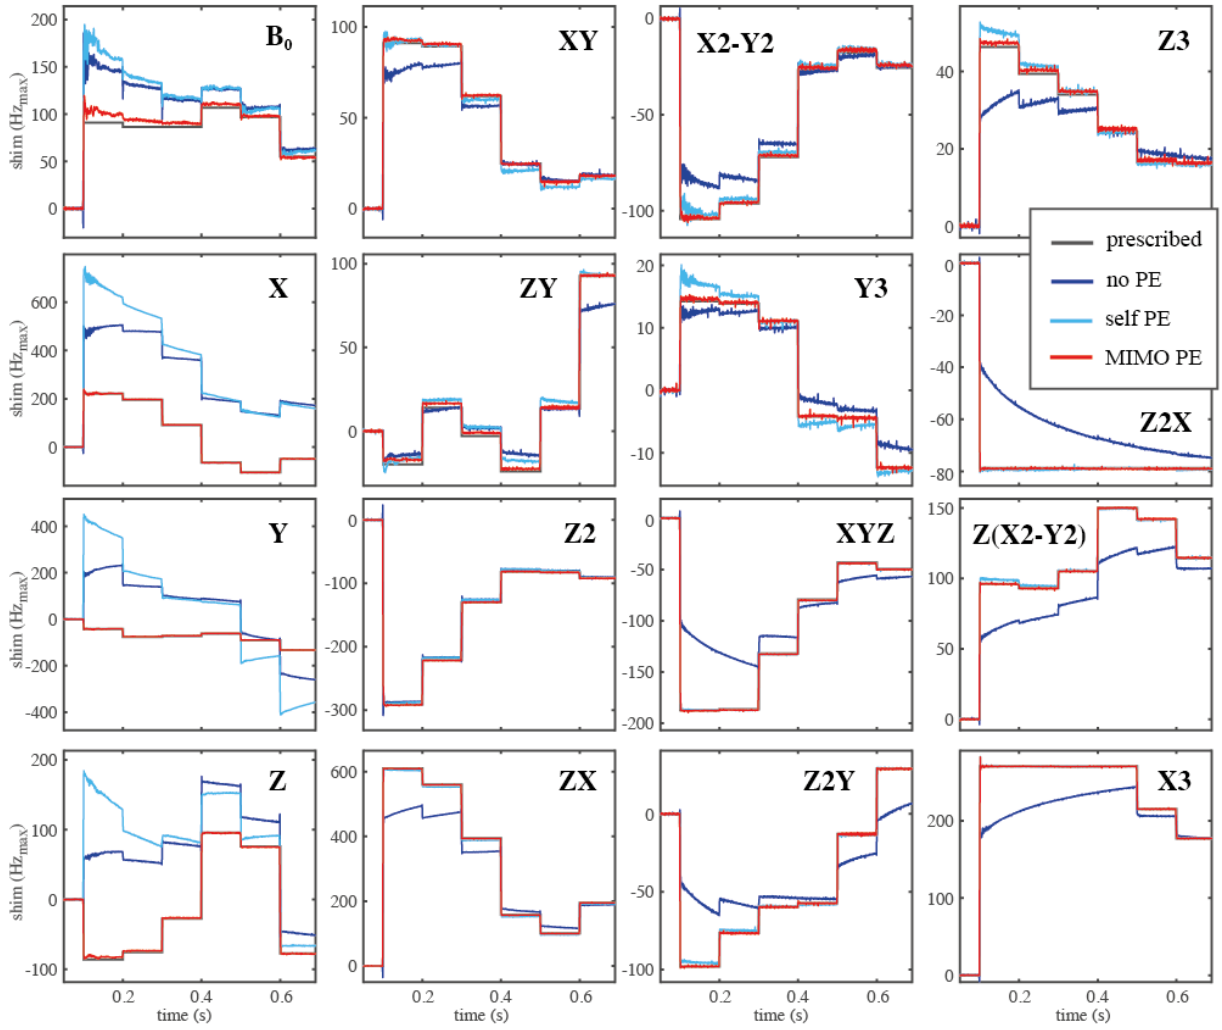


Fig. 7: Measured field output during dynamic shimming (update every 100 ms), performed with no pre-emphasis (no PE), self-term pre-emphasis only (self PE), and full MIMO pre-emphasis (MIMO PE).




Fig. 8: Phantom EPI images. Four selected slices acquired with **a)** no shim, **b)** slice-wise shimming without pre-emphasis (SW, no PE), **c)** slice-wise shimming with self-term pre-emphasis only (SW, self PE) and **d)** slice-wise shimming with full MIMO pre-emphasis (SW, MIMO PE). The red lines indicate edges extracted from a spin-warp image of the same geometry.

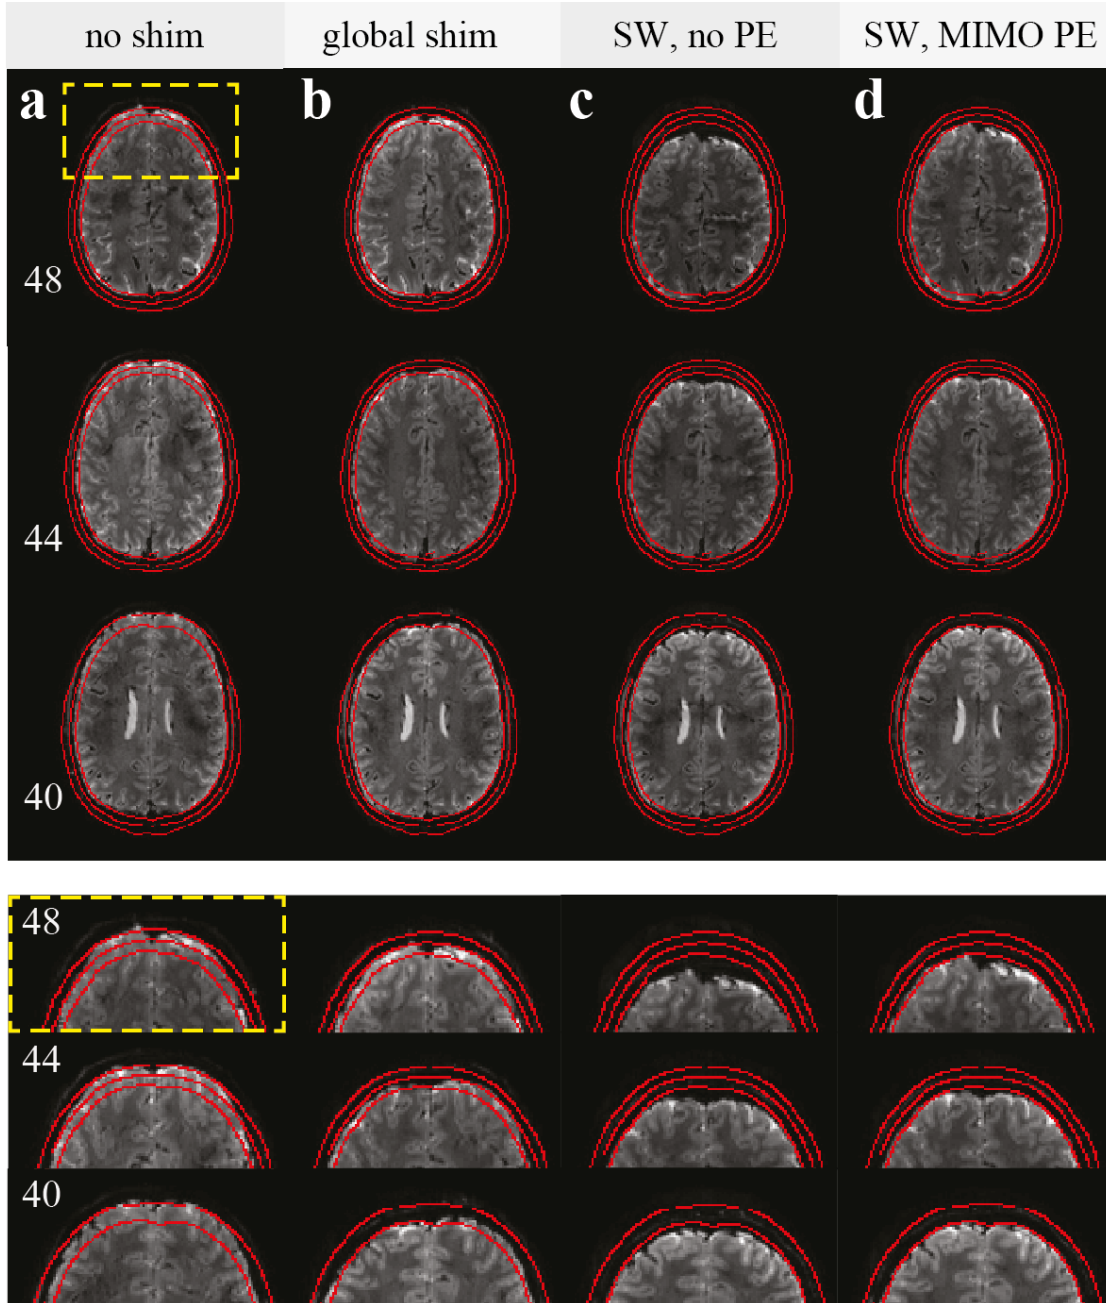


Fig. 9: In-vivo EPI images. Three selected slices acquired with **a)** no shim, **b)** global shim settings, **c)** slice-wise shimming without pre-emphasis (SW, no PE) and **d)** slice-wise shimming with full MIMO pre-emphasis (SW, MIMO PE). The red lines indicate edges extracted from a spin-warp image of the same geometry.