SUPPORTING METHODS

Calculating flip angle progression for FLEET-VFA

The flip angle evolution can be calculated from the Bloch equations and is a function of tissue T_1 value and recovery time τ. If we assume that the flip angle of the last pulse, α_{N_s}, is 90°, then (ignoring T_2^* effects) the expression for the flip angle schedule [McKinnon, 1993] is

$$\tan\left(\alpha_{i-1}\right) = \frac{M_i \sin\left(\alpha_i\right) \exp\left(-\tau/\mathrm{T}_1\right)}{\left[M_i - M_0\left(1 - \exp\left(-\tau/\mathrm{T}_1\right)\right)\right]}$$

where M_0 represents the equilibrium magnetization and M_i represents the magnetization remaining after the i-th segment. Given that the T_1 values in the brain are much longer than the recovery times used in practice for consecutive-segment EPI, we can assume that $T_1 \gg \tau$, which yields a simplified expression [Mansfield, 1984; Chapman *et al.*, 1987; McKinnon, 1993],

$$\alpha_{i-1} = \tan^{-1}\left(\sin\left(\alpha_i\right)\right), \tag{1}$$

with which to calculate the variable flip angle schedule. For example, if the flip angle of the last pulse is 90° (to capture all remaining longitudinal magnetization), the flip angle of the pulse before should be 45°, and the flip angle of the pulse before that should be set to 35.3°. This provides a 45°, 90° sequence (with each shot's signal consisting of 71% of the magnetization) for a two-shot segmented ACS acquisition and a 35.3°, 45°, 90° sequence (with each shot's signal consisting of 60% of the magnetization) for a three-shot segmented ACS acquisition.

Calculating flip angle progression for FLEET

While the Ernst angle (about 8–18° for 3T gray and white matter T_1 values and typical τ values of 20–40 ms) is one candidate for the constant flip angle because it maximizes the signal level in the steady state, typically the steady state will not be reached in FLEET acquisitions with only 2–6 segments. Furthermore, higher signal levels can be achieved in the approach to steady state with slightly larger (*e.g.*, 16–22°) flip angles [Hänicke *et al.*, 1990], however these larger flip angles produce steeper transients and therefore will increase the difference in signal levels across segments. The choice of flip angle is driven by balancing the constraints of having a moderately-high signal level to maintain adequate SNR and producing uniform signal levels across segments so as to not cause artifacts that could interfere with the GRAPPA kernel training. Given the complexity of predicting the tolerance of the GRAPPA training for signal level jumps across segments, in this study we empirically determined several FLEET flip angle values that performed well.

Image evaluation: tSNR and ghost level calculation

Rigid-body motion correction was performed with the AFNI (http://afni.nimh.nih.gov/) command `3dvolreg` [Cox, 1996; Cox & Jesmanowicz, 1999] using the middle time-point as a reference. Time-series SNR (tSNR) was calculated from the motion-corrected data as the temporal mean divided by the temporal standard deviation after linear detrending [Triantafyllou *et al.*, 2011]. In two runs of one 3T session, periods of excessive motion greater than 5 mm displacement were detected and discarded—in one case the last 25% of frames were discarded, in another the first 40% of frames were discarded—but in both cases because of the long duration of the runs there were sufficient numbers of time frames for estimating tSNR.

Ghost levels (*i.e.*, residual aliasing levels) were quantified by calculating the average signal intensity in the ghost region outside of the brain as a percentage of the brain signal intensity. The brain signal was defined as the average intensity within a brain mask defined automatically using the FSL (http://www.fmrib.ox.ac.uk/fsl/) command `bet` [Smith, 2002]. The ghost signal was defined as the average signal intensity within a manually drawn mask outside of the brain along the phase-encoding axis (while avoiding voxels containing chemical shift artifact from the fat layer around the head); a single mask was drawn for each subject based on the images reconstructed from the single-shot EPI ACS data.

GRAPPA g-factor analysis

In addition to SNR maps, geometry factor or “g-factor” maps were computed to characterize the noise enhancement in the accelerated parallel imaging reconstructions [Pruessmann *et al.*, 1999; Breuer *et al.*, 2009]. The g-factor is given by the ratio of the SNR of the accelerated acquisition with the SNR of the unaccelerated acquisition (normalized by the intrinsic $\sqrt{R}$ loss due to the reduced number of samples in the accelerated acquisition). Because of the differential distortions in the accelerated and non-accelerated EPI data, calculating a direct ratio between the corresponding SNR maps is not possible. Therefore g-factor maps were calculated using a Monte Carlo simulation approach [Robson *et al.*, 2008; Triantafyllou *et al.*, 2011]. Briefly, the 10 time-points in the image series identified as exhibiting the least motion relative to the reference time-point (as calculated during the motion correction preprocessing) were averaged together in k-space to derive a high-SNR template representing the accelerated, undersampled EPI data. This k-space template was then replicated 60 times, and to each replicate an independent sample of complex-valued Gaussian white noise with the same channel noise covariance measured from a noise-only (*i.e.*, 0 V RF excitation) thermal noise prescan was added. This synthetic time-series of raw k-space data was then written to a file and fed back into the online image reconstruction system on the scanner. Image SNR was then calculated directly from the reconstructed synthetic data by dividing the mean of each voxel by its standard deviation across the replicates.

GRAPPA artifact map analysis

While the g-factor provides a useful characterization of the noise enhancement resulting from a GRAPPA reconstruction, it does not provide a direct characterization of residual aliasing artifacts that can arise for high acceleration factors. To quantify any residual aliasing artifacts, we employed a previously described procedure for characterizing errors in the GRAPPA reconstruction: first the fully-sampled ACS data were undersampled to emulate the accelerated acquisition, then the GRAPPA reconstructions of the ACS data were compared to the fully sampled ACS data

[Polimeni *et al.*, 2008]. This procedure provides a map of the GRAPPA kernel error across k-space that captures any residual aliasing that cannot be resolved by the GRAPPA reconstruction, which can then be transformed into the image domain to produce an artifact map. If the fitted GRAPPA kernel were perfectly consistent with the ACS data, there would be no coherent discrepancy between the GRAPPA-reconstructed ACS data and the fully-sampled ACS data. Therefore, this “GRAPPA inconsistency” map captures the *inconsistency* between the fitted GRAPPA kernel and the same ACS data used to train it.

Regularizing effect of noise

To test the impact of ACS data SNR on the resulting GRAPPA reconstruction, we conducted two additional tests on 3T phantom data. The data was acquired with the same 3T *R*=4 protocol used in the in vivo experiments. For these tests, only the conventional consecutive-slice ACS data was required. For both tests, the image reconstruction was computed offline in MATLAB.

In the first test, we acquired three sets of consecutive-slice ACS data with different flip angles: 90°, 20°, and 5° (corresponding to the same flip angles used in the FLEET data acquisition). GRAPPA kernels were fit to each of these ACS data sets, and these three GRAPPA kernels were used to reconstruct the *R*=4 accelerated EPI data acquired with a 90° flip angle—that is, the flip angle for the accelerated time-series data was held at 90° so that the only difference between the three sets of reconstructed time-series data was the signal level of the ACS acquisition.

In the second test, the SNR of the consecutive-slice ACS data was also varied, but this time it was controlled by adding noise to the raw k-space data. Specifically, we added complex-valued Gaussian white noise to the ACS data (using a procedure identical to that used for our g-factor Monte Carlo analysis described above) prior to fitting the GRAPPA kernel. In this test, 11 noise levels ranging from 0 to 20 were added to the data. Here the noise level is parameterized by the scaling of the thermal noise that is added to the image data, so if the ACS data naturally exhibited a tSNR of s/σ

then a noise level of 1 would yield a tSNR of *s*/2σ and a noise level of 10 would yield a tSNR of *s*/11σ.

To evaluate the effects of ACS SNR on the GRAPPA reconstructions, we examined the reconstructed image quality, calculated the image tSNR, and computed the GRAPPA inconsistency map to quantify the GRAPPA artifact level. In addition, for the second test we also calculated the GRAPPA kernel condition number, which is directly related to the noise enhancement imparted by the application of the GRAPPA kernel and so is akin to an average GRAPPA g-factor. Together these evaluations provided measures of both artifact level and noise level as a function of the ACS data SNR.

SUPPORTING DISCUSSION

ACS SNR and the regularizing effect of noise on GRAPPA kernel training

Although both the FLASH-based and FLEET-based ACS schemes are "consecutive-segment" approaches, and both eliminate the discontinuous SNR across slices, the acquisitions differ in terms of the SNR of the ACS data. Both the SNR_0 and tSNR are expected to differ between the acquisition schemes. Low-SNR ACS data may affect the GRAPPA kernel training and could lead to artifacts and noise enhancement in the final images. A comparison of the SNR_0 between the ACS schemes tested, for R=2 and R=3 acceleration factors, is presented in Supporting Fig. 3. The consecutive-slice ACS image exhibits the highest SNR_0 of all methods tested (mean SNR_0: 148,137,151 for the R=2, 3, 4 data, respectively), yet the tSNR of the resulting reconstructions using these ACS data is the lowest (see Fig. 4); therefore, as expected, the SNR_0 of the ACS data is not the sole determinant of final image quality. (Also residual ghosting appears in the R=3 consecutive-slice ACS data, presumably due to phase errors between k-space segments, which likely gives rise to the low-quality image reconstructions generated from these ACS data.) The SNR_0 is considerably lower for the FLEET20-ACS data (75,78,81) and lower still for the FLEET5-ACS data (22,24,26), with the FLASH-ACS

SNR_0 (42,44,43) being intermediate between the FLEET-based methods. The relationship between SNR_0 of the ACS data and the resulting accelerated image quality is therefore not straightforward. However, when using either FLASH- or FLEET-based approaches, several protocol parameters for the ACS data can be adjusted independently from the image data (*e.g.*, the flip angle), which provides some flexibility to increase SNR_0 of the ACS data as needed.

Although the reordered acquisition employed in the FLEET, FLEET-VFA, and FLASH data results in a lower SNR_0 compared to the consecutive-slice acquisition, the reordering provides GRAPPA training data that is more robust to motion and dynamic B_0 changes and yields improved tSNR in the accelerated reconstruction. The relationship between the SNR0 of the this training data and the image reconstruction performance of the derived kernels is complicated in part by the regularizing effect that noise can have on the kernel training [Sodickson, 2000]. In cases where the ACS data are not completely consistent with the GRAPPA equations, such that no one kernel accurately captures the relationship between local k-space data in a shift-invariant way across all of k-space, the addition of noise can help to make the system of linear equations formed in the GRAPPA training more consistent, and can act to improve the condition number. This may partly explain the results shown in Fig. 9, where even in phantom data (where no motion or respiration is present) the tSNR achieved when using FLASH or FLEET-ACS data is improved compared to what is achieved when using consecutive-slice ACS data.

To better understand the relationship between ACS signal and noise levels and the resulting tSNR generated by the corresponding GRAPPA kernel, we performed two additional phantom experiments (see Supporting Methods). First, we tested whether simply reducing the flip angle *in a conventional consecutive-slice ACS acquisition* would impact the resulting tSNR in the GRAPPA-reconstructed time-series data; for this test the flip angle for the accelerated time-series data was held at 90° so only the signal level of the ACS acquisition was varied. The results are summarized in Supporting Fig. 4, where it is clear that the artifact level worsens and tSNR improves as the ACS flip angle is reduced, suggesting a regularization effect due to the increased noise level.

Second, we examined the effect of ACS data SNR by incrementally adding noise to the ACS data, and the results are shown in Supporting Fig. 5. In agreement with the results of Supporting Fig. 4, artifact levels worsen and tSNR improves with increasing noise (Supporting Fig. 5A–B). In addition, as the noise level of the ACS data is increased, we observe a marked improvement in the numerical conditioning of the GRAPPA kernel training which leads directly to an improved g-factor (Supporting Fig. 5C). The summary, shown in Supporting Fig. 5D, highlights how this effect of ACS SNR on GRAPPA reconstruction artifact and tSNR represents a classic bias-variance trade-off. In other words, while low SNR in the ACS data can regularize the GRAPPA kernel fit and provide higher tSNR in the GRAPPA-reconstructed time-series, it can also cause increased image artifacts.

There are several implications of this behavior. The lower SNR seen in FLEET (and FLASH) ACS data may contribute to the higher tSNR seen in the corresponding GRAPPA reconstructions, thus the lower flip angles employed in these acquisitions may not actually be a disadvantage—to a point. However, while simply reducing the flip angle on a consecutive-slice ACS acquisition might regain some tSNR, it is expected to increase the artifact levels from the corresponding GRAPPA reconstructions, and our results indicate that the artifact levels seen in the FLEET-ACS-based GRAPPA reconstructions are already as low or lower than those seen in the consecutive-slice ACS-based GRAPPA reconstructions (see Fig. 5). Furthermore, the FLEET-ACS strategy provides critical insensitivity to motion and dynamic B_0 changes during the acquisition.

In summary, the relationship between ACS data SNR and final reconstruction quality is complex, and ACS SNR represents an additional parameter that can be adjusted when optimizing an acquisition protocol to control of the trade-offs between resulting artifact and noise levels.

Removal of discontinuous tSNR through breath-holding

The spatially-varying instability manifested as discontinuous tSNR across slices, was removed when a breath-hold was performed during the ACS acquisition. This observation points to dynamic B_0 changes during the ACS acquisition—which generate phase errors across EPI segments—causing a loss of temporal stability. These phase jumps between segments are presumably imparted on by the B_0 offset driven by the changes in chest cavity volume during respiration, which can reach up to 2.3 Hz at 3T [Van de Moortele *et al.*, 2002]. Not only did the dynamic B_0 changes cause a discontinuous tSNR across slices, but the overall tSNR was seen to be lower in the free-breathing run relative to the breath-hold run shown in Fig. 7, suggesting that (in this example at least) none of the slices were free of phase errors across segments, therefore all slices exhibited reduced tSNR, yet some slices contained more severe phase errors than others. Note that even if the slices were acquired sequentially there would still be a problematic slice-dependent phase error, but the pattern of tSNR losses would be different.

Sensitivity of GRAPPA to phase jumps across segments versus signal-level jumps across segments

The precise dependence of the GRAPPA kernel fitting on the level of phase errors in multi-shot EPI ACS data is unclear. In some cases direct image reconstructions of the consecutive-slice multi-shot EPI ACS data themselves were generated, and a low level of ghosting was apparent (*e.g.*, the R=2 case shown in Supporting Fig. 3), yet a loss of SNR was seen in the accelerated image reconstructions based on this ACS data. It appears therefore that GRAPPA is quite sensitive to these phase errors, and that potentially the requirements for admissible multi-shot EPI ACS data should be stricter than multi-shot EPI data used for other purposes. While some attempt has been made to equalize the signal levels across segments, especially in the case of our FLEET-VFA data (see *Methods*), GRAPPA appears to be more sensitive to ACS phase errors. Further investigations will be required to better understand this sensitivity, and perhaps to adapt the kernel training to be less influenced by phase errors.

Optimal flip angles in FLEET acquisitions used for GRAPPA kernel training

The choice of flip angle for the FLEET-ACS data must also balance between maximizing the signal level and maintaining a uniform signal level across segments. A demonstration of the approach to steady state following a train of constant flip angle excitations is provided in Supporting Fig. 1 assuming the approximate T_1 value of cortical gray matter at 3T (1600 ms) and the recovery time τ for the *R*=4 FLEET-ACS acquisition (19.3 ms). Although the approach to steady state can be controlled by the choice of flip angle, it is unclear how uniform the signal levels must be across segments to provide suitable ACS data for GRAPPA kernel training. The signal levels across the segments can be partially equalized during image reconstruction preprocessing [Kim *et al.*, 1996], yet for most in vivo imaging scenarios there will still be differing levels of T_1 contrast across segments which will be difficult to remove through normalization. Despite these complications, empirically the GRAPPA kernel training appears to have been successful over the range of FLEET flip angles tested. Further work will be necessary to determine the optimal flip angle for FLEET-ACS data used for GRAPPA calibration.

ACS data distortion-matched to the accelerated EPI

The proposed FLEET-based ACS has the same echo-spacing as the effective echo-spacing of the accelerated data, and therefore the GRAPPA kernel is trained to data that have the same geometric distortion as the accelerated image data to which it is applied. Our results indicate that at 7T and in the presence of deliberate B_0 shim offsets at 3T this approach provides higher image quality than approaches where the ACS data have either more (as in the case of the single-shot EPI ACS data) or less (as in the case of the FLASH-ACS data) distortion than the accelerated EPI data. One way to understand the impact of differential distortion between the GRAPPA kernel and the accelerated image data is to consider the k-space representation. A susceptibility gradient alters the spacing of the k-space data (locally increasing or decreasing the image resolution and scale in the phase-encode direction). The size of the k-space spacing change is dependent on the velocity in k-space in the phase-encode direction. If the ACS and image data are not matched in this respect, essentially a kernel is fit that is inappropriate for the image data.

The 7T experiments, summarized in Fig. 8, largely confirm the observations made at 3T. Fig. 8A demonstrates that, for a 3-mm isotropic voxel size, the tSNR from the images reconstructed with consecutive-slice multi-shot EPI ACS data is dramatically lower than the tSNR from images reconstructed with FLEET-ACS data or FLASH-ACS data. The effect of ACS data acquisition scheme appears to be stronger for this voxel size than seen in the 3T results (Fig. 3), perhaps due to the stronger off-resonance effects during respiration at 7T. However for the experiments with higher-resolution data, the discrepancy in tSNR between the various ACS approaches is more subtle. This may be partly caused by increased thermal noise dominance over physiological noise in the small voxels [Triantafyllou *et al.*, 2005, 2011], which may render the ACS data less sensitive to phase errors due to respiration. Furthermore, in the data acquired with the highest resolution tested, 1-mm isotropic, the residual aliasing seen in the images reconstructed with FLASH-ACS data appears to be reduced relative to the images reconstructed with FLEET-ACS data. This could be caused by the reduced geometric distortion of the R=4 accelerated EPI data and the resulting reduced mismatch between it and the effectively undistorted FLASH-ACS data.

While the reconstructions from the FLEET-ACS data perform well in the presence of both naturally-arising and deliberately-induced off-resonance, a high-frequency edge artifact is seen in the reconstructions as shown in the example of Fig. 5. The exact cause of this unresolved aliasing artifact—or the stronger, low-frequency artifact in the reconstructions based on the FLASH-ACS data—is currently not completely understood but is a topic of future investigation.

Acquisition time of ACS data

The conventional consecutive-slice multi-shot ACS approach is made vulnerable to motion artifacts by the time interval imposed between the acquisition adjacent segments within a slice. This vulnerability is increased with the number of slices and with the acceleration factor, such that, for example, a whole-brain high-resolution protocol with R=4 acceleration and a 5 s TR results in a time interval of 15 s between the acquisition of the first and last segments in each slice, making any head movement or breathing during this 15 s period able to corrupt the GRAPPA kernel training. In comparison, when the segments are re-ordered, all ACS segments for a given slice are acquired in less than 0.2 s. Thus the standard segmented ACS acquisition is a particularly vulnerable period. Any movements within this time window create artifacts in the ACS data and infect the entire time-series by repeatedly applying the corrupted GRAPPA kernel.

Table 1 shows that, while the FLASH protocol was optimized to yield short acquisition durations, still the FLEET protocols yield consistently shorter acquisition durations per slice. For FLEET the duration is a function of the number of shots (which is equivalent to the acceleration factor) and the number of ACS preparation pulses (determined based on the ACS flip angle), whereas the FLASH scan is not, so for extreme acceleration factors (*e.g.*, $R \geq 8$) the two will have comparable durations.

Alternate approaches to achieve robust auto-calibration

Other approaches to remedy the vulnerability of multi-shot, segmented EPI to motion and B_0 changes across segments include direct alignment of phase between segments in post-processing, although typically these methods require extra information or additional scans [Hoge *et al.*, 2010; Chen *et al.*, 2013]. Bulk, rigid head motion could potentially be addressed by either image-based [Tisdall *et al.*, 2012] or external sensor-based [Zaitsev *et al.*, 2006; Schulz *et al.*, 2012] motion tracking along with either real-time feedback to the gradient system or a model for how the motion impacts the information acquired across the multiple, interleaved segments. However, independent motion (e.g., of the eyes) during ACS acquisition may not be easily incorporated into this framework. Dynamic B_0 changes in the brain driven by respiration tend to exhibit a spatially inhomogeneous distribution and vary spatially within and across subjects [Van de Moortele *et al.*, 2002; van Gelderen *et al.*, 2007], and therefore are challenging to remove, although dynamic shimming methods attempt to remove these changes [Morrell & Spielman, 1997; van Gelderen *et al.*, 2007; Juchem *et al.*, 2011]. Field probes have been introduced to measure respiratory phase changes which can be recorded for subsequent correction [Vannesjo *et al.*, 2014]. Additionally, sophisticated methods have been developed to remove the two-dimensional pattern of phase offsets that are generated between segments due to B_0 changes [Chen & Wyrwicz, 2004; Xu *et al.*, 2010; Zur, 2011], although simpler strategies employing one-dimensional navigators are known to perform relatively poorly [Barry & Menon, 2005; Barry *et al.*, 2008]. Other potential strategies to reduce ACS corruption include training subjects either to be especially still during the ACS acquisition, including maintaining eye fixation to avoid eye movements, and requesting that subjects perform a breath-hold. These strategies may be limited to healthy, motivated subjects and are prone to break down if the subject is not fully compliant.

REFERENCES

Barry, R. L., Klassen, L. M., Williams, J. M., & Menon, R. S. (2008). Hybrid two-dimensional navigator correction: a new technique to suppress respiratory-induced physiological noise in multi-shot echo-planar functional MRI. *NeuroImage*, *39*(3), 1142–50.

Barry, R. L., & Menon, R. S. (2005). Modeling and suppression of respiration-related physiological noise in echo-planar functional magnetic resonance imaging using global and one-dimensional navigator echo correction. *Magnetic Resonance in Medicine*, *54*(2), 411–8.

Breuer, F. A., Kannengiesser, S. A. R., Blaimer, M., Seiberlich, N., Jakob, P. M., & Griswold, M. A. (2009). General formulation for quantitative G-factor calculation in GRAPPA reconstructions. *Magnetic Resonance in Medicine*, *62*(3), 739–46.

Chapman, B., Turner, R., Ordidge, R. J., Doyle, M., Cawley, M., Coxon, R., Glover, P., & Mansfield, P. (1987). Real-time movie imaging from a single cardiac cycle by NMR. *Magnetic Resonance in Medicine*, *5*(3), 246–54.

Chen, N., & Wyrwicz, A. M. (2004). Removal of EPI Nyquist ghost artifacts with two-dimensional phase correction. *Magnetic Resonance in Medicine*, *51*(6), 1247–53.

Chen, N.-K., Guidon, A., Chang, H.-C., & Song, A. W. (2013). A robust multi-shot scan strategy for high-resolution diffusion weighted MRI enabled by multiplexed sensitivity-encoding (MUSE). *NeuroImage*, *72*, 41–7.

Cox, R. W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, *29*(3), 162–73.

Cox, R. W., & Jesmanowicz, A. (1999). Real-time 3D image registration for functional MRI. *Magnetic Resonance in Medicine*, *42*(6), 1014–8.

Hänicke, W., Merboldt, K. D., Chien, D., Gyngell, M. L., Bruhn, H., & Frahm, J. (1990). Signal strength in subsecond FLASH magnetic resonance imaging: the dynamic approach to steady state. *Medical Physics*, *17*(6), 1004–10.

Hoge, W. S., Tan, H., & Kraft, R. A. (2010). Robust EPI Nyquist ghost elimination via spatial and temporal encoding. *Magnetic Resonance in Medicine*, *64*(6), 1781–91.

Juchem, C., Nixon, T. W., McIntyre, S., Boer, V. O., Rothman, D. L., & de Graaf, R. A. (2011). Dynamic multi-coil shimming of the human brain at 7 T. *Journal of Magnetic Resonance*, *212*(2), 280–8.

Kim, S. G., Hu, X., Adriany, G., & Uğurbil, K. (1996). Fast interleaved echo-planar imaging with navigator: high resolution anatomic and functional images at 4 Tesla. *Magnetic Resonance in Medicine*, *35*(6), 895–902.

Mansfield, P. (1984). Spatial mapping of the chemical shift in NMR. *Magnetic Resonance in Medicine*, *1*(3), 370–86.

McKinnon, G. C. (1993). Ultrafast interleaved gradient-echo-planar imaging on a standard scanner. *Magnetic Resonance in Medicine*, *30*(5), 609–16.

Morrell, G., & Spielman, D. (1997). Dynamic shimming for multi-slice magnetic resonance imaging. *Magnetic Resonance in Medicine*, *38*(3), 477–83.

Polimeni, J. R., Wiggins, G. C., & Wald, L. L. (2008). Characterization of artifacts and noise enhancement introduced by GRAPPA reconstructions. *Proc Intl Soc Mag Reson Med*, *16*, 1286.

Pruessmann, K. P., Weiger, M., Scheidegger, M. B., & Boesiger, P. (1999). SENSE: sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine*, *42*(5), 952–62.

Robson, P. M., Grant, A. K., Madhuranthakam, A. J., Lattanzi, R., Sodickson, D. K., & McKenzie, C. A. (2008). Comprehensive quantification of signal-to-noise ratio and g-factor for image-based and k-space-based parallel imaging reconstructions. *Magnetic Resonance in Medicine*, *60*(4), 895–907.

Schulz, J., Siegert, T., Reimer, E., Labadie, C., Maclaren, J., Herbst, M., Zaitsev, M., & Turner, R. (2012). An embedded optical tracking system for motion-corrected magnetic resonance imaging at 7T. *Magnetic Resonance Materials in Physics, Biology and Medicine*, *25*(6), 443–53.

Smith, S. M. (2002). Fast robust automated brain extraction. *Human Brain Mapping*, *17*(3), 143–55.

Sodickson, D. K. (2000). Tailored SMASH image reconstructions for robust in vivo parallel MR imaging. *Magnetic Resonance in Medicine*, *44*(2), 243–51.

Tisdall, M. D., Hess, A. T., Reuter, M., Meintjes, E. M., Fischl, B., & van der Kouwe, A. J. W. (2012). Volumetric navigators for prospective motion correction and selective reacquisition in neuroanatomical MRI. *Magnetic Resonance in Medicine*, *68*(2), 389–99.

Triantafyllou, C., Hoge, R. D., Krueger, G., Wiggins, C. J., Potthast, A., Wiggins, G. C., & Wald, L. L. (2005). Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters. *NeuroImage*, *26*(1), 243–50.

Triantafyllou, C., Polimeni, J. R., & Wald, L. L. (2011). Physiological noise and signal-to-noise ratio in fMRI with multi-channel array coils. *NeuroImage*, *55*(2), 597–606.

Van de Moortele, P.-F., Pfeuffer, J., Glover, G. H., Ugurbil, K., & Hu, X. (2002). Respiration-induced B0 fluctuations and their spatial distribution in the human brain at 7 Tesla. *Magnetic Resonance in Medicine*, *47*(5), 888–95.

Van Gelderen, P., de Zwart, J. A., Starewicz, P., Hinks, R. S., & Duyn, J. H. (2007). Real-time shimming to compensate for respiration-induced B0 fluctuations. *Magnetic Resonance in Medicine*, *57*(2), 362–8.

Vannesjo, S. J., Wilm, B. J., Duerst, Y., Gross, S., Brunner, D. O., Dietrich, B. E., Schmid, T., Barmet, C., & Pruessmann, K. P. (2014). Retrospective correction of physiological field fluctuations in high-field brain MRI using concurrent field monitoring. *Magnetic Resonance in Medicine*, In press.

Xu, D., King, K. F., Zur, Y., & Hinks, R. S. (2010). Robust 2D phase correction for echo planar imaging under a tight field-of-view. *Magnetic Resonance in Medicine*, *64*(6), 1800–13.

Zaitsev, M., Dold, C., Sakas, G., Hennig, J., & Speck, O. (2006). Magnetic resonance imaging of freely moving objects: prospective real-time motion correction using an external optical motion tracking system. *NeuroImage*, *31*(3), 1038–50.

Zur, Y. (2011). Two-dimensional phase correction method for single and multi-shot echo planar imaging. *Magnetic Resonance in Medicine*.

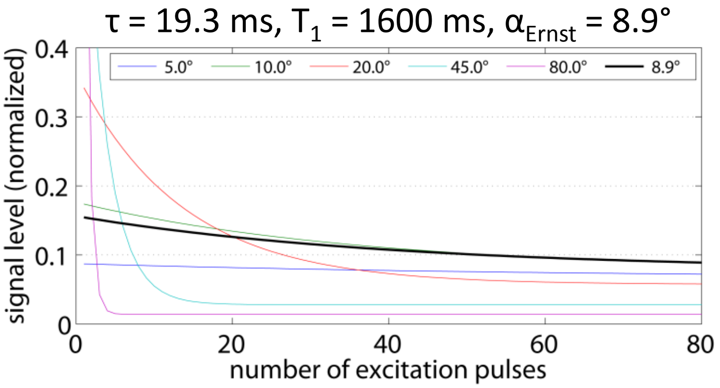

Supporting Fig. 1: Approach to steady-state magnetization with a train of excitation pulses. Transverse signal levels were calculated, using the well-known analytic expression, for a range of constant flip angles assuming the approximate T_1 value of cortical gray matter at 3 T (1600 ms) and the recovery time τ for the *R*=4 FLEET ACS acquisition (19.3 ms). Signal levels are normalized to the equilibrium magnetization and correspond to transverse signal immediately following a series of excitations. Normalized signal level curves are plotted for RF pulse flip angles 5° (blue line), 10° (green line), 20° (red line), 45° (cyan line), and 80° (magenta line), as well as the Ernst angle, 8.9° (dashed black line) for these recovery time and tissue T_1 values. While the highest flip angle of 80° converges after the least number of excitation pulses, the signal level at this steady-state value is the lowest of all flip angles tested. An intermediate flip angle of 20° yields high signal levels after the initial few excitations but in this initial phase the signal level changes steeply after each additional excitation and is slow to converge. The signal level produced by excitations at the Ernst angle (8.9°) converges to the highest signal level of all flip angles tested, however during the initial phase the signal levels change by about 20%. The low flip angle of 5° achieves a relatively flat signal level during the initial few excitations and at a moderately high signal level, and therefore may represent the best choice for maintaining a constant signal over a range of about 5–10 excitations. These curves demonstrate the three-way trade-off between: the speed with which the final steady state can be achieved, the signal level during the transient approach to steady state, and rate of the absolute signal change from one excitation to the next.

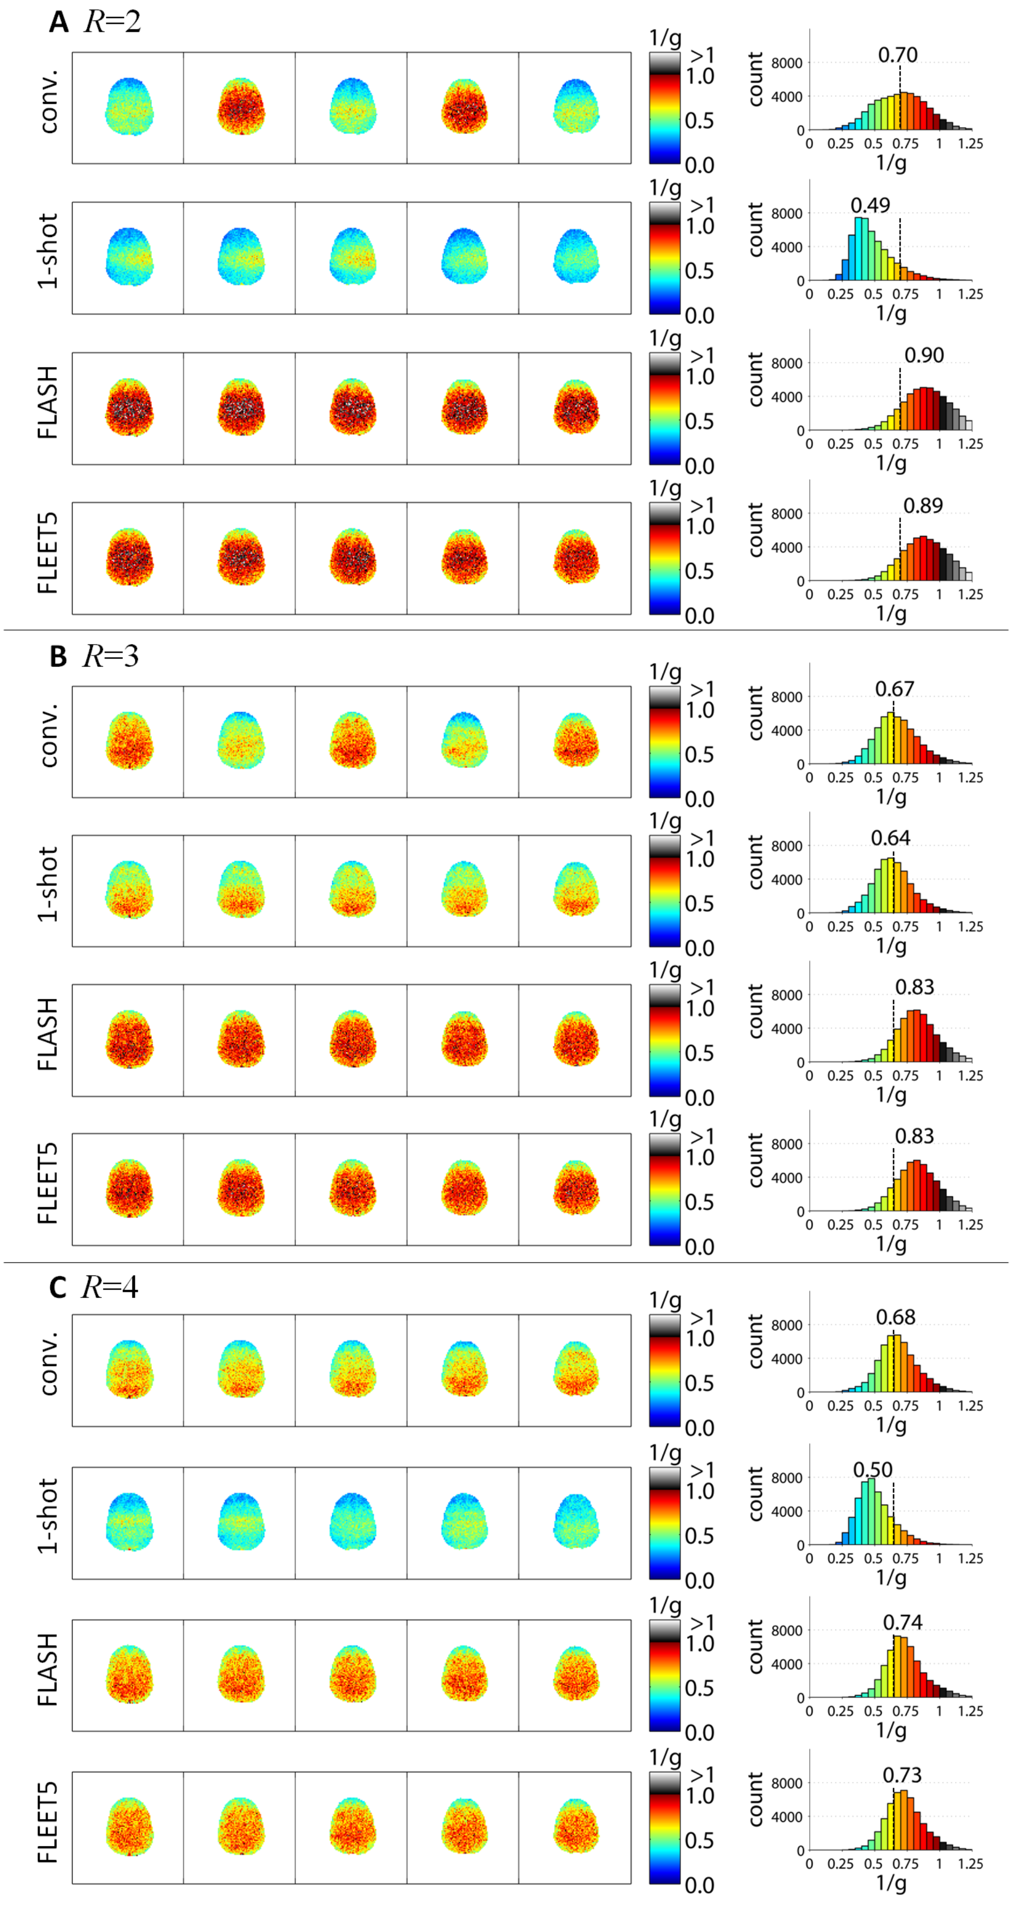

Supporting Fig. 2: Noise enhancement or GRAPPA “g-factor” maps (plotted as 1/*g* for visualization) for one example subject, generated from Monte Carlo analysis and offline image reconstruction. Reciprocal g-factor maps, equivalent to “normalized percent image SNR retained” maps, shown from a set of contiguous slices (LEFT). Map depicts 1/*g* values below 1.0 with a color scale and values above 1.0 with a gray scale to differentiate regions where the g-factor is less than 1 (see text). Note that the frames surrounding each map denote the image Field of View. Corresponding histograms of 1/*g* values, taken from brain mask region extending across all slices in acquisition (RIGHT). Above each histogram is the corresponding mean 1/*g* value from the distribution within the mask. (**A**) Maps from *R*=2 data. (**B**) Maps from *R*=3 data. (**C**) Maps from *R*=4 data.

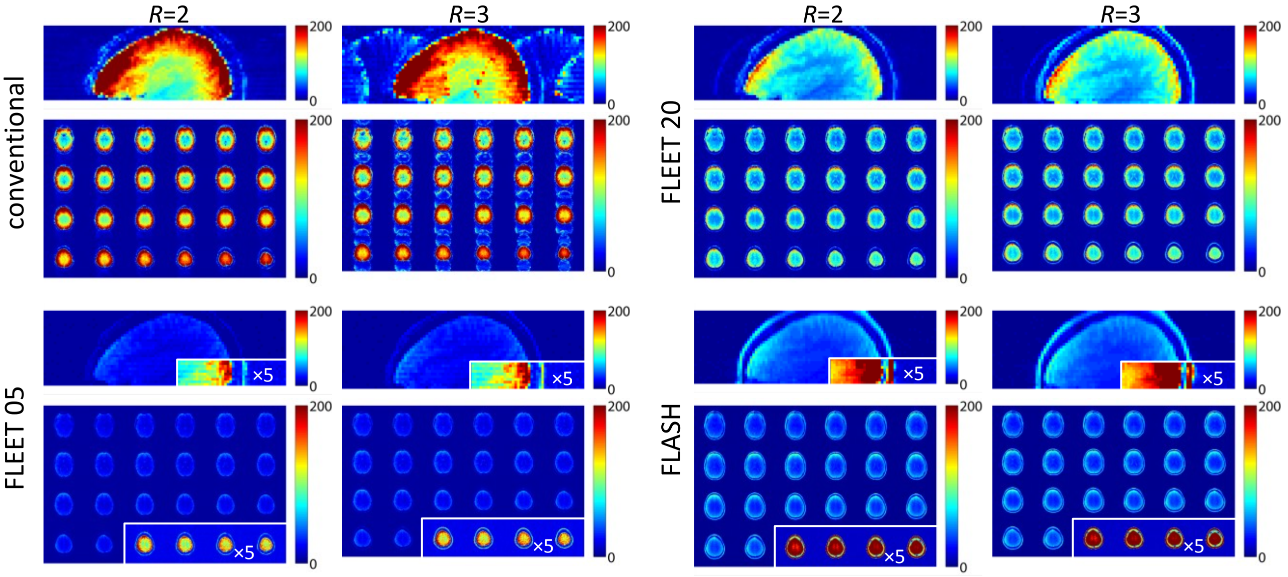

Supporting Fig. 3: Image SNR maps of the ACS data used in the GRAPPA training. The image SNR is highest for the conventional consecutive-slice multi-shot EPI (reconstructed using only linear phase alignment between segments), but ghosting artifacts due to phase errors can be easily seen, and these phase errors are likely the cause of the resulting poor GRAPPA reconstructions. The FLEET EPI images with α_i = 20° exhibits more image SNR than the FLASH images, and FLEET EPI images with α_i = 5° exhibit the lowest image SNR. (Insets show maps scaled by a factor of 5, for better visualization.)

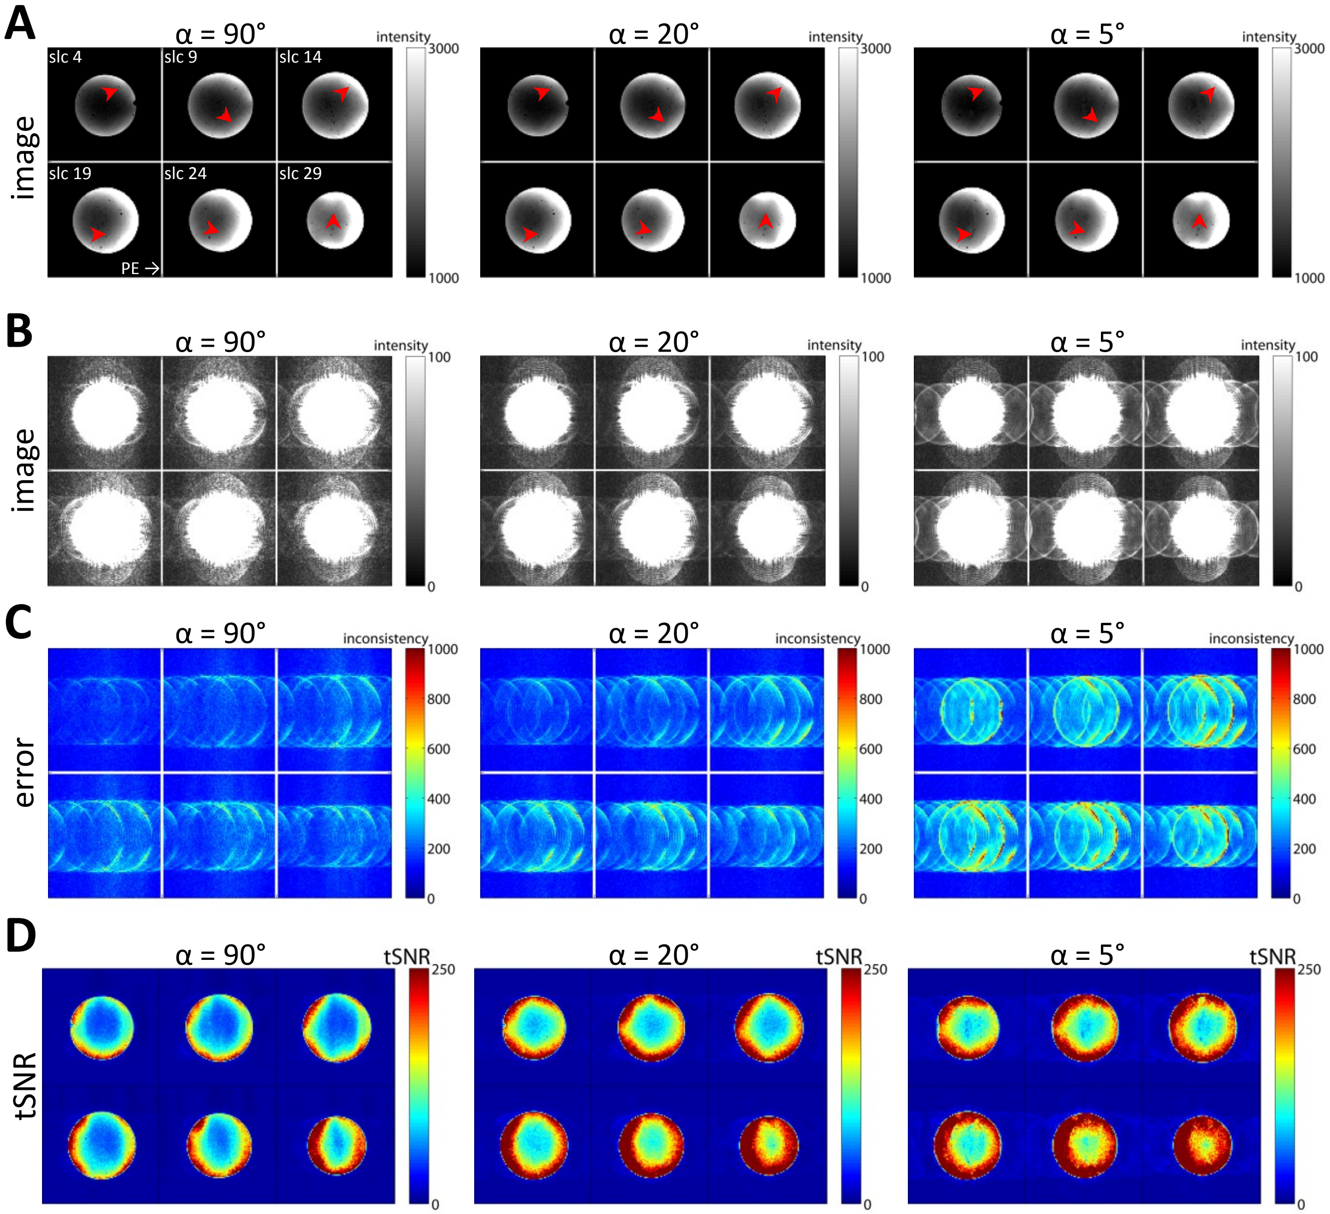

Supporting. Fig 4: Effect of ACS flip angle and signal level on GRAPPA reconstruction image artifact levels demonstrated with an *R* = 4 accelerated agar phantom acquisition, shown across six image slices. (**A**) Image reconstructions using conventional consecutive-slice ACS acquisitions with an ACS flip angle of 90°, 20°, and 5°. The accelerated image acquisition was set to 90° for all reconstructions, only the flip angle for the ACS acquisition was allowed to vary. As the ACS flip angle is decreased, stronger aliasing artifacts (mostly appearing as aliased replicated of phantom edge) appear in image reconstructions. Red arrowheads indicates locations where strong aliasing is apparent in reconstruction based on 5° ACS data but is not apparent in reconstruction based on 90° ACS data. (Slice numbers are and phase-encode direction are indicated in first panel. Slice ordering is consistent throughout figure.) (**B**) Same image reconstructions shown in (A) windowed to highlight artifacts. As the ACS flip angle is decreased, the aliasing artifacts seen in the image background appear to sharpen and become more spatially structured—the aliasing appears to be more structured in the reconstruction using ACS data with a flip angle of 5°, whereas the aliasing artifacts in the reconstruction using ACS data with a flip angle of 90° are more spatially incoherent. (**C**) This trend towards stronger aliasing artifact with decreasing ACS flip angle is verified in the corresponding inconsistency maps. The systematic error resulting from aliasing is strongest and most coherent in the reconstruction using ACS data with a flip angle of 5°.

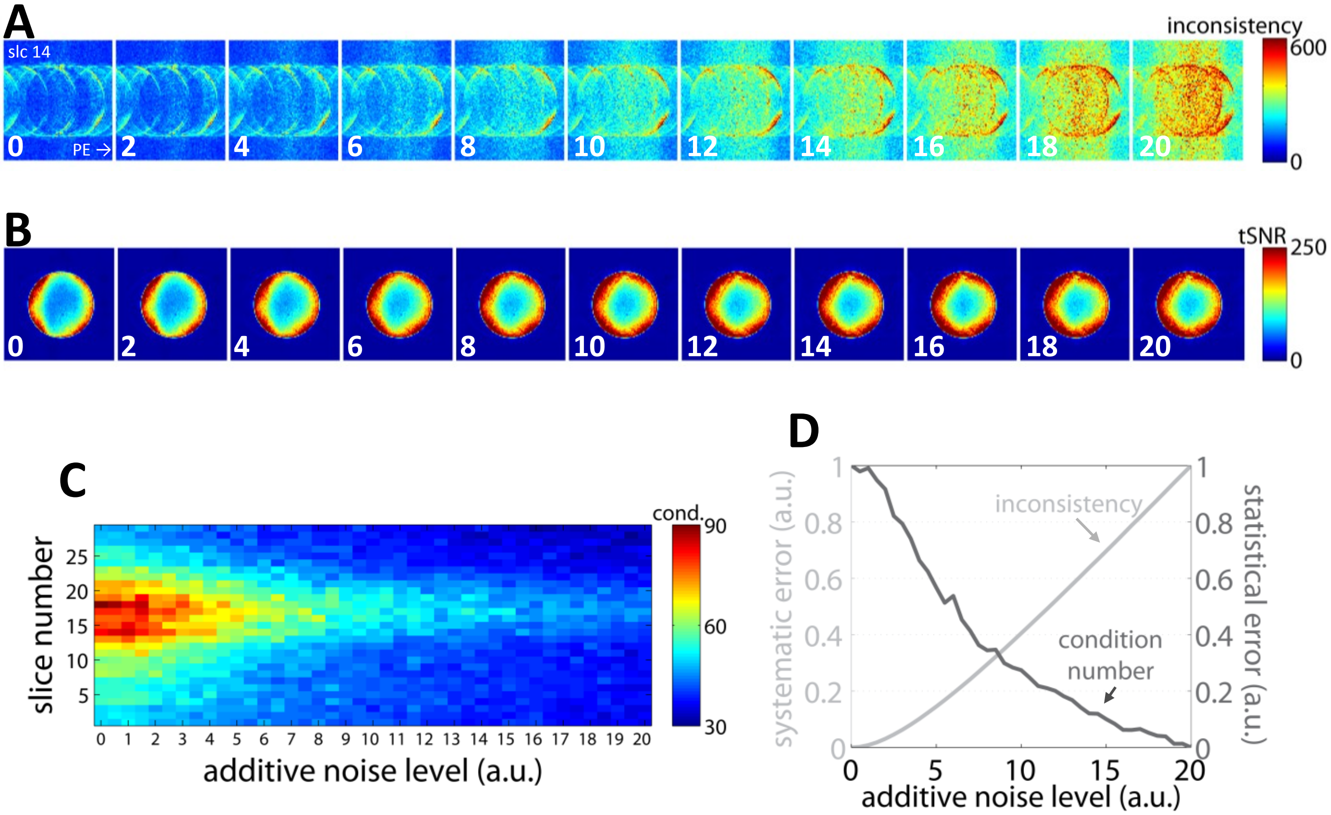

Supporting Fig. 5: Regularizing effect of noise demonstrated by adding increasing levels of noise to conventional consecutive-slice ACS data for an $R = 4$ accelerated agar phantom acquisition. (**A**) GRAPPA inconsistency maps depict increasing levels of systematic error with increasing additive noise levels in the ACS data. Additive noise level is indicated in each panel. (The noise level is parameterized by the scaling of the thermal noise that is added to the image data, so if the ACS data naturally exhibited a tSNR of s/σ then a noise level of 1 would yield an tSNR of $s/2\sigma$ and a noise level of 10 would yield a tSNR of $s/11\sigma$.) Slice number and phase-encode direction are indicated on first panel. The images reconstructed with GRAPPA kernels fit to the noisiest ACS data exhibited the strongest artifacts. (**B**) Corresponding tSNR maps from the resulting GRAPPA-reconstructed time-series data across the additive noise levels. The tSNR clearly increases as noise is added to the ACS data. (**C**) Summary of the effect of additive noise on the conditioning of GRAPPA. The condition number is related to the g-factor, with low numerical conditioning yielding lower g-factors and thus higher tSNR. The condition number varies smoothly across slices but decreases substantially with increasing noise level, indicating that lower noise levels will provide higher tSNR. (**D**) Measured trade-off between artifacts and noise in image reconstruction as a function of noise level. The average inconsistency (calculated from the maps shown in (A)) increases steadily with additive noise level, whereas the condition number decreases. Therefore, while adding noise (or reducing signal) in the ACS data may improve tSNR, it will also lead to an increase in aliasing artifacts.