Enhancement-Constrained Acceleration: A Robust Reconstruction Framework in Breast DCE-MRI

Easley, T.[1], Ren, Z.[2], Kim, B.[3], Karczmar, G.[2], Barber, RF.[3], Pineda, F.[3]

September 28, 2020

1. Washington University in St. Louis, McKelvey School of Engineering
2. University of Chicago, Department of Radiology
3. University of Chicago, Department of Statistics

Abstract

In patients with dense breasts or at high risk of breast cancer, dynamic contrast enhanced MRI (DCE-MRI) is a highly sensitive diagnostic tool. However, its specificity is highly variable and sometimes low; quantitative measurements of contrast uptake parameters may improve specificity and mitigate this issue. To improve diagnostic accuracy, data need to be captured at high spatial and temporal resolution. While many methods exist to accelerate MRI temporal resolution, not all are optimized for the conditions of breast DCE-MRI. We propose a novel, flexible, and powerful framework for the reconstruction of highly-undersampled DCE-MRI data: enhancement-constrained acceleration (ECA). Enhancement-constrained acceleration relies on (a) an assumption of smooth enhancement over small time-scales and (b) somewhat precise knowledge of per-frequency acquisition times. This method is tested in silica with physiologically realistic virtual phantoms, simulating state-of-the-art ultrafast acquisitions at 3.5s temporal resolution reconstructed at 0.25s temporal resolution (demo code available here). Virtual phantoms were developed from real patient data and parametrized in continuous time with arterial input function (AIF) models and lesion enhancement functions. Enhancement-constrained acceleration was compared to standard ultrafast reconstruction in estimating the bolus arrival time and initial slope of enhancement from reconstructed images. We found that the ECA method reconstructed images at 0.25s temporal resolution with no significant loss in image fidelity and a significant reduction in the error of bolus arrival time estimation in both lesions ($p \leq 0.05$) and arteries ($p \leq 0.02$). Our results suggest that ECA is a powerful and versatile tool for breast DCE-MRI.

Introduction

Dynamic contrast enhanced MRI (DCE-MRI) is an important tool for the diagnosis of breast cancer. MRI detects cancers that other screening methods fail to detect. DCE-MRI is particularly important for patients with dense breasts or at high risk for breast cancer. DCE-MRI is highly sensitive (93% [1]) to invasive cancers, and has a variable and sometimes high false positive rate. One 2016 meta-analysis puts the specificity at 71% [1], while another puts it between 78% and 94% [2]; individual studies have reported specificities as low as 37% [3]. These results suggest a need for acquisition and analysis methods that increase the diagnostic accuracy of DCE-MRI. Quantitative measurement of the parameters that describe contrast uptake kinetics offers one route to improved specificity [4], [5], but accurate measurement of these parameters can prove challenging.

Clinical standard-of-care focuses on morphological analysis of DCE-MRI images, as well as evaluation of the overall kinetics of the contrast uptake and washout. As a result, clinical MRI protocols produce post-contrast-injection images at very high spatial resolution; these show patient anatomy in exquisite detail but require long scan times. In the standard-of-care setting, DCE-MRI images are acquired at temporal resolutions of 60-90 seconds. These temporal resolutions are too low to accurately measure kinetic parameters, especially in early uptake, when signal changes rapidly, particularly in cancers. Findings in recent years indicate that lesion conspicuity is highest immediately after contrast uptake [6], [7], so it is especially important to faithfully capture early-uptake kinetics. Other modes of analysis, even primarily morphological ones, also benefit from increased temporal resolution. Some groups have found that texture features, often used to characterize the morphology of lesions and classify them into benign and malignant subgroups, become more accurate with the inclusion of kinetic data [8]–[11]. Thus, high temporal resolution DCE-MRI may offer significant advantages in diagnostic accuracy.

Pharmacokinetic lesion analysis requires an accurate quantitative measurement of the arterial input function (AIF), which requires high temporal resolution measurements. Henderson et al. [12] found that a temporal resolution of 1s is necessary to capture the AIF. Parker et al. [13] opt instead for a population-average AIF, which they calculate from 5s/image data. Estimates of the optimal temporal resolution for pharmacokinetic analysis vary based on the underlying assumptions used to model tissue behavior. For example, Kershaw et al. showed that the standard compartmental (AATH) model requires a temporal resolution of at least 1.5s for accurate diagnosis [14], even when a population AIF is used. In small mammals, which require small fields-of-view, researchers have been able to characterize AIFs with significantly higher sampling frequency; Yankeelov et al. [15] measured an AIF at 0.9s/image in mice, while Kershaw et al. measured an AIF at 0.44s/image in rabbits [16]. Current state-of-the-art in ultrafast breast DCE-MRI produces full 3D bilateral breast scans with temporal resolution 2.7-3.8s [4], [6], [17], [18], well above the desired threshold of temporal resolution.

In fact, the thresholds of 1s and 1.5s offered here represent necessary conditions for only a subset of analytic approaches. Fractional-second temporal resolutions in breast DCE-MRI may allow new modes of kinetic characterization, including detailed local measurements of arterial blood flow and effects of the cardiac cycle, interstitial pressure and vessel permeability, and the initial time and early morphology of lesion enhancement. These parameters have potential as indicators of malignancy [4]. However, these approaches have not been adequately explored, since they require data with high resolution in both the spatial and temporal domains. High spatio-temporal resolution data could offer significant advances in the characterization of tumor physiology and access to biomarkers previously unavailable through established techniques. In order to characterize these vascular properties and assess their utility as malignancy biomarkers, we must first develop and validate methods for the acquisition and reconstruction of high spatio-temporal resolution DCE-MRI data.

Many groups have proposed techniques to increase temporal resolution in MRI, each with their own sets of trade-offs and optimal use cases. While many methods straddle categories, most algorithms fit approximately into one of the groups below:

Ultrafast methods [5], [6], [17], [18] tend to rely on parallel imaging (and partial Fourier sampling) techniques to accelerate scans with greatly reduced coverage and/or spatial resolution. While easy to implement and well-suited to kinetic analysis, the images produced at reduced coverage/resolution do not always contain enough morphological data to be clinically interpretable.

Parallel Imaging techniques [19]–[22] make use of multiple receiver coils to acquire imaging data "in parallel," with data from each coil constraining the image reconstruction. These approaches are powerful and ubiquitous, and they can be used in tandem with many other techniques. However, they suffer from highly nonlinear artifacts at high accelerations. The impact of these artifacts on pharmacokinetic analysis is not well-characterized. They also rely heavily on coil sensitivity maps, which can be difficult to measure precisely.

View-sharing methodologies [6], [23]–[29] accelerate acquisitions by sampling k-space with non-uniform densities, sampling low frequencies much more often than high spatial frequencies. This category includes many common acquisition sequences, including DISCO, TWIST, TRICKS, 4D-TRAK, and most keyhole methods. While these methods do an excellent job of categorizing large-scale enhancement patterns (e.g., average enhancement within a lesion), they sample different spatial frequencies at different temporal resolutions. This could create errors in quantitative analysis. When low spatial frequencies are sampled more often than high ones, it is difficult to reliably interpret the enhancement kinetics of small, sharp structures like blood vessels and the edges of lesions. These structures are critical for accurate clinical diagnosis.

Compressed sensing [30]–[35] strategies capitalize on the sparse enhancement of DCE-MRI in the early uptake phase to create L^1-constrained image reconstructions from very highly undersampled data. This provides high spatial and temporal resolution. However, because these approaches require sampling incoherence, they are susceptible to artifacts from non-uniform k-space sampling. These artifacts are greatest when the signal is changing rapidly, as in the critical phase of early contrast uptake.

Learning-based [10], [36]–[40] reconstruction methods treat image reconstruction from k-space data as a process that can be learned from repeated attempts over large datasets. Especially over the past few years, these methods have become immensely popular for their power and versatility. However, many of these methods do not adequately incorporate the known physics of MRI acquisition or physiological information governing contrast media uptake. Because the reconstruction process is a “black box”, it is difficult characterize or interpret the artifacts introduced by a learned algorithm. Unless a model is specifically trained to reconstruct pharmacokinetic data, it may introduce artifacts that are difficult to account for or even detect.

While all of these methods are powerful, they may not be well-suited to the task of precisely recovering early enhancement kinetics in breast DCE-MRI.

We propose a novel, flexible, and powerful framework for the reconstruction of highly-undersampled DCE-MRI data: enhancement-constrained acceleration (ECA). If raw k-space data is available from the scanner and the acquisition sequence used to produce that data is known, then data can be re-partitioned into almost-arbitrarily small intervals. The data sampled during each intervals can then be used to reconstruct a new set if images with a temporal resolution equal to the interval length. We use the term "sweep time" to refer to the equivalent temporal resolution of a conventional (fully-Nyquist) Fourier-sampled scan at the same spatial resolution as the reconstructed image. In a recent study, bilateral ultrafast scans with complete Fourier sampling had "sweep times" between 3.4 and 4.1s [41]. When the temporal intervals used for image reconstruction are small compared to the "sweep time," this reconstruction problem is (highly) underdetermined. To fully constrain the reconstruction problem, we introduce a penalty function that requires approximately smooth enhancement on the short timescale of the temporal resolution of the reconstructed image. In this setting, our reconstruction problem reduces to the minimization of a constrained quadratic penalty.

Though the work shown here implements a sampling scheme taken from a Cartesian grid, we emphasize that the framework presented is highly general and can be used to invert data from arbitrary sampling schemes. Furthermore, if the time course of the sampling scheme used to acquire the data is either (a) deterministic with known acquisition parameters or (b) recorded as metadata, this framework allows the retroactive reconstruction of existing data.

In the investigation presented here, we (1) develop physiologically realistic breast phantoms from patient data, (2) simulate a virtual scanner with custom pulse sequences to acquire data, and (3) compare the time-tagged reconstructions of phantom data to a "gold standard" conventional Fourier-sampled ultrafast acquisition. To reflect the current state-of-the-art in conventional-Fourier ultrafast acquisitions, we simulate data with a sweep time of 3.5 seconds and reconstruct at a temporal resolution of 0.25 seconds. Code that reproduces some of the examples discussed here is openly available at <github.com/tyo8/ECA_Demo>.

Theory and Methods

Virtual Phantoms

To provide flexible, programmable, and quantitative ground truths for reconstruction testing, virtual breast phantoms modeling realistic pharmacokinetic behavior were created from patient data (**Figure 1**). Five (5) virtual breast phantoms were created from paired ultrafast and high-spatial resolution DCE-MRI datasets representing a range of pathologies (**Table 1**). Acquisition parameters for the original patient scans are shown in **Table 2.**

Fig 1. A maximum-intensity projection of a sample phantom.

The red dashed circle (upper left) shows a sample voxel enhancing via the AIF; the blue dashed circle (middle right) shows a sample voxel enhancing via the EMM.

Table 1: List of pathologies shown in 5 cases selected to use as phantoms

	Pathology
Case 1	Invasive ductal carcinoma (IDC), Grade III
Case 2	IDC, Grade II, ductal carcinoma in situ (DCIS), intermediate and high grade, solid type with necrosis and calcifications
Case 3	Invasive lobular carcinoma, Grade II, lobular carcinoma in situ classic type
Case 4	IDC, Grade III, DCIS, high grade with necrosis and calcifications
Case 5	No abnormal enhancement (control)

Table 2: Scan parameters of the source MR image datasets

	HIGH-SPATIAL RESOLUTION	**ULTRAFAST**
TR/TE (ms)	4.8/2.4	3.2/1.6
Acquisition Voxel Size (mm^3)	$0.8 \times 0.8 \times 1.6$	$1.5 \times 1.5 \times 3$
Temporal Resolution Range (s)	60 - 70	2.8 – 3.6
Flip Angle	10°	10°
Field of View (mm^3)	300-380mm (X,Y),	300-380mm (X,Y),

	180-220mm (Z)	180-220mm (Z)
Number of Slices	120-250	80-100

Virtual phantoms were created as a set of parameters (e.g., bolus arrival time and an uptake rate constant) and functions that call those parameters (e.g., the Parker AIF) to estimate whole-image signal at an input time point. We use these enhancement models to compute the time-evolution of the virtual phantom during the scanning process. The phantoms used in this experiment modeled noisy acquisition with signal from three different classes of sources; vessel, lesion, and background. Vessel and lesion signal components come from distinct models but a single procedure; all signal components are additive.

Vessels and lesions were segmented from patient images and perfusion maps were created from constraining flow equations with ultrafast data, using a method developed by Wu et. al [4]. First, ultrafast and high-spatial resolution image sets undergo non-rigid registration for motion correction. Next, a total 3D vasculature is segmented from high-spatial resolution image sets by a Hessian filtering process. Lesions were segmented by hand. Using the segmented lesions and vessels, ultrafast image sets were used to fit a fluid dynamics-based contrast perfusion model, which created a map of bolus arrival times in the breast. These maps of contrast perfusion times parametrized Parker et al.'s population-based AIF model [13] in vessels and an empirically-derived enhancement model [42] in lesions. Additional parameters of the empirical model were fit from ultrafast image sets.

Background signal in the phantom is static except for fluctuations caused by measurement noise (noise modeling is discussed in the "Virtual Scanner" section). Motion artifacts are not directly simulated in this initial study. Each high-resolution image set contains a single pre-contrast image; these were used to compute the background signal for each phantom.

This breast phantom construction emphasizes the characteristic "sparse plus low-rank" nature of early-enhancement breast DCE-MRI [43], [44] while including features that have high resolution in **both** spatial and temporal domains. In addition, this virtual phantom design is highly modular and can flexibly incorporate staggered changes to both functions and parameter sets; different morphologies, perfusion models, and enhancement functions can be swapped around with relative ease.

Representation of Sampling Schemes

Here we clarify a notion of k-space sampling trajectories, which plays a role in both reconstruction and data acquisition.

A known sampling trajectory can be parametrized as a path k_t through k-space as a function of time t. It specifies which Fourier coefficients are measured and the time at which each that measurement is taken. Though the present work only implements reconstruction on trajectories embedded on a Cartesian grid, any k-space trajectory that can be parametrized as k_t can be reconstructed using the method here; any deterministic path can be simulated. Furthermore, if timestamps from a stochastic measurement are recorded during the scan (or known within the precision of reconstruction time resolution), then these data can be used for enhancement-constrained image reconstruction.

We also define a few parameters that will be useful in describing the sampling and reconstruction processes. By N, we denote the total number of k-space points acquired during a given scan. We use T to represent the number of time-points in an image set. Finally, V is the number of voxels in the image volume at any single time-point.

In the standard IFFT reconstruction of a Nyquist-sampled, Cartesian-acquisition k-space dataset, $N = VT$. On the other hand, when we form an "accelerated" (undersampled) set of images from N measurements in k-space, we have $N < VT$. The acceleration factor is then $\alpha = \frac{VT}{N}$. In this work, we reconstruct dynamic image sets of $VT \sim 10^9$ total voxels from $N \sim 10^8$ measurements (acceleration factor $\alpha = 14$).

Virtual Scanner

A virtual scanner was developed to simulate the acquisition of data from the phantoms described above under varying k-space acquisition paths.

To simulate signal evolution during the acquisition window, phantom signal was recomputed for each k-space measurement $(k_{t_1}, \ldots, k_{t_N})$ made by the scanner. For computational efficiency, however, full re-evaluations of the virtual phantom contrast functions were only made every 50ms of scan time; signal updates between full re-evaluations were computed as linear interpolations between updates. Linearly interpolating between full updates allowed us to capture and quantify mid-scan changes in contrast dynamics, which typically occur at too fine of time-resolutions to be differentiated in standard acquisitions. These are precisely the types of signal changes we are hoping to detect. However, using linear interpolations between function evaluations implicitly encodes the assumption that contrast concentration curves are approximately linear in 50ms windows. and each window has sufficient SNR. While current compartmental, population, and empirical models of enhancement suggest this assumption is reasonable (see **Figure 2** for a visual representation of this assumption on the Parker AIF) [42], [45], [46] and introduces negligible error to the models used in our phantoms, contrast perfusion curves are not sufficiently characterized to fully vet this assumption

Fig 2. Sample concentration curves paired with their 50ms-interval linear approximations.

The AIF curve (used in vessels) is shown on the left and the EMM curve (used in lesion voxels) is shown on the right.

Virtual breast phantoms generated signal data under similar scan parameters as an ultrafast acquisition. A spoiled gradient-echo signal model was used (FA = 10˚, TR/TE = 3.2/1.6 ms, resolution = 1 mm^3) for each phantom acquisition. Noise was modeled as independent Gaussian distributions in k-space with variances computed from pre-contrast ultrafast data. To estimate interference from noise and acquisition artifacts, temporal variances were computed at each k-space point across five (5) pre-contrast ultrafast images for each phantom. Each point's temporal variance parameterized a Gaussian distribution, from which noise realizations were produced. Over all cases and 100 noise realizations, this method of noise generation produced data with an average PSNR of 37dB.

Though the virtual scanner does not have parallel imaging (e.g., SENSE or GRAPPA) [19], [21] or Partial Fourier [47] implementations, path times computed from these scan parameters were appropriately scaled to match standard ultrafast time resolutions. Each scan sequence completed a Nyquist-complete k-space sample of V points every 3.5s and was later reconstructed at a temporal resolution of 0.25s. These k-space acquisition and reconstruction times were chosen to demonstrate the high accelerations this method can achieve (an acceleration factor of $\alpha = \frac{VT}{N} = 14$, from 3.5s to 250ms), while also allowing an investigation of the minimum time resolution necessary to fully resolve enhancement dynamics of clinical interest.

A simple but (to our knowledge) novel k-space trajectory was used to simulate acquisitions in this study (**Figure 3**). By **Un**dersampling **W**ith **R**epeated **A**dvancing **P**hase (UnWRAP), we allow our choice of reconstructed temporal resolution to inform the design of our sampling trajectory. Splitting each group of V acquisitions into f disjoint subsets, we acquire the first line of each subset before moving to the second line of the first subset: we continue in this way until all V acquisitions have been made. This ensures that a uniform distribution of k-space frequency bands determine each reconstructed image, which makes the subsequent reconstruction both (a) robust to noise and (b) sensitive to fast changes in sharp features.

Fig 3. Cross-section of UnWRAP sequence in the $k_y k_z$ plane.

In this scan sequence, k-space is divided into 14 sections, which are each separated into 14 sheaves. Each section must have one sheaf scanned before any section can have another sheaves scanned. This scheme ensures a nearly uniform distribution of high and low spatial frequencies are present in each reconstructed image, while still satisfying the (spatial) Nyquist criterion when all scan data are combined.

The virtual scanner pipeline is summarized in **Figure 4**.

Fig 4. A flowchart summarizing the virtual scanner pipeline.

IFFT Reconstruction

The "standard" IFFT reconstruction is used to benchmark the ECA reconstruction. This "reconstruction" is simply the inverse fast Fourier transform applied to the k-space dataset output by the virtual scanner. This procedure represents the "ideal" ultrafast scan: it assumes a (spatially) Nyquist complete sample is acquired at 3.5s temporal resolution using typical ultrafast acquisition parameters (see **Table 2**). Because ultrafast often relies on other methods (like SENSE and Partial Fourier) to achieve such high temporal resolution [5], [6], [18], we are essentially benchmarking against an assumed perfect SENSE and Partial Fourier reconstruction.

ECA Reconstruction

Our enhancement-constrained acceleration (ECA) reconstruction method penalizes sharp enhancement between reconstructed time-points and requires that new images match the measured k-space data. The formal details of this reconstruction algorithm may be found in **Appendix A** and **Appendix B**, which can be summarized as follows:

1. Partition k-space data into time intervals of equal length; require that each measurement constrains only the image reconstructed in the time interval containing that measurement's time-tag **(Figure 5)**.
2. Denote the reconstructed image by the timeseries $X = (X_1, \ldots, X_T)$ and its spatial Fourier transform as $\tilde{X} = (\tilde{X}_1, \ldots, \tilde{X}_T)$. Define
   (a) A convex, quadratic smoothness penalty on $X = (X_1, \ldots, X_T)$ applied separately to each voxel $v = 1, \ldots, V$ in the spatial domain. When a voxel enhances smoothly, the penalty is small; when it doesn't, the penalty is large. The penalty can be weighted differently on each voxel v, reflecting spatial variation in desired enhancement smoothness.
   (b) A data fidelity constraint on $\tilde{X} = (\tilde{X}_1, \ldots, \tilde{X}_T)$, requiring that, for each t, any subset of $\tilde{X}_t$ measured during time interval t must exactly match the data acquired during that interval.
3. Solve for the image $X = (X_1, \ldots, X_T)$ that minimizes the smoothness penalty and satisfies the data fidelity constraint on $\tilde{X} = (\tilde{X}_1, \ldots, \tilde{X}_T)$.

Fig 5. An illustration of the k-space partitioning process.

(Top) The sampling scheme displayed in Figure 3 is overlaid on sample k-space data. All k-space measurements in the same time interval constrain the reconstruction of a single time-point in the accelerated reconstruction. (Bottom) These k-space points form the "measured" partition of the reconstructed dataset. The temporal resolution of the reconstruction determines the size of the measured partition of k-space points used for each reconstruction. The higher the acceleration factor α, the shorter the duration of each measured partition and the more underdetermined the reconstruction problem.

Intuitively, we can think of this optimization as a search for the smoothest set of enhancement curves that are consistent with our measured k-space data. Formulating the reconstruction in this way relies heavily on one key assumption: enhancement is smooth on the *timescale of the reconstruction's temporal resolution*. Requiring smoothness on short timescales does not limit the ability of this algorithm to accurately measure sharp spatio-temporal changes as in the AIF, since these changes occur on longer timescales. We chose a target temporal resolution of 0.25 seconds in response to speculations in the literature [12-16] about optimal temporal resolutions for pharmacokinetic analysis in breast DCE-MRI.

For a formal description of the partitioning process invoked above, see **Appendix A**. Since the smoothness penalty optimized during reconstruction is a positive-definite quadratic form, the reconstruction optimization is convex and has a unique solution. While this solution can be defined analytically, we calculate it iteratively via conjugate gradient descent. See **Appendix B** and **Appendix C** for further details. Finally, the computation of image updates requires some amount of regularization to converge; for a discussion on choice of regularization parameter, see **Appendix D**. Documentation and demos for the phantom, scanner, and reconstruction pipelines are available at <github.com/tyo8/ECA_Demo>.

Data Analysis

As an initial investigation of the ECA reconstruction framework, we compared images and enhancement curves recovered from ECA and standard IFFT reconstructions. Two parameters, bolus arrival time (BAT) and initial enhancement slope, were extracted from the signal enhancement curve of vessel and lesion voxels by ECA and standard IFFT methods.

BAT was measured from time of peak enhancement in vessel voxels. In lesion voxels, BAT was calculated as the earliest time at which voxels reached or exceeded 20% of their maximum enhancement over baseline.

To calculate initial slope in vessel voxels, each voxel timeseries was interpolated by a modified Akima method [49]. The initial slope was the maximum first derivative of the interpolated AIF curve. In lesion voxels, percent signal enhancement (PSE) versus time was fitted to a piecewise empirical mathematical model (EMM):

$$PSE(t) = \begin{cases} A * \dfrac{\left(\alpha(t-t_0)\right)^2}{1+(\alpha(t-t_0))^2}, & (t \geq t_0) \\ 0, & (t < t_0) \end{cases}$$

where t_0 is the BAT in lesion voxels, A is the upper limit of percent enhancement, and α is the uptake rate; thus, $A\alpha$ is the initial enhancement slope.

To assess image preservation, voxel-wise image fidelity was also compared between the two methods. Ground-truth images are computed by evaluating the signal function at the center of the temporal window surrounding each reconstructed time point. The distribution of absolute voxel-wise signal differences between reconstructed and ground-truth images is then computed and summarized.

Results

Bolus Arrival Time (BAT)

Sample BAT maps, computed from both the IFFT and ECA reconstructions, are shown in **Figure 6**. The images computed from an ECA reconstruction show more accurate and precise bolus arrival time estimate than do the images computed from an IFFT reconstruction (**Figure 7a/b**). Absolute error distributions of BAT estimate from ECA and IFFT reconstructions were compared via a two-sample Kolmogorov-Smirnov test, and BAT estimates were found to be significantly more accurate in ECA reconstructions ($0.01 \leq p \leq 0.02$ in vessels and $0.04 \leq p \leq 0.05$ in lesions). BAT estimation error distributions are shown (for lesion and vessel voxels) for all cases in **Figure 7a/b**. Summary statistics over all cases are shown in **Table 3**.

Fig 6. Bolus arrival time computed from the case 4 image set.

From left to right: IFFT reconstruction, ECA reconstruction, and ground truth. Times shown on the color bar are measured in seconds.

Fig 7. Error distributions are shown for all cases.

Distributions for IFFT and ECA are shown in different colors on the same plot. (a) and (b) show errors in the estimation of the bolus arrival time in milliseconds; (c) and (d) show the distribution of the proportional voxel error.

Table 3. Median absolute error values over all cases.

BAT Error (ms)

	Lesion	Vessel
IFFT	1701 ms	904 ms
ECA	267 ms	64.6 ms

Voxel Error (%)

	Lesion	Vessel
IFFT	0.007%	0.39%
ECA	0.47%	0.56%

BAT differences are shown separated by case in **Figure 8**. While the standard IFT reconstruction showed no substantial difference in bolus arrival time by case, ECA predicted bolus arrival time substantially better in cases 1, 3, and 4 than in cases 2 and 5.

Fig 8. Box and whisker plots of the error in bolus arrival time, by case number and reconstruction method.

Red boxplots (right) are from ECA-reconstructed data; blue boxplots (left) are from standard IFFT reconstructions. Separate sets of plots are shown for lesion and vessel voxels.

Images reconstructed from the ECA algorithm show much greater precision in estimation of bolus arrival time. Errors in BAT have much smaller spread (median absolute deviation) for ECA reconstructions than for IFFT reconstructions, especially in vessel voxels. Furthermore, the BAT error distribution is clustered much nearer to 0 in ECA reconstructions than in IFFT reconstructions (**Figure 7a/b**), especially in vessel voxels.

Overall, ECA reconstructions were much more successful in recovering bolus arrival times than IFFT reconstructions. Because bolus arrival time estimates were more accurate with ECA, we conclude that ECA reconstruction allows for more accurate and more precise bolus tracking than traditional ultrafast methods.

Initial Slope

Enhancement curves recovered from reconstructed images closely match simulated enhancement from the phantoms. **Figure 9** plots ground truth versus estimate values for the initial slope, as derived from both ECA (left) and IFFT (right) reconstructions. Compared to standard IFFT, ECA reconstruction more accurately recovers the initial slope of the enhancement curve in both vessel and lesion voxels. To see this, first note that the coefficient of determination in both sets of truth-estimate fits is larger for ECA than IFFT; therefore, ECA produces lower-variance estimates of initial slope than IFFT does. Next, compare the slopes and offsets of the truth-estimate fits. In vessel voxels, IFFT and ECA have similar fit slopes and offsets, and therefore introduce similar amounts of bias; in lesion voxels, ECA introduces much less bias than IFFT. Although both methods exhibit greater error when recovering enhancement curves with larger initial enhancement slope, ECA estimates the initial slope more accurately than standard IFFT.

Fig 9. Scatter plot between ground truth initial slope and estimated initial slope.

The panels show ECA and standard IFFT in (a)(b) vessels voxels and (c)(d) lesion voxels. The red lines and blue lines represent the linear correlations and black dashed lines show unity.

Image Fidelity

ECA-reconstructed images are highly similar to ground-truth images. A sample error map is overlaid on a phantom in **Figure 10**. Voxel-wise error statistic summaries are shown in **Figure 11** for the 5 phantoms tested, and a sample enhancing-voxel error distribution is shown in **Figure 7c/d**. Voxel-wise errors in the IFFT reconstruction were generally smaller than in the ECA reconstruction, though fidelity errors were very small in both methods. The voxel intensity error distributions shown in **Figure 10** show that the increased temporal resolution comes with at most a negligible cost in voxel intensity accuracy.

Fig 10. Proportional intensity error per voxel for case 4.

Proportional intensity error is shown from the mean projection over time and through the volume. From left to right: IFFT reconstruction error, ECA reconstruction error, and the ground truth image.

Fig 11. Box and whisker plots of the proportional enhancement error, by case and reconstruction method.

Red boxplots (right) are from ECA-reconstructed data; blue boxplots (left) are from standard IFFT reconstructions. Separate plots are shown for lesion and vessel voxels.

Comparison with Standard Methods

Figure 12 juxtaposes a median-quality curve from the ECA reconstruction with the IFFT reconstruction (which assumes perfect SENSE and Partial Fourier reconstructions) of the same voxel. Sample curves are shown for constant, vessel, and lesion voxels.

Fig 12. Comparison of a sample constant-signal from standard IFFT and ECA reconstruction.

The ECA is more sensitive to noise than the IFFT, but the noise is still small with respect to the signal. IFFT reconstruction estimates bolus arrival time and peak signal less accurately than ECA reconstruction.

Overall, the ECA reconstruction captured bolus arrival times in enhancing voxels more accurately than the IFFT reconstruction, suffering only a small loss of accuracy in estimating the per-time point image (**Table 3, Figure 7c/d, Figure 11**). While ECA proved uniformly more accurate in recovering the BAT in the vessel, ECA and IFFT estimated lesion BAT with similar bias in two cases; in the other three, ECA estimated the BAT with lower bias and variance. Even in cases where ECA and IFFT produced similarly biased estimations of BAT, the ECA estimated the BAT with lower variance (**Table 3, Figure 7a/b, Figure 8**).

Discussion

The results from realistic phantoms reported here demonstrate that sparse uniform samples of k-space can be used to reconstruct DCE-MRI breast images with high fidelity and very high temporal resolution. This allows more accurate arterial bolus tracking and more accurate measurement of lesion enhancement parameters such as the bolus arrival time and initial enhancement slope. These important diagnostic parameters have been used to improve cancer diagnosis ([5], [11], [50], [51]). Since the early phase of enhancement is critical for distinguishing cancers from background parenchymal enhancement [6], [52], high fidelity high temporal resolution images produced with ECA may significantly improve identification and characterization of small cancers.

The ECA method introduced here is based on two primary principles.

(1) If k-space data is partitioned into small subsets by acquisition time, each subset retains important kinetic information.

Especially when enhancement is sparse (as in the early phase of contrast uptake), even highly sub-Nyquist acquisitions contain sufficient information to almost fully constrain the evolution of contrast kinetics. The UnWRAP sequence used in this study demonstrates this principle in action, using simple uniform undersampling to sample a representative bandwidth of spatial frequencies. We believe the UnWRAP k-space ordering scheme to be a good choice for sampling the early phase of contrast uptake, but we emphasize that this principle is applicable to any known/deterministic k-space sampling trajectory.

(2) DCE-MRI enhancement is approximately smooth in small time intervals.

Provided kinetic processes are slow on fractional-second timescales and samples are acquired with sufficient SNR and bandwidth, very few measurements of k-space are needed to “tie together” the time-evolution of an image set. This is especially true when all of the partial k-space measurements taken together form a Nyquist-complete set. We designed the UnWRAP acquisition sequence to maximally leverage this principle, but it is applicable to many undersampled reconstruction methods in DCE-MRI.

The UnWRAP method introduced here maintains relatively high SNR over each subset of k-space by sampling a mixture of high and low spatial frequencies. The extent to which both reconstruction methods preserved voxel-wise intensity suggests the UnWRAP k-space trajectory chosen offers some advantages over a standard sequential acquisition. Because it maintains a uniform frequency density in the scan, the UnWRAP sequence samples the k-space center often enough to preserve signal intensity and the k-space edges often enough to correctly assign signal to spatial features. As is true for any acceleration method, effective application of the UnWRAP method requires adequate SNR during each measurement interval.

The results summarized here demonstrate that ECA combined with UNWRAP sampling has promise for improving breast cancer screening and diagnosis. However, this study had some limitations:

- This was a simulation study, and it will be critical to test these results in vivo. These tests are currently underway.
- Motion artifacts were not included in this work. It will be critical to evaluate effects of motion in future simulations as well as in in vivo studies.
- Neither heart nor background enhancement were modeled in these simulations. It will be critical to assess the capacity of ECA to reconstruct diagnostically useful enhancement in the presence of background enhancement.
- T2* effects were not simulated. These effects are significant during the early phase of contrast media uptake, especially in arteries.
- Other sampling trajectories were not tested; because not all scanners can implement all undersampling trajectories, it will be important to test ECAs performance with other types of accelerated acquisitions.

In addition to addressing the study limitations listed above, we suggest several further avenues of future investigation. First and foremost, ECA requires a thorough characterization of its performance across a wide range of noise levels. Second, because of the ubiquity of partial Fourier, parallel imaging, and ML acceleration methods, we will integrate our acceleration algorithm with popular implementations of these. Finally, we hope to test ECA on a wide variety of sampling trajectories and use this process to evaluate the optimality of both ECA and these trajectories in a wider context of DCE-MRI acceleration methods.

Taken as a whole, the data presented in this work constitute an argument that "sparse + smooth enhancement" characterize contrast kinetics in breast DCE-MRI to very high precision during the early phase of contrast media uptake. Smooth enhancement is a stringent condition to impose on DCE-MRI data and, on its own, encodes a great deal of physiological structure. Within such a constraint, even a small number of well-chosen measurements can closely characterize early enhancement in the breast. ECA reconstruction provides a robust framework to increase diagnostic accuracy and improve understanding of hemodynamics in normal breast and cancers.

Appendix A: Formal Description of Partition Constraints

Suppose a total of N measurements are taken over the course of the scan: denote these as $Y = (y_1, \dots, y_N)$. Choose a reconstruction temporal resolution such that we have T time points in the reconstructed image. We may identify each measurement with the time $t_n \in \{1, \dots, T\}$ at which it was taken. For example, if we partition the time duration of the scan into T intervals, then $t_n = 1$ indicates the nth measurement was acquired during the first time interval. We will also write $v_n \in \{1, \dots, V\}$ to denote the k-space voxel (i.e., spatial frequency) at which measurement n was acquired. Our method does not require that we observe an equal number of measurements during each time interval $t = 1, \dots, T$, but we typically expect to have N/T measurements for each t.

Let $X = (X_1, \dots, X_T)$ be a dynamic timeseries image with T time-points, where each static image X_t has V voxels; this is the unknown sequence of true images (up to spatiotemporal discretization) that we aim to reconstruct. In a fully sampled regime ($N = VT$), we would observe the complete k-space data $\tilde{X}_t = \mathcal{F}X_t$ at each time $t = 1, \dots, T$, where $\mathcal{F}$ is the $V \times V$ discrete Fourier transform matrix. In other words, we would measure (a noisy version of) the sequence

$$\tilde{X} = \left(\tilde{X}_1, \dots, \tilde{X}_T\right) = (\mathcal{F}X_1, \dots, \mathcal{F}X_T) = (I_T \otimes \mathcal{F})X,$$

where $\otimes$ denotes the Kronecker product. (Abusing notation, we will interpret X and $\tilde{X}$ as either $V \times T$ matrices or vectors of length VT, depending on the context.) When reconstructing at accelerated time resolution ($N < VT$), the measured data is a proper subset of the full timeseries (i.e., $Y \subsetneq \tilde{X}$). Define

$$\Omega = \left((t_1, v_1), \dots, (t_N, v_N)\right)$$

as the sequence of **(time-index, spatial frequency-index)** pairs at which the k-space measurements $y_1, \dots, y_N$ were taken (i.e., each entry in Ω lies in the set $\{1, \dots, T\} \times \{1, \dots, V\}$). Then $\tilde{X}_\Omega$ is the observed part of the Fourier transform of the dynamic image sequence. For each n, our observation y_n is the voxel v_n from the image $\tilde{X}_{t_n}$, giving the relation $y_n = \left(\tilde{X}_{t_n}\right)_{v_n} +$ noise. The remaining entries in $\tilde{X}$ (i.e., those in $\tilde{X}_{\Omega^c}$) are left unobserved and must be reconstructed by our algorithm.

Our data fidelity constraint stipulates that the observed k-space data Y must remain unaltered by the reconstruction. Therefore, any reconstruction $\hat{X}$ of X must satisfy

$$\left[(I_T \otimes \mathcal{F})\hat{X}\right]_{\Omega} = Y. \tag{1}$$

Appendix B: The Penalty Function

After requiring that each reconstructed image match the k-space points in the partition corresponding to its interval, we minimize over a weighted and regularized smoothness penalty to determine a unique reconstruction. The mathematical details of this process are laid out in this section.

First, we will define the loss function minimized by our reconstruction optimization: the discretized curvature. For a given signal $x = (x_1, \dots, x_T)$ composed of T time points, define the smoothness penalty function S as the l^2-norm of its discrete second derivative:

$$S(x) = \sum_{t=2}^{T-1} |(x_{t+1} - x_t) - (x_t - x_{t-1})|^2.$$

Since S is a quadratic function of x, we may write it in terms of a linear operator D satisfying

$$S(x) = x^* D x.$$

As a matrix, D takes the form

$$D = \begin{pmatrix} 1 & -2 & 1 & 0 & \dots & 0 & 0 & 0 & 0 \\ -2 & 5 & -4 & 1 & \dots & 0 & 0 & 0 & 0 \\ 1 & -4 & 6 & -4 & \dots & 0 & 0 & 0 & 0 \\ 0 & 1 & -4 & 6 & \dots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & 6 & -4 & 1 & 0 \\ 0 & 0 & 0 & 0 & \dots & -4 & 6 & -4 & 1 \\ 0 & 0 & 0 & 0 & \dots & 1 & -4 & 5 & -2 \\ 0 & 0 & 0 & 0 & \dots & 0 & 1 & -2 & 1 \end{pmatrix}$$

The operator D is poorly conditioned. In fact, it can be shown that the respective supremum and infimum of eigenvalues of D (over all values of T) are $\sigma_{max}(D) = 16$ and $\sigma_{min}(D) = 0$.

Since D is (nearly) degenerate and we will later need to invert it, we add a small parameter λ (we chose $\lambda = 10^{-5}$) to regularize the operator D:

$$D_\lambda = D + \lambda I_T,$$

where I_T is the T-dimensional identity. Thus, in the reconstruction process, we use the regularized smoothness penalty

$$S_\lambda(x) = x^* D_\lambda x.$$

If $X = (X_1, \dots, X_T)$ is a dynamic image composed of T static images, each with V voxels, then we can extend S_λ to act on X by

$$S_\lambda(X) = \sum_{v=1}^{V} x_v^* D_\lambda \, x_v,$$

where x_v is the timeseries signal at the v^{th} voxel.

It may also be desirable to penalize a lack of smoothness in some voxels more than in others. To enforce such a prioritization of "interesting" voxels, we may also include voxel-wise weighting terms w_v, generating our full loss function $\mathcal{L}$:

$$\mathcal{L}(X) = \sum_{v=1}^{V} w_v x_v^* D_\lambda \, x_v.$$

Setting W as the $V \times V$ diagonal matrix with entries w_v, we can write the reconstruction optimization problem as

$$\hat{X} = \underset{X}{argmin} \ \{ \langle W, X D_\lambda X^* \rangle \mid [(I_T \otimes \mathcal{F}) X]_\Omega = Y \}. \tag{2}$$

Appendix C: Solutions to the Optimization Problem

General Case: Arbitrary Voxel Weights

In the following section, we will show that the optimization problem posed by the reconstruction has the following unique solution:

$$\hat{X} = [D_\lambda^{-1} \otimes W^{-1}\mathcal{F}^*]_{*,\Omega} \cdot \left([D_\lambda^{-1} \otimes \mathcal{F}W^{-1}\mathcal{F}^*]_{\Omega,\Omega}\right)^{-1} \cdot Y, \tag{3}$$

where, for a $VT \times VT$ matrix M, $M_{*,\Omega}$ denotes the $VT \times |\Omega|$ submatrix of M with columns belonging to Ω; similarly, $M_{\Omega,\Omega}$ is the $|\Omega| \times |\Omega|$ submatrix of M consisting of rows and columns in Ω.

Since computing this solution for $\hat{X}$ requires a large matrix inversion ($N \times N \sim 10^{16}$), we implement this computation iteratively. We initialize the solution in k-space, estimating the Fourier-transformed image sequence $\tilde{X}$ by zero-filling around the N k-space measurements of Y ($\tilde{X}_\Omega = Y$ and $\tilde{X}_{\Omega^c} = 0$). We then descend along the conjugate gradient of the smoothness penalty until we converge to (3). See code for implementation: https://github.com/tyo8/ECA_Demo.

We will now check our solution to the target optimization problem. First, we will show that (3) satisfies the problem constraints.

$$\begin{aligned}
\left[(I_T \otimes \mathcal{F})\hat{X}\right]_\Omega &= \left[(I_T \otimes \mathcal{F})[D_\lambda^{-1} \otimes W^{-1}\mathcal{F}^*]_{*,\Omega} \cdot \left([D_\lambda^{-1} \otimes \mathcal{F}W^{-1}\mathcal{F}^*]_{\Omega,\Omega}\right)^{-1} \cdot Y\right]_\Omega \\
&= [D_\lambda^{-1} \otimes \mathcal{F}W^{-1}\mathcal{F}^*]_{\Omega,\Omega} \cdot \left([D_\lambda^{-1} \otimes \mathcal{F}W^{-1}\mathcal{F}^*]_{\Omega,\Omega}\right)^{-1} \cdot Y \\
&= Y.
\end{aligned}$$

Thus, $\hat{X}$ is a feasible solution to the optimization problem.

Next, we will show that (3) satisfies first-order optimality conditions. Since D_λ and W are both positive-definite matrices, an optimal solution is unique. To see that $\hat{X}$ is optimal, we must show that the gradient of the loss function lies in the span of the gradient of the constraints.

$$\begin{aligned}
\nabla_X \mathcal{L}(\hat{X}) &= (D_\lambda \otimes W)\,\hat{X} \\
&= (D_\lambda \otimes W) \cdot [D_\lambda^{-1} \otimes W^{-1}\mathcal{F}^*]_{*,\Omega} \cdot \left([D_\lambda^{-1} \otimes \mathcal{F}W^{-1}\mathcal{F}^*]_{\Omega,\Omega}\right)^{-1} \cdot Y \\
&= \left[(D_\lambda \otimes W) \cdot \left(D_\lambda^{-1} \otimes W^{-1}\mathcal{F}^*\right)\right]_{*,\Omega} \cdot \left([D_\lambda^{-1} \otimes \mathcal{F}W^{-1}\mathcal{F}^*]_{\Omega,\Omega}\right)^{-1} \cdot Y \\
&= [I_T \otimes \mathcal{F}^*]_{*,\Omega} \cdot \left([D_\lambda^{-1} \otimes \mathcal{F}W^{-1}\mathcal{F}^*]_{\Omega,\Omega}\right)^{-1} \cdot Y
\end{aligned}$$

Since $[I_T \otimes \mathcal{F}^*]_{*,\Omega}$ is the gradient of the constraint function, we've shown that $\nabla_X \mathcal{L}(\hat{X})$ lies in the span of the constraint gradients. It follows that $\hat{X}$ is the unique solution to the reconstruction optimization problem.

Special Case: Uniform Voxel Weights

When all voxels are uniformly weighted ($W \propto I_V$), the solution simplifies significantly:

$$\hat{X} = (I_T \otimes \mathcal{F}^*)\bar{Y},$$

where $\bar{Y} = [D_\lambda^{-1} \otimes I_V]_{*,\Omega} \cdot \left([D_\lambda^{-1} \otimes I_V]_{\Omega,\Omega}\right)^{-1} \cdot Y$.

Appendix D: Regularization in the Optimization Problem

Recall the penalty function defined above:

$$S(x) = x^* D x.$$

Since the operator D is not invertible (and must be inverted to efficiently compute the optimal solution), we add a small diagonal element λI_T to D to get the regularized penalty function

$$S_\lambda(x) = x^* D_\lambda x.$$

As with any regularization parameter, different values of λ offer different trade-offs between the optimality of the solution and the speed at which it is reached. We hope to choose λ large enough that D_λ is well-conditioned and small enough that $S_\lambda - S$ is small compared to S. Equivalently, we require that

$$\lambda^{-1} \gg \frac{\langle W, XX^* \rangle}{\langle W, XDX^* \rangle}$$

and $-\log_{10} \lambda \leq k$ for some integer k (which should be chosen empirically based on problem size, system requirements, and the conditioning of the weights matrix W). We chose $\lambda = 10^{-5}$ because it converged sufficiently quickly and did not significantly alter the penalty function computed in our experiment. To illustrate the regularization-vs-performance tradeoff, a pair of plots for different values of λ in a small-scale ($VT \sim 10^6$) reconstruction are shown below (**Figure 13**). As regularization increases, we converge more quickly to a solution, but that solution is less accurate.

Fig 13. Reconstruction error and convergence speed as a function of regularization.

(Blue) Error, measured here by normalized mean-square error (MSE), increases with regularization strength. (Red) Computation time, measured in number of iterations, decreases with regularization strength.

References

[1] L. Zhang, M. Tang, Z. Min, J. Lu, X. Lei, and X. Zhang, "Accuracy of combined dynamic contrast-enhanced magnetic resonance imaging and diffusion-weighted imaging for breast cancer detection: A meta-analysis," Acta radiol., vol. 57, no. 6, pp. 651–660, Jun. 2016, doi: 10.1177/0284185115597265.

[2] J. Melnikow et al., "Supplemental Screening for Breast Cancer in Women with Dense Breasts: A Systematic Review for the U.S. Preventive Services Task Force HHS Public Access," Ann Intern Med, vol. 164, no. 4, pp. 268–278, 2016, doi: 10.7326/M15-1789.

[3] L. W. Turnbull, "Dynamic contrast-enhanced MRI in the diagnosis and management of breast cancer," NMR in Biomedicine, vol. 22, no. 1. NMR Biomed, pp. 28–39, 2009, doi: 10.1002/nbm.1273.

[4] C. Wu, F. Pineda, D. A. Hormuth, G. S. Karczmar, and T. E. Yankeelov, "Quantitative analysis of vascular properties derived from ultrafast DCE-MRI to discriminate malignant and benign breast tumors," Magn. Reson. Med., vol. 81, no. 3, pp. 2147–2160, Mar. 2019, doi: 10.1002/mrm.27529.

[5] S. J. Lee, K. H. Ko, H. K. Jung, J. E. Koh, and A. Y. Park, "The additional utility of ultrafast MRI on conventional DCE-MRI in evaluating preoperative MRI of breast cancer patients," Eur. J. Radiol., vol. 124, Mar. 2020, doi: 10.1016/j.ejrad.2020.108841.

[6] F. D. Pineda et al., "Ultrafast Bilateral DCE-MRI of the Breast with Conventional Fourier Sampling: Preliminary Evaluation of Semi-quantitative Analysis," Acad. Radiol., vol. 23, no. 9, pp. 1137–1144, Sep. 2016, doi: 10.1016/J.ACRA.2016.04.008.

[7] S. A. Jansen, X. Fan, G. S. Karczmar, H. Abe, R. A. Schmidt, and G. M. Newstead, "Differentiation between benign and malignant breast lesions detected by bilateral dynamic contrast-enhanced MRI: A sensitivity and specificity study," Magn. Reson. Med., vol. 59, no. 4, pp. 747–754, Apr. 2008, doi: 10.1002/mrm.21530.

[8] B. J. Woods et al., "Malignant-lesion segmentation using 4D co-occurrence texture analysis applied to dynamic contrast-enhanced magnetic resonance breast image data," J. Magn. Reson. Imaging, vol. 25, no. 3, pp. 495–501, Mar. 2007, doi: 10.1002/jmri.20837.

[9] S. C. Agner et al., “Textural kinetics: A novel dynamic contrast-enhanced (DCE)-MRI feature for breast lesion classification,” J. Digit. Imaging, vol. 24, no. 3, pp. 446–463, Jun. 2011, doi: 10.1007/s10278-010-9298-1.

[10] M. Mardani et al., “Deep generative adversarial neural networks for compressive sensing MRI,” IEEE Trans. Med. Imaging, vol. 38, no. 1, 2019, doi: 10.1109/TMI.2018.2858752.

[11] J. Milenković, M. U. Dalmış, J. Žgajnar, and B. Platel, “Textural analysis of early-phase spatiotemporal changes in contrast enhancement of breast lesions imaged with an ultrafast DCE-MRI protocol,” Med. Phys., vol. 44, no. 9, pp. 4652–4664, Sep. 2017, doi: 10.1002/mp.12408.

[12] E. Henderson, B. K. Rutt, and T. Y. Lee, “Temporal sampling requirements for the tracer kinetics modeling of breast disease,” Magn. Reson. Imaging, vol. 16, no. 9, pp. 1057–1073, 1998, doi: 10.1016/S0730-725X(98)00130-1.

[13] G. J. M. Parker et al., “Experimentally-derived functional form for a population-averaged high-temporal-resolution arterial input function for dynamic contrast-enhanced MRI,” Magn. Reson. Med., 2006, doi: 10.1002/mrm.21066.

[14] L. E. Kershaw and H. L. M. Cheng, “Temporal resolution and SNR requirements for accurate DCE-MRI data analysis using the AATH model,” Magn. Reson. Med., vol. 64, no. 6, pp. 1772–1780, Dec. 2010, doi: 10.1002/mrm.22573.

[15] T. E. Yankeelov et al., “Comparison of a reference region model with direct measurement of an AIF in the analysis of DCE-MRI data,” Magn. Reson. Med., 2007, doi: 10.1002/mrm.21131.

[16] L. E. Kershaw and H. L. M. Cheng, “A general dual-bolus approach for quantitative DCE-MRI,” Magn. Reson. Imaging, vol. 29, no. 2, pp. 160–166, Feb. 2011, doi: 10.1016/j.mri.2010.08.009.

[17] A. Chatterjee et al., “Performance of Ultrafast DCE-MRI for Diagnosis of Prostate Cancer,” Acad. Radiol., vol. 25, no. 3, pp. 349–358, Mar. 2018, doi: 10.1016/J.ACRA.2017.10.004.

[18] F. D. Pineda, T. O. Easley, and G. S. Karczmar, “Dynamic field-of-view imaging to increase temporal resolution in the early phase of contrast media uptake in breast DCE-

MRI: A feasibility study," Med. Phys., vol. 45, no. 3, pp. 1050–1058, 2018, doi: 10.1002/mp.12747.

[19] M. A. Griswold et al., "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," Magn. Reson. Med., vol. 47, no. 6, pp. 1202–1210, Jun. 2002, doi: 10.1002/mrm.10171.

[20] Z. Xiao, W. S. Hoge, R. V Mulkern, L. Zhao, G. Hu, and W. E. Kyriakos, "Comparison of Parallel MRI Reconstruction Methods for Accelerated 3D Fast Spin-Echo Imaging," 2008, doi: 10.1002/mrm.21679.

[21] S. Ramani and J. A. Fessler, REGULARIZED PARALLEL MRI RECONSTRUCTION USING AN ALTERNATING DIRECTION METHOD OF MULTIPLIERS. 2011.

[22] M. Lustig and J. M. Pauly, "SPIRiT: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space," Magn. Reson. Med., vol. 64, no. 2, p. n/a-n/a, Aug. 2010, doi: 10.1002/mrm.22428.

[23] B. Platel, R. Mus, T. Welte, N. Karssemeijer, and R. Mann, "Automated Characterization of Breast Lesions Imaged With an Ultrafast DCE-MR Protocol," IEEE Trans. Med. Imaging, vol. 33, no. 2, pp. 225–232, Feb. 2014, doi: 10.1109/TMI.2013.2281984.

[24] Y. Le, R. Kroeker, H. D. Kipfer, and C. Lin, "Development and evaluation of TWIST Dixon for dynamic contrast-enhanced (DCE) MRI with improved acquisition efficiency and fat suppression," J. Magn. Reson. Imaging, 2012, doi: 10.1002/jmri.23663.

[25] R. P. Lim et al., "3D time-resolved MR angiography (MRA) of the carotid arteries with time-resolved imaging with stochastic trajectories: Comparison with 3D contrast-enhanced bolus-chase MRA and 3D time-of-flight MRA," Am. J. Neuroradiol., 2008, doi: 10.3174/ajnr.A1252.

[26] F. R. Korosec, R. Frayne, T. M. Grist, and C. A. Mistretta, "Time-Resolved Contrast-Enhanced 3D MR Angiography," 1996.

[27] J. Hennig, K. Scheffler, J. Laubenberger, and R. Strecker, "Time-Resolved Projection Angiography after Bolus Injection of Contrast Agent," 1997.

[28] M. Saranathan, D. W. Rettmann, B. A. Hargreaves, S. E. Clarke, and S. S. Vasanawala,

"DIfferential Subsampling With Cartesian Ordering (DISCO): A High Spatio-temporal Resolution Dixon Imaging Sequence for Multiphasic Contrast Enhanced Abdominal Imaging," J. Magn. Reson. IMAGING, vol. 35, pp. 1484–1492, 2012, doi: 10.1002/jmri.23602.

[29] W. A. Willinek et al., "4D Time-Resolved MR Angiography With Keyhole (4D-TRAK): More Than 60 Times Accelerated MRA Using a Combination of CENTRA, Keyhole, and SENSE at 3.0T," J. Magn. Reson. Imaging, vol. 27, pp. 1455–1460, 2008, doi: 10.1002/jmri.21354.

[30] R. Liu, L. Shi, S. Yu, N. Xiong, and D. Wang, "Reconstruction of Undersampled Big Dynamic MRI Data Using Non-Convex Low-Rank and Sparsity Constraints," Sensors, vol. 17, no. 3, p. 509, Mar. 2017, doi: 10.3390/s17030509.

[31] D. L. Donoho, "Compressed Sensing," IEEE Trans. Inf. THEORY, vol. 52, no. 4, 2006, doi: 10.1109/TIT.2006.871582.

[32] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," Magn. Reson. Med., 2007, doi: 10.1002/mrm.21391.

[33] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing MRI: A look at how CS can improve on current imaging techniques," IEEE Signal Processing Magazine. 2008, doi: 10.1109/MSP.2007.914728.

[34] G. Adluru, T. Tasdizen, M. C. Schabel, and E. V. R. DiBella, "Reconstruction of 3D dynamic contrast-enhanced magnetic resonance imaging using nonlocal means," J. Magn. Reson. Imaging, vol. 32, no. 5, pp. 1217–1227, Nov. 2010, doi: 10.1002/jmri.22358.

[35] O. N. Jaspan, R. Fleysher, and M. L. Lipton, "Compressed sensing MRI: A review of the clinical literature," British Journal of Radiology, vol. 88, no. 1056. 2015, doi: 10.1259/bjr.20150487.

[36] K. Hammernik et al., "Learning a variational network for reconstruction of accelerated MRI data," Magn. Reson. Med., 2018, doi: 10.1002/mrm.26977.

[37] A. Bora, A. Jalal, E. Price, and A. G. Dimakis, "Compressed Sensing using Generative Models," 2017, doi: 10.5829/idosi.wasj.2013.23.04.368.

[38] G. Yang et al., “DAGAN: Deep De-Aliasing Generative Adversarial Networks for Fast Compressed Sensing MRI Reconstruction,” IEEE Trans. Med. Imaging, vol. 37, no. 6, 2018, doi: 10.1109/TMI.2017.2785879.

[39] M. A. Dedmari, S. Conjeti, S. Estrada, P. Ehses, T. Stöcker, and M. Reuter, “Complex fully convolutional neural networks for mr image reconstruction,” in Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 2018, vol. 11074 LNCS, doi: 10.1007/978-3-030-00129-2_4.

[40] J. C. Ye, Y. Han, and E. Cha, “Deep Convolutional Framelets: A General Deep Learning Framework for Inverse Problems,” 2017, doi: 10.1137/17M1141771.

[41] F. Pineda, D. Sheth, H. Abe, M. Medved, and G. S. Karczmar, “Low-dose imaging technique (LITE) MRI: initial experience in breast imaging,” Br. J. Radiol., vol. 92, no. 1103, p. 20190302, Nov. 2019, doi: 10.1259/bjr.20190302.

[42] X. Fan et al., “New model for analysis of dynamic contrast-enhanced MRI data distinguishes metastatic from nonmetastatic transplanted rodent prostate tumors,” Magn. Reson. Med., vol. 51, no. 3, pp. 487–494, Mar. 2004, doi: 10.1002/mrm.10737.

[43] R. Otazo, E. Candès, and D. K. Sodickson, “Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components,” Magn. Reson. Med., vol. 73, no. 3, 2015, doi: 10.1002/mrm.25240.

[44] F. Ong and M. Lustig, “Beyond low rank + sparse: Multi-scale low rank matrix decomposition,” in ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings, 2016, vol. 2016-May, doi: 10.1109/ICASSP.2016.7472561.

[45] D. K. Woolf et al., “Arterial input functions in dynamic contrast-enhanced magnetic resonance imaging: Which model performs best when assessing breast cancer response?,” Br. J. Radiol., 2016, doi: 10.1259/bjr.20150961.

[46] D. M. McGrath, D. P. Bradley, J. L. Tessier, T. Lacey, C. J. Taylor, and G. J. M. Parker, “Comparison of model-based arterial input functions for dynamic contrast-enhanced MRI in tumor bearing rats,” Magn. Reson. Med., 2009, doi: 10.1002/mrm.21959.

[47] G. Mcgibney, M. R. Smith, S. T. Nichols, and A. Crawley, “Quantitative Evaluation of Several Partial Fourier Reconstruction Algorithms Used in MRI,” 1993. Accessed: Nov. 07, 2018. [Online]. Available: http://mriquestions.com/uploads/3/4/5/7/34572113/partial_fourier_methods_2010_pfi.pdf.

[48] M. R. Hestenes and E. Stiefel, “Methods of Conjugate Gradients for Solving Linear Systems 1,” 1952.

[49] H. Akima, “A New Method of Interpolation and Smooth Curve Fitting Based on Local Procedures,” *J. ACM*, vol. 17, no. 4, pp. 589–602, Oct. 1970, doi: 10.1145/321607.321609.

[50] R. D. Mus et al., “Time to enhancement derived from ultrafast breast MRI as a novel parameter to discriminate benign from malignant breast lesions,” *Eur. J. Radiol.*, vol. 89, pp. 90–96, Apr. 2017, doi: 10.1016/j.ejrad.2017.01.020.

[51] A. Ohashi et al., “Diagnostic performance of maximum slope: A kinetic parameter obtained from ultrafast dynamic contrast-enhanced magnetic resonance imaging of the breast using k-space weighted image contrast (KWIC),” *Eur. J. Radiol.*, vol. 118, pp. 285–292, Sep. 2019, doi: 10.1016/j.ejrad.2019.06.012.

[52] S. A. Jansen et al., “Characterizing early contrast uptake of ductal carcinoma in situ with high temporal resolution dynamic contrast-enhanced MRI of the breast: A pilot study,” *Phys. Med. Biol.*, vol. 55, no. 19, Oct. 2010, doi: 10.1088/0031-9155/55/19/N02.