APPLIED SCIENCES AND ENGINEERING

Compressive 3D ultrasound imaging using a single sensor

Pieter Kruizinga,[1,2]* Pim van der Meulen,[3] Andrejs Fedjajevs,[3] Frits Mastik,[1] Geert Springeling,[4] Nico de Jong,[1,2] Johannes G. Bosch,[1] Geert Leus[3]

Three-dimensional ultrasound is a powerful imaging technique, but it requires thousands of sensors and complex hardware. Very recently, the discovery of compressive sensing has shown that the signal structure can be exploited to reduce the burden posed by traditional sensing requirements. In this spirit, we have designed a simple ultrasound imaging device that can perform three-dimensional imaging using just a single ultrasound sensor. Our device makes a compressed measurement of the spatial ultrasound field using a plastic aperture mask placed in front of the ultrasound sensor. The aperture mask ensures that every pixel in the image is uniquely identifiable in the compressed measurement. We demonstrate that this device can successfully image two structured objects placed in water. The need for just one sensor instead of thousands paves the way for cheaper, faster, simpler, and smaller sensing devices and possible new clinical applications.

INTRODUCTION

Ultrasound imaging is a widely used technique in medical decision-making, mainly due to the fact that ultrasound devices use nonionizing radiation and are relatively low-cost. By carefully aiming the transducer, the operator can use ultrasound to obtain real-time images of specific cross sections of the body. The contrast in ultrasound images is due to sound speed and density differences between tissues. This is why, for example, a baby's skull in a watery womb forms a clear picture in ultrasound imaging. Almost all ultrasound images are made with transducers composed of many tiny sensors (between 64 and 10,000) that can transmit short ultrasonic bursts (typically 1 to 40 MHz). These bursts are then reflected by the tissue and recorded by the same sensor array. The time delay between transmission of the burst and detection of the echoes defines where the tissue is located. The strength of the reflected echo contains information on the density of the tissue.

The sensors that make up these arrays are made mainly from piezoelectric crystals. When a voltage is applied to the material, the crystal vibrates and produces a short ultrasonic burst of a few cycles (1). When the sensors are activated with appropriate delays, the emitted ultrasound wave travels along a straight line (in the form of a narrow beam) and can be focused on a point of interest. With these focused beams, a whole volume can be scanned (Fig. 1A). In terms of receiving the signal reception, similar delays can be used to selectively enhance the signals from scatterers along the beam of interest. This type of focusing that uses delayed signals in transmitting and/or receiving is named beamforming and has been a source of active academic research for many years (2). One advantage of applying beamforming in receive mode is that it can be carried out digitally after the received ultrasound field has been recorded (Fig. 1B). An alternative technique applied in ultrasound imaging is to use an unfocused ("plane") wave during transmission. The focus is then regained upon reception, allowing for a faster frame rate because the medium is not "scanned" line by line with a focused beam.

The majority of ultrasound imaging is performed using transducer arrays that have their sensors distributed along one dimension (Fig. 1, A and B), thereby offering a two-dimensional (2D) cross-sectional view of the inside of the human body. To obtain a full 3D view, one needs to mechanically move or rotate the 1D array or to create a 2D array that allows beam steering in two directions. These 2D arrays have been shown to produce stunning high-resolution 3D images, for example, of faces of fetuses (3, 4). However, 3D ultrasound imaging is a long way from achieving the popularity of 2D ultrasound imaging apart from in certain medical disciplines, such as obstetrics, cardiology, and image-guided intervention. This is surprising given that a full 3D view should in almost all cases allow for a better assessment of the specific medical question. One of the main reasons for this low acceptance of 3D ultrasound imaging is that 3D imaging requires a high hardware complexity (>1000 sensors on a small footprint, integrated electronics, high data rates, etc.). Most of this is necessary to cope with the constraints imposed by the Nyquist theorem when sampling the received radiofrequency ultrasound signals.

Fortunately, as shown by recent discoveries in statistics (5–7), the classical idea that digitization of analog signals (such as ultrasound waves) demands uniform sampling at rates twice as high as the highest frequency present (known as the Nyquist theorem) is no longer a rule carved in stone. This discovery has opened up a very active field of research known as compressive sensing (CS) (8). One of the leading thoughts in CS is that signals generally contain some natural structure. This structure can be exploited to compress the signal to one that is many times smaller than the original, as is done, for example, in jpeg compression of images. However, until recently, this size reduction was always done after the signal had been acquired at full Nyquist rate. CS now allows the compression to be merged with the sensing. As an immediate consequence, the number of measurements required to recover the signal of interest can be drastically reduced, potentially leading to cheaper, faster, simpler, and smaller sensing devices. In the case of ultrasound imaging, this has already led to the development of successful new strategies for sampling and image reconstruction (9–12).

One of the most striking implementations of CS has been the design of the single-pixel camera by Duarte et al. (13). Here, the authors managed to reconstruct images using an adjustable mirror mask and a single photodetector. By projecting the image onto the photodetector using random mask patterns, they were able to recover the actual image from
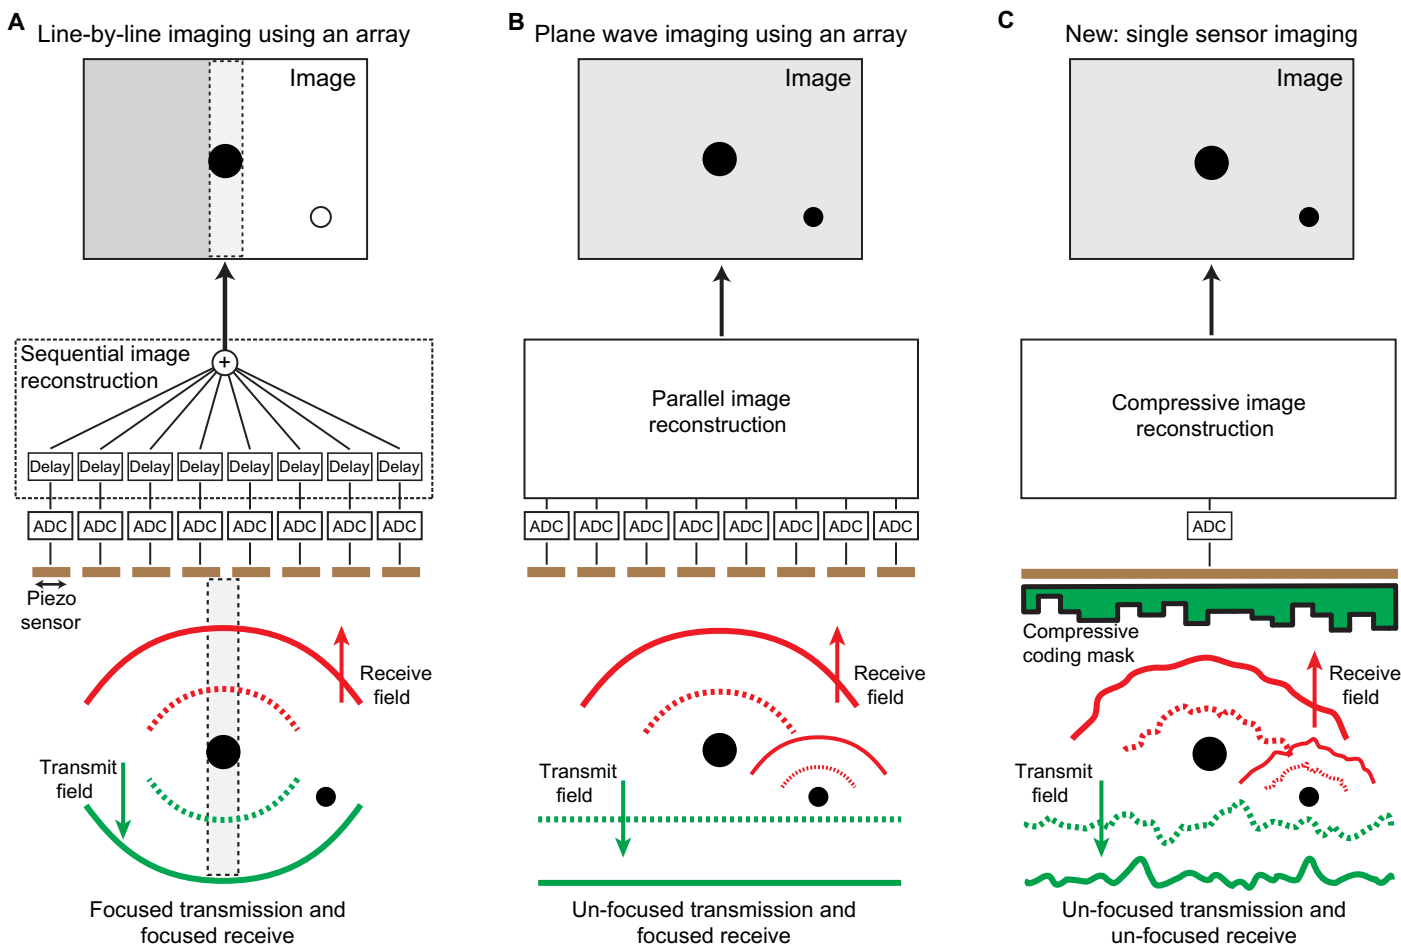
[1]Thoraxcenter–Biomedical Engineering, Erasmus Medical Center, 3000 CA Rotterdam, Netherlands. [2]Faculty of Applied Sciences–Imaging Physics, Delft University of Technology, 2600 AA Delft, Netherlands. [3]Faculty of Electrical Engineering, Mathematics and Computer Science–Micro Electronics, Delft University of Technology, 2628 CD Delft, Netherlands. [4]Department of Experimental Medical Instruments, Erasmus Medical Center, 3015 CN Rotterdam, Netherlands.
*Corresponding author. Email: p.kruizinga@erasmusmc.nl

Fig. 1. The three techniques of ultrasound imaging formation. Conventional ultrasound imaging requires an array of sensors to (**A**) focus the transmitted and received wave field or (**B**) only the received wave field by postprocessing the individual signals after analog-to-digital conversion (ADC). We propose using only one sensor without focusing in transmission or receive (**C**).

a number of measurements much less than the number of image pixels, hence the name compressive imaging. This work has inspired many others to try and use the idea in a similar way for different applications (*14–16*), including within the ultrasound research community.

Our work follows the concept central to compressive imaging: that of projecting the object/image information through a set of incoherent functions onto a single measurement. A successful strategy for obtaining these compressed measurements is to apply a so-called "coded aperture" (*17–19*). In this case, the information—whether it entails different light directions, different frequency bands, or any other information that is conventionally needed to build up an image—is coded into the available aperture and associated measurement. Smart algorithms are then needed to decode the retrieved measurements to form an image.

Along these lines, here we introduce for the first time in ultrasound imaging a simple compressive imaging device that can create 3D images. Our device contains one large piezo sensor that transmits an ultrasonic wave through a simple plastic coding mask. Local variations in the mask thickness (Fig. 1C) cause local delays, which scrambles the phase of the wave field. Similar techniques were demonstrated by Melde *et al.* (*20*) and Brown *et al.* (*21*) to produce holograms that generate predefined (photo)acoustic fields. This enables a complex interference pattern to

propagate inside the volume, removing ambiguity among echoes from different pixels, as illustrated in Fig. 2. The interference pattern propagates through the medium, scatters from objects within the medium, and then propagates back through the coding mask onto the same ultrasound sensor, providing a single compressed ultrasound measurement of the object.

Using our device, we aim to mitigate the hardware complexity associated with conventional 3D ultrasound. The manufacturing costs of our device will be much lower. A simple, cheap, single-element transducer is used, and the plastic coding mask can be produced for less than a euro. These lower costs enable broader use of these 3D compressive imaging devices, for example, for long-term patient monitoring. Because the imaging is 3D, finding and maintaining the proper 2D view does not require a trained operator. One could also envision other applications, such as minimally invasive imaging catheters that are too thin to accommodate the hundreds or thousands of electrical wires currently needed for 3D ultrasound imaging. Here, we discuss this new imaging concept and demonstrate experimentally that a simple setup with a coding mask allows for 3D ultrasound imaging—thereby paving the way for an entirely new method of imaging in which the complexity is shifted away from the hardware and toward computing power.
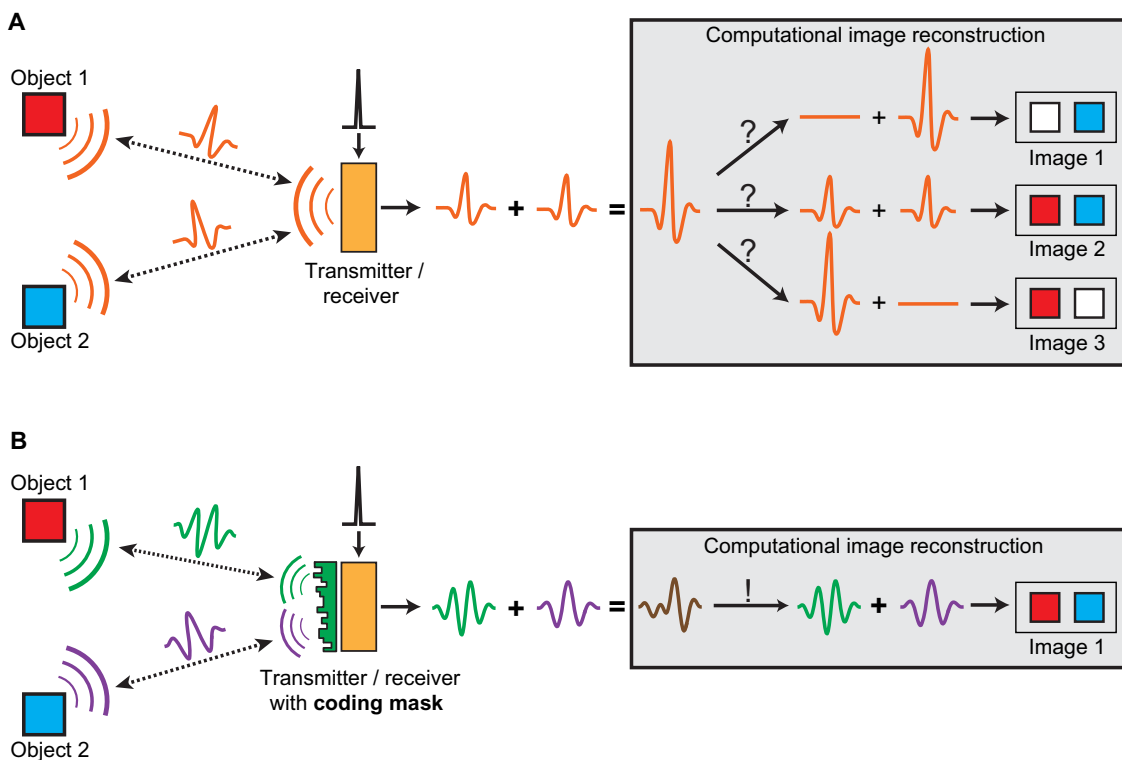
Fig. 2. A spatiotemporally varying ultrasound field allows for compressive imaging. (**A**) One sensor (transmitter and receiver) cannot distinguish between two objects when the transmitted wave field has no spatiotemporal diversity. (**B**) A simple coding mask can introduce a spatiotemporally varying wave field, which allows unique signal separation between two objects.

RESULTS

Ultrasound wave field diversity using a coded aperture

For single-sensor compressive imaging, the 3D object information needs to be projected onto a set of incoherent basis functions via a coded aperture. This is accomplished by breaking the phase uniformity of the ultrasound wave in both transmission and reception. Conventionally, the loss of phase uniformity (for example, through aberration) is unwanted because it leads to blurring and artifacts in the image. In our case, however, the pattern of this interference is known and is used to our benefit.

The coding mask placed in front of the ultrasound sensor consists of a plastic material that supports ultrasound propagation with a speed of around 2750 m/s, in contrast to the speed in water, which is 1480 m/s at room temperature. Varying the mask thickness consequently varies the time the ultrasound wave spends within the plastic before entering the medium, in our case water. For an ultrasound wave with a central frequency of 5 MHz to be delayed by a full wavelength with respect to the rest of the wave requires a local thickness increase of 641 µm. When the thickness of the plastic varies spatially, the phase uniformity of the propagating wave is broken, initiating a random deterministic interference pattern in the medium. The mask covers the complete aperture of the sensor (diameter, 12.7 mm), and the thickness randomly varies between 0.1 and 1 mm. For more details, see the Materials and Methods section. It is important to note that the mask generates an interference pattern that essentially stretches in four dimensions (three spatial and one temporal). So for every location in 3D space (every pixel), we have one unique ultrasound signal that originates from that location. Figure 3 and movie S1 show how the natural focusing of the ultrasound sensor is altered after we place our coding mask in front of the sensor.

After excitation with a short high-voltage pulse, the piezo sensor used to insonify the medium produces a short ultrasound burst. Although short, the limited bandwidth ensures the burst has a certain length (time) and contains several (positive-negative) cycles. After propagating through the mask, the pulse length increases further due to the distortion of the wavefront by the coding mask (Fig. 4A). These waves will propagate spherically in the medium and will interfere constructively and destructively, a process that stretches out longer in time than an undistorted wavefront does. Consequently, the available spatial bandwidth increases compared to an undistorted burst by a mask-less sensor. This effect is highlighted in Fig. 4B, which shows two spatial frequency spectra: one at the beginning of the pulse and one spectrum a few microseconds later. Here, we assume that most of the acoustic energy is directly transmitted through the coding mask; yet, in reality, some of the energy will be reflected back from the mask-water interface and possibly redirected in the medium through the mask in a second pass after bouncing back from the mask-transducer interface, thereby increasing the spatial variability even more.

Any ultrasound imaging problem requires information about the spatial ultrasound field, typically acquired by array sampling. By modeling the aperture mask as a collection of point sensors, each having different transmit/receive delays, the mask can still be regarded as a sensor array. However, all these point sensor signals are subsequently summed by the piezo sensor just after they have passed though the mask, resulting in a single compressed measurement, as depicted in the right panel of Fig. 1 (for a further comparison between a sensor array and a single sensor with a coding mask, see text S2 and figs. S1 and S2). The delays produced by the mask create complex spatiotemporal interference patterns that ensure that each pixel generates a unique
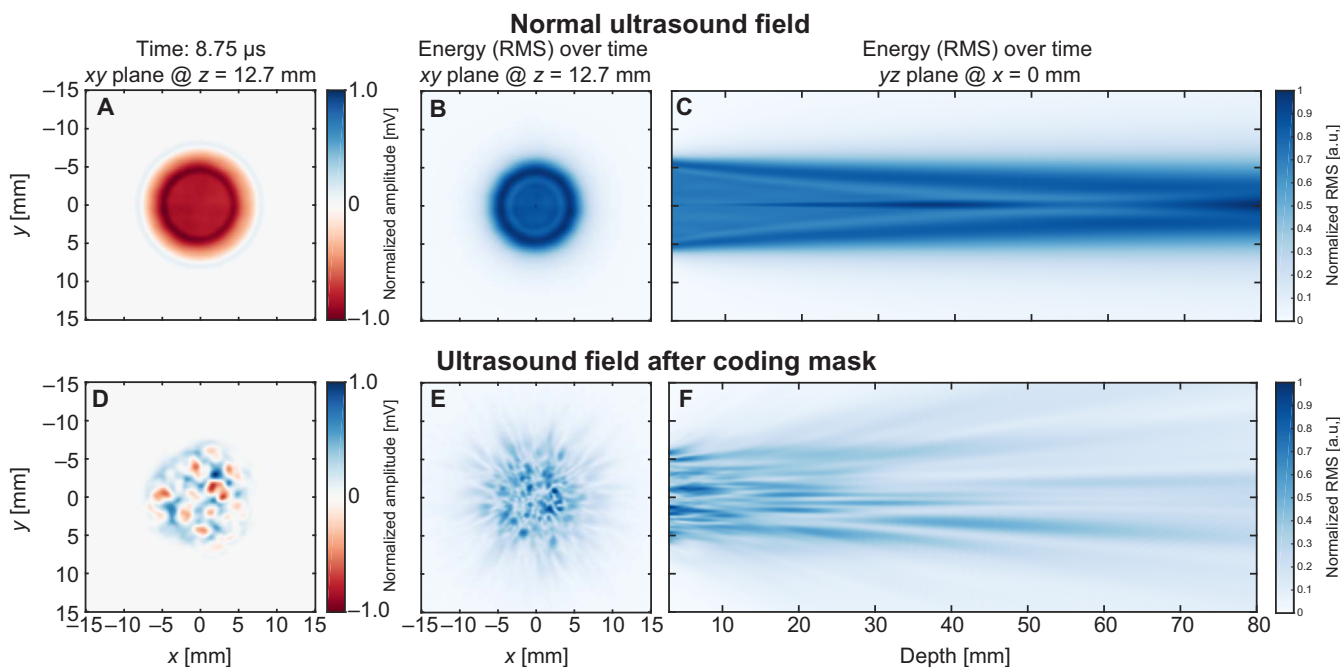
Fig. 3. A random delay coding mask breaks the phase uniformity of the ultrasound transmission to enable compressive imaging. The top panel shows the ultrasound field transmitted from a normal piezoelectric sensor recorded by a small microphone (hydrophone) in front of the sensor. Note that the transmitted wave focuses naturally into a narrower bundle after propagation. The bottom panel shows the same transmitted wave field after propagation through the plastic coding mask. Here, the phase uniformity of the wave is completely lost. The images in (**A**) and (**D**) show one time sample (at 8.75 μs) of the transmitted wave field in a 2D plane (12.7 mm away from the sensor surface) perpendicular to the propagation axis, before and after the addition of the coding mask. The images in (**B**) and (**E**) show a projection of the total ultrasound energy over time in the same 2D plane that is used for (A) and (D). The images (**C**) and (**F**) show the energy projection over time in a 2D plane parallel to the ultrasound propagation axis.

temporal signal in the compressed measurement. This unique pixel signature enables direct imaging without the need for uncompressed spatial measurements.

Because of the limitations in temporal bandwidth, as well as in mask thickness distribution, it is not guaranteed that all pixel echo signals are uncorrelated. To introduce more diversity between the pixels, we chose to rotate the mask in front of the sensor such that the interference pattern is rotated along with it, thereby obtaining additional measurements containing new information. Mask translations (in x, y, or z dimension) would provide similar pixel diversity. Note that for a sufficiently good mask, mask rotation or translation is not necessarily required but significantly decreases interpixel correlations. The result on the ultrasound field as received by every pixel in one xy plane is depicted in Fig. 4C, where the effect of rotation in terms of temporal variety is shown at two points in time for three different pixels in the field of view.

Pixel diversity

If we model the ultrasound field of the coded aperture by approximating the mask aperture as a collection of point sources and sensors (see Materials and Methods section), we can analyze the coding performance of the aperture. To this end, we compute the pulse-echo signals for pixels in the central (xy plane) area at a fixed depth (z) relative to the sensor and then compute their cross-correlations. Ideally, all signals are uncorrelated, so that any superposition of pulse-echo signals can be uniquely broken down into its composing individual pulse-echo signals (as in Fig. 2), in which case the signals for each pixel can be unambiguously resolved. Figure 5B shows the histogram of interpixel correlations for a plane parallel to the sensor surface at a depth

of 12.7 mm (aperture diameter size). We computed these correlations without a mask for several rotations. Because the wave field phase without using a mask is highly uniform, most pulse-echo signals are highly correlated. However, adding a mask to the single sensor causes the distribution to be zero-mean and removes the high correlations around +1 and −1. As expected, these results illustrate that adding more measurements by rotating the mask causes correlations to be distributed more narrowly around zero; pulse-echo signals become more orthogonal. This tells us that the use of rotation removes ambiguities and consequently increases resolvability of scatterers. Here, we do not include the signal correlations over the depth dimension, because generally speaking, the decorrelation in the acoustic propagation direction comes naturally with the pulse-echo delay, or in other words, using only one sensor, we can estimate the depth of an object, but we cannot see whether the object is located left or right from the sensor.

Instead of looking at the pixel-pixel correlation in an entire xy imaging plane, we can also observe the local correlations. Such observations are demonstrated in Fig. 5D, which visualizes the correlations between the center pixel and all other local pixels. A higher correlation between pixels results in lower resolution due to their increased ambiguity. These figures show that the resolution in x and y worsens as the depth increases and that the resolution in z is better than that in xy, which is often the case for ultrasound imaging.

For a more detailed analysis, we can also use the Cramer-Rao lower bound (CRLB) and apply it in the aforementioned approximate model (22). The CRLB is a bound on the best attainable estimation variance for any unbiased estimator (for example, the estimated parameter can be the image itself). That is, the CRLB is related to the average squared
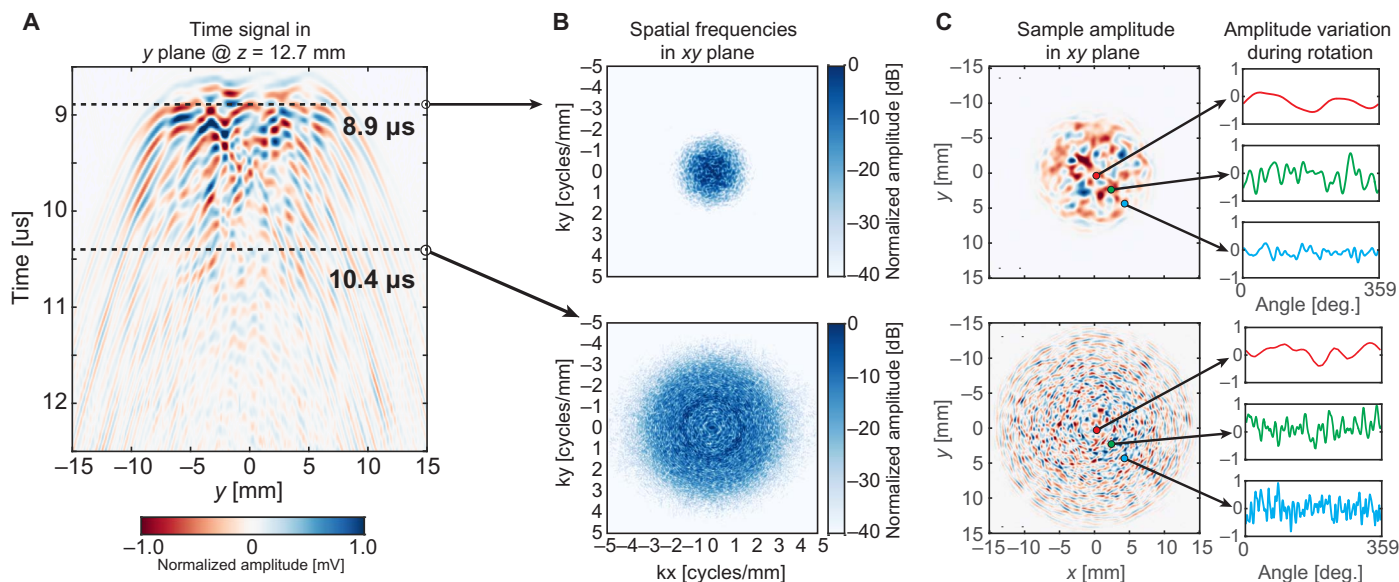
Fig. 4. Spatiotemporal diversity as well as wave field rotation provides a unique signal for every pixel. (**A**) A 2D slice (time and y dimension) of the 4D ultrasound field (time and three spatial dimensions), where $x = 0$ and $z = 12.7$ mm. Every time sample of this 4D wave field has unique spatial properties, as is shown in (**B**) and (**C**). (**B**) The spatial ultrasound field at a later point in time contains higher spatial frequencies, which is desirable for reconstructing objects with high resolution. The possibility of rotating the mask offers another dimension of signal variation to be used for image reconstruction. (**C**) Analysis of two samples based on their spatial properties with respect to coding mask rotation. For pixels further away from the center, the local variance is higher than for pixels closer to the center.

estimation error. Hence, a lower CRLB suggests that the best obtainable resolution is lower as well. The actual estimator (beamformer) used in this paper is described in the next section and does not necessarily achieve the CRLB. Consequently, the error covariance of the actual estimator is typically greater than the CRLB, but the CRLB is a good measure to study trends in the resolution. The CRLB is derived for the case of estimating the position of a single scatterer given a pulse-echo measurement with additive Gaussian noise (see the Materials and Methods section for more details). For a 3D problem, this results in estimating three unknowns (x, y, and z coordinates) from N measurements, where N is the total number of samples. Although position estimation and imaging are not the same, one could argue that the best value obtained for positional error spread (for example, 3 SD of error) nevertheless is indicative for the achievable imaging resolution.

To analyze the behavior and performance of the mask, we compute the CRLB for several points in 3D space, comparing the bounds for measurements obtained with no and four rotations (Fig. 5E). We distinguish two main factors that contribute to the CRLB. First, rotation of the mask significantly reduces the CRLB, and a very large performance gain is obtained by using more than one rotation. This effect is apparent when comparing the CRLB for several rotations in the graphs in Fig. 5E. Second, when we analyze the CRLB for no rotation, we see that the geometry of the problem determines the performance, which is visible in the second panel in Fig. 5E. That is, the performance for a small radial distance is dependent on the specific mask layout and worsens as the received energy decreases inversely proportional to the distance from the sensor. For computing the CRLB graphs, we used a constant noise energy such that the signal-to-noise ratio of a pixel at $(x, y, z) = (0, 0, 12.7$ mm) is equal to 10 dB.

Signal model and image reconstruction
Because only one sensor is used, there is no possibility of applying standard beamforming techniques typically used in conjunction with

sensor arrays. These techniques mainly consider the geometry of the problem and do not require detailed information on the ultrasound transmission itself. In our case, it is essential that this detailed information is available because we only have one sensor observation. Thus, instead of applying a standard geometric operation to multiple sensor observations, we attempt to explain the received signal as a linear combination of point scatterer echo signals (*23*, *24*). Without loss of generality, let us represent the unknown true image by a vector **v**, containing N pixels at a desired resolution. In our case, we use a pixel size smaller than two times the smallest wavelength detected (0.120 mm). We found that this discretization size was a good trade-off between model accuracy and memory usage. Furthermore, we assume that the data obtained by the measurement device are linearly related to the image **v** and can be modeled as

$$\mathbf{u} = \mathbf{Hv} + \mathbf{n}$$

where **u** is an M-dimensional vector containing the sampled pulse-echo signals, which are concatenated when using multiple rotations, such that M is equal to the number of time samples times the number of rotations. The M-dimensional vector **n** represents the additive noise. Note that in this model, the system matrix or so-called observation matrix **H** is determined by the measurement device (hardware) and contains all pulse-echo signals of all the pixels in **v**. For more details on the dimensions of **H**, a comparison to multisensor arrays, and the relation to CS, see text and figs. S2 and S3.

Having defined **u**, **H**, and **v**, and assuming that a linear propagation and scattering model holds, imaging can be accomplished using a wide variety of linear inversion methods. To demonstrate the system, we imaged two plastic letters that were positioned at a distance of 17 and 23 mm away from the sensor. Here, **H** contained pulse-echo signals for all pixels within a 3D volume containing the scattering objects and was
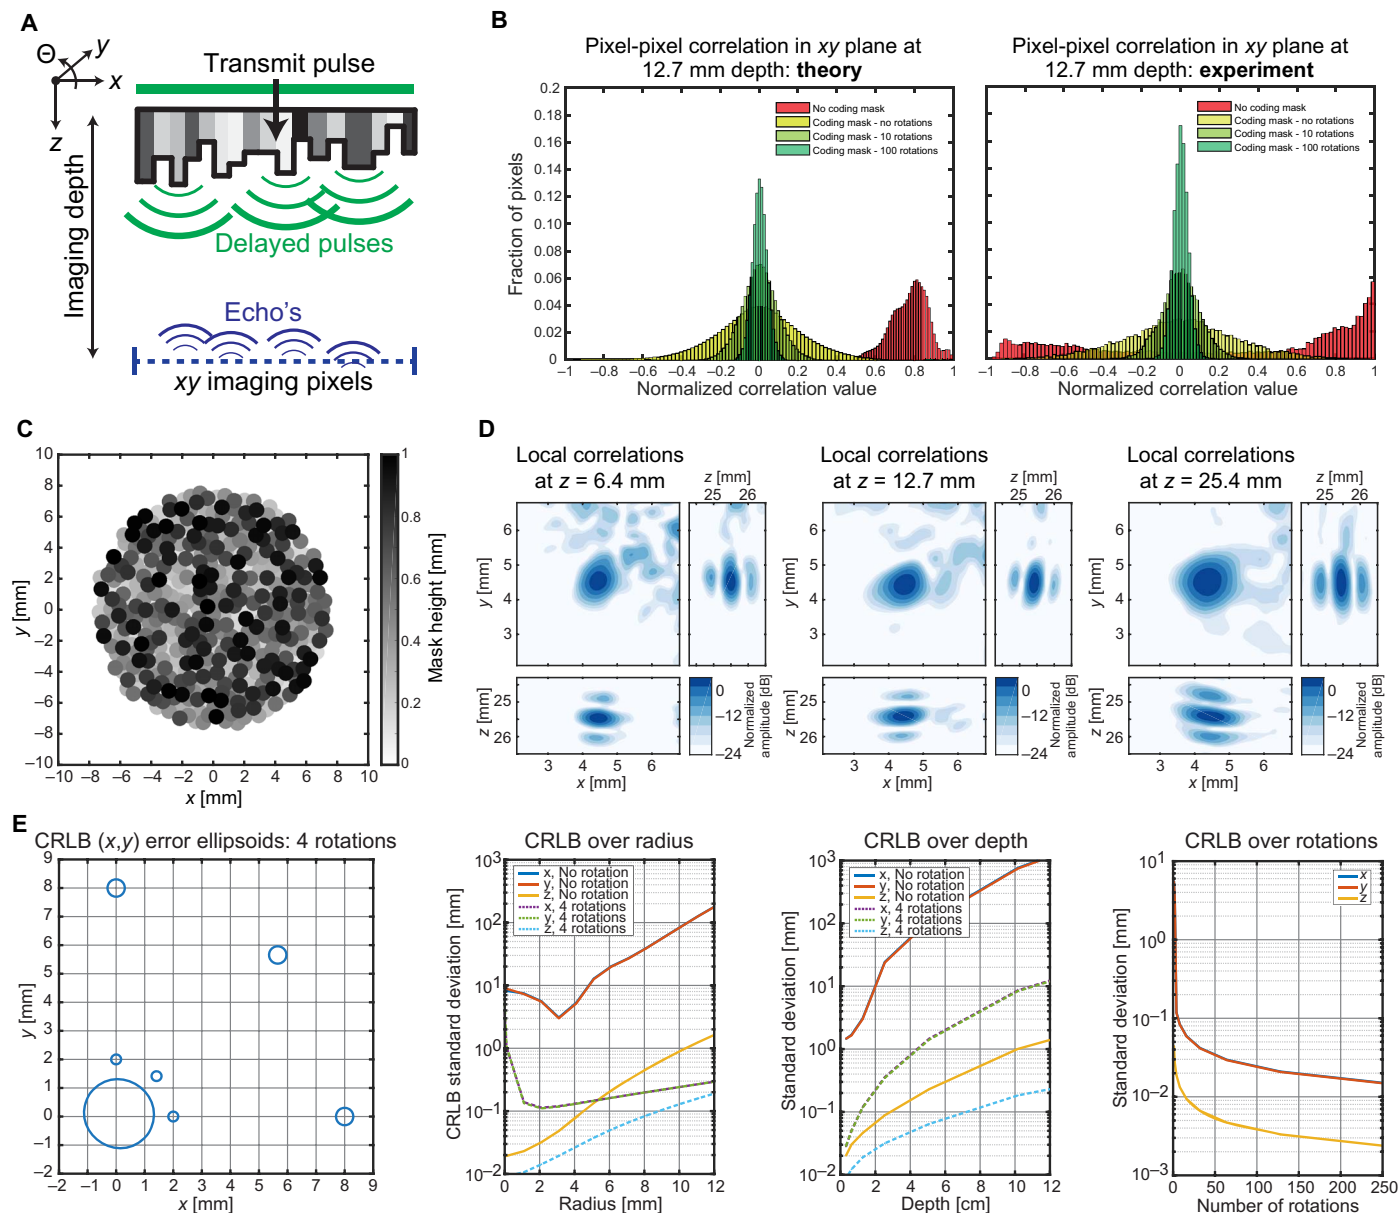
Fig. 5. Random delays in the coding mask generate signal variability for unique pixel reconstruction. (**A**) Schematic drawing of the approximated ultrasound field in a plane inside the medium parallel to the sensor surface. For every pixel in this plane, we want the echo signals to be as uncorrelated/orthogonal as possible. (**B**) The histograms show the cross-correlation values of a large set of pixels in the plane, obtained using an approximate simulation model. When the plastic delay mask is rotated, the cross-correlation values distribute closer to zero, which suggests better image reconstruction. (**C**) Mask layout where the gray values indicate the thickness variations, causing local delays of the ultrasound field. (**D**) Local correlations for a pixel at a fixed (x, y) position at several depths using 72 mask rotations. (**E**) CRLB analysis. The leftmost panel shows the level curves for one SD of position estimation error for a depth z of 12.7 mm. The center graphs show the SD in x, y, and z over radius (distance to the center of rotation) at depth 12.7 mm, and over depth for a pixel at a distance of 3 mm from the center. The rightmost graph shows how the CRLB changes as more measurements are added by rotation for a pixel at a distance of 3 mm from the center and at a depth of 12.7 mm.

constructed using the calibration method discussed in the next section. As seen in Fig. 6 (B, D, and F), both letters were successfully reconstructed, demonstrating the 3D imaging capabilities of the presented methods. For Fig. 6 (B and D), we imaged a small volume containing each letter and applied regularized least squares inversion using the LSQR algorithm (25) with limited iterations to regularize the problem, which is known to be equivalent to a filtered singular value decomposition inversion (26).

For the full 3D reconstruction shown in Fig. 6F, we made use of the sparsity of these two letters in water by applying the sparsity-promoting basis pursuit denoising (BPDN) algorithm (27). As can be seen, this prior knowledge about the image could be effectively exploited to improve image quality, significantly improving the dynamic range from 15 to 40 dB. Unlike in CS reconstructions, sparsity is not necessary for linear inversion due to the local correlations of the encoding. However, sparsity can be exploited to improve the reconstruction as discussed in fig. S3.
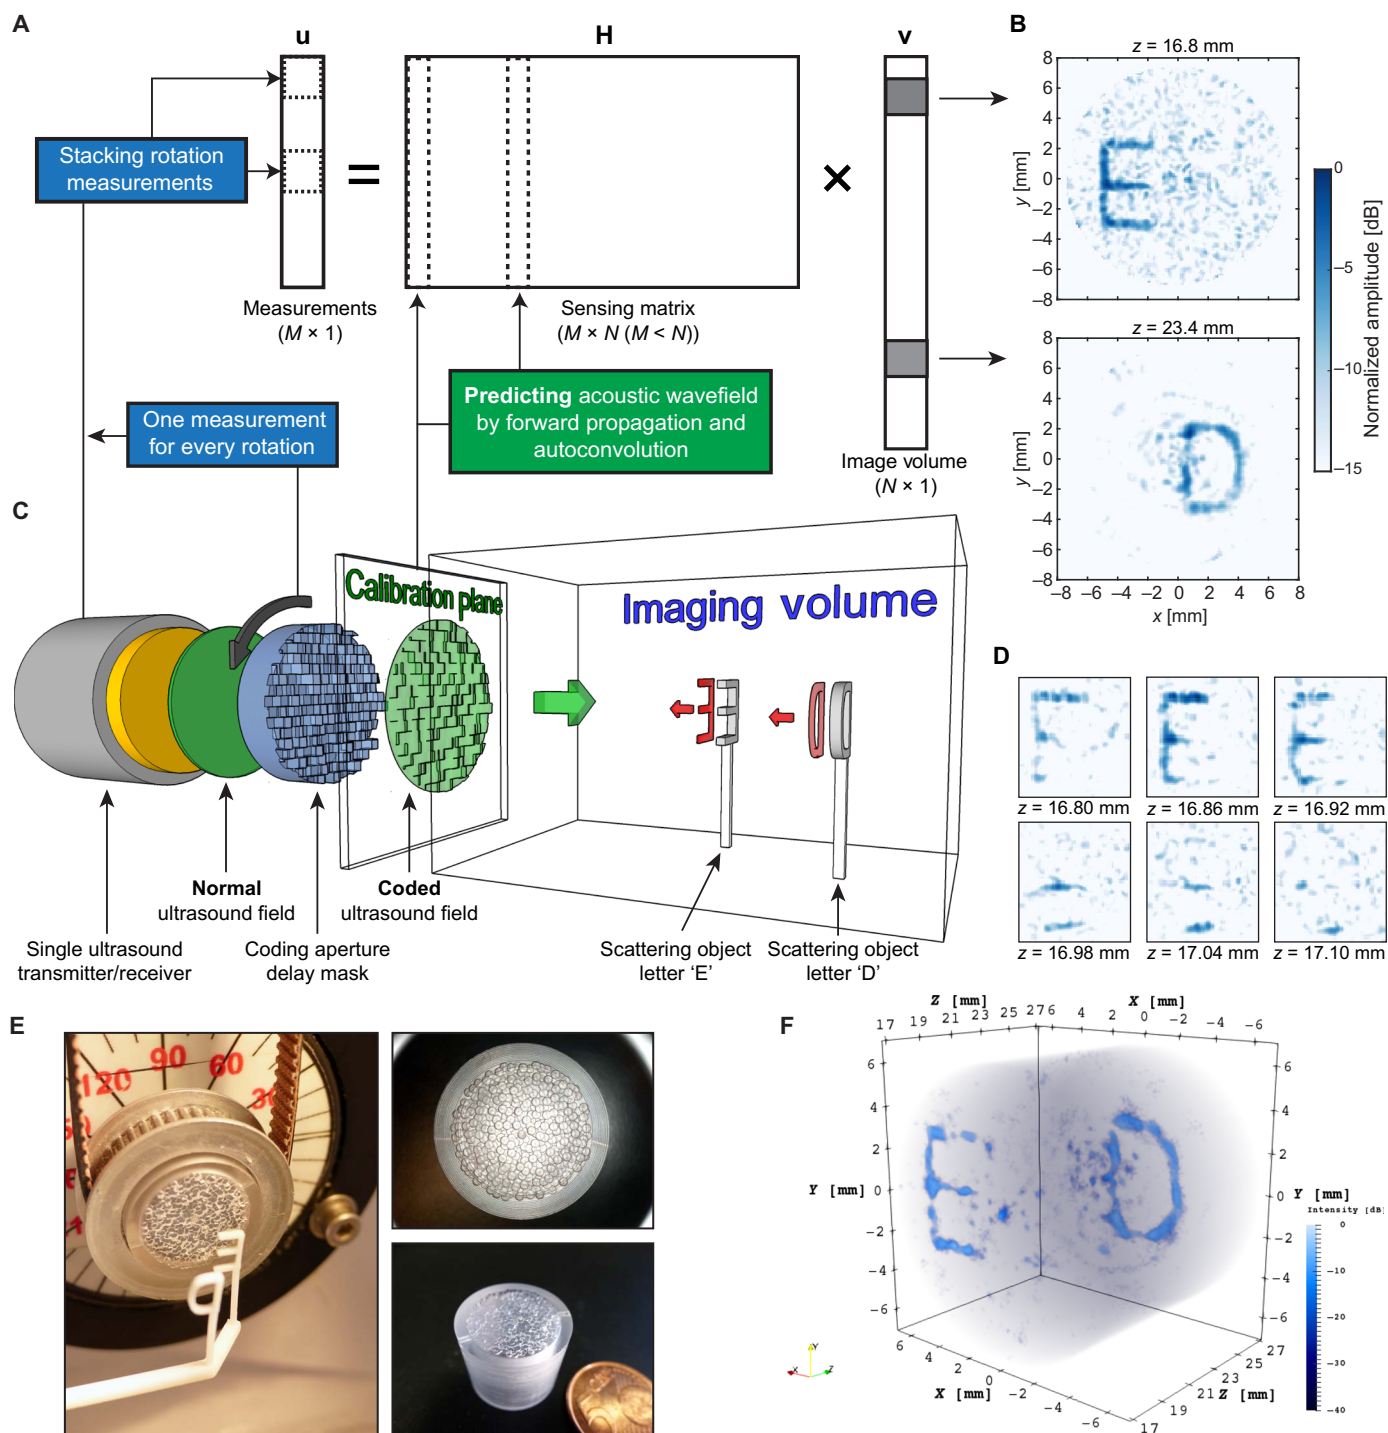
Fig. 6. Compressive 3D ultrasound imaging using a single sensor. (**A**) Schematic sketch of the signal model involved in this type of compressive imaging. Each column of the observation matrix **H** contains the ultrasound pulse-echo signal that is associated with a pixel in 3D space, which is contained in the image vector **v.** By rotating the coding mask in front of the sensor, we obtain new measurements that can be stacked as additional entries in the measurement vector **u** and additional rows in **H.** (**B**) Result of solving part of the image vector **v** using an iterative least squares technique. The two images are mean projections of six pixels along the z dimension [individual z slices shown in (D)]. (**C**) Schematic overview of the complete imaging setup. A single sensor transmits a phase uniform ultrasound wave through a coding mask that enables the object information (two plastic letters "E" and "D") to be compressed to a single measurement. Rotation of the mask enables additional measurements of the same object. (**D**) Reconstruction of the letter "E" in six adjacent z slices. A small tilt of the letter (from top left corner to bottom right corner) can be observed, demonstrating the potential 3D imaging capabilities of the proposed device. (**E**) Image showing the two 3D-printed letters and the plastic coding mask with a rubber band for rotating the mask over the sensor. The two right-hand panels show close-ups of the plastic coding mask. (**F**) 3D rendering of the complete reconstructed image vector **v**, obtained by BPDN. The images shown in (B) and (D) were obtained using 72 evenly spaced mask rotations, and the full 3D image in (F) was obtained using only 50 evenly spaced rotations to reduce the total matrix size.

Constructing the system matrix H

For computing the column signals of **H**, it is crucial to know the entire spatiotemporal wave field. To this end, we adopt a simple calibration measurement in which we spatially map the impulse signal (using a small hydrophone and a translation stage) in a plane close to the mask surface and perpendicular to the ultrasound propagation axis (Fig. 6C). This recorded wave field can then be propagated to any plane that is parallel to the recorded plane using the angular spectrum approach (28, 29). Note that this procedure is required only once and is unique for the specific sensor and coding mask. According to the reciprocity theorem, we are able to obtain the pulse-echo signal by performing an autoconvolution of the propagated hydrophone signal with itself (30). Using these components, we can then compute the pulse-echo ultrasound signal for every point in 3D space and subsequently populate the columns of **H**. The additional rotation of the coding mask can be regarded as a rotation of the ultrasound field inside the medium. Because every rotation provides a new measurement of the same image **v**, we can make use of this information by stacking these new signals as additional rows in the matrix **H**. The more independent rows we add to **H**, the better conditioned our image reconstruction problem becomes.

More formally, let $\mathbf{u}_i$ denote the ith sampled pulse-echo signal of length K, where K is the number of temporal samples obtained from one pulse-echo measurement and i indexes the various measurements obtained from different rotation angles. Let $\mathbf{H}_i$ denote the measurement matrix for rotation i of size $K \times N$, where N is the number of pixels. Then, if R rotations are performed, **u** and **H** are defined as follows: $\mathbf{u} = [\mathbf{u}_0^T \ \mathbf{u}_1^T \ \dots \ \mathbf{u}_{R-1}^T]^T$ and $\mathbf{H} = [\mathbf{H}_0^T \ \mathbf{H}_1^T \ \dots \ \mathbf{H}_{R-1}^T]^T$, where **u** is length M, with $M = KR$, and **H** has size $M \times N$. Furthermore, because each $\mathbf{H}_i$ is obtained by propagating and rotating the same calibration measurement (see "Acoustic calibration procedure" section in the Materials and Methods), each $\mathbf{H}_i$ can be computed from a reference measurement matrix $\tilde{\mathbf{H}}$: $\mathbf{H}_i = f(\tilde{\mathbf{H}}, \phi_i)$, where ϕ_i is the angle of rotation corresponding to measurement i, and the function $f(\mathbf{H}, \phi_i)$ computes the matrix $\mathbf{H}_i$ from the reference matrix $\tilde{\mathbf{H}}$ and input angle ϕ_i by appropriately rotating the acoustic field implicitly stored in $\tilde{\mathbf{H}}$. Consequently, **H** is expressed as $\mathbf{H} = [f(\tilde{\mathbf{H}}, \phi_1)^T \ f(\tilde{\mathbf{H}}, \phi_2)^T \ \dots \ f(\tilde{\mathbf{H}}, \phi_{R-1})^T]^T$.

DISCUSSION

We have demonstrated that compressive 3D ultrasound imaging can be successfully performed using just a single sensor and a simple aperture coding mask, instead of the thousand or more sensors conventionally required. The plastic mask breaks the phase uniformity of the ultrasound field and causes every pixel to be uniquely identifiable within the received signal. The compressed measurement contains the superimposed signals from all object pixels. The objects can be resolved using a regularized least squares algorithm, as we have proved here by successfully imaging two letters ("E" and "D"). Furthermore, we have provided an in-depth analysis of the qualities of our coding aperture mask for compressive ultrasound imaging. We believe that this technique will pave the way for an entirely new means of imaging in which the complexity is shifted away from the hardware side and toward the power of computing.

The calibration procedure that we use and our imaging device both share common ideas with the time-reversal work conducted by Fink and co-workers (30–34). Time-reversal ultrasound entails the notion that any ultrasound field can be focused back to its source by re-emitting the recorded signal back into the same medium. Consequently, the ultrasonic waves can be focused onto a particular point

in space and time if the impulse signal of that point is known. Using a reverberant cavity to create impulse signal diversity and a limited number of sensors, Fink et al. have shown that a 3D medium can be imaged by applying transmit focusing with respect to every spatial location. Hence, they need as many measurements as there are pixels, resulting in an unrealistic scenario for real-time imaging. Our work differs in the fact that rather than focusing our ultrasound waves to a single point, we encode the whole 3D volume onto one spatiotemporal varying ultrasound field, from which we then reconstruct an image based on solving a linear signal model. This has the advantage of requiring only a few measurements to image an entire 3D volume, in contrast to time-reversal work. This difference in acquisition time is crucial for medical imaging where the object changes over time, for example, in a beating heart. Another important distinction is the small impedance difference between our coding mask and imaging medium (less than a factor of 2). As a consequence, there is very little loss of acoustic energy as compared to reverberant cavities that inherently require high impedance mismatch to ensure signal diversity (33, 34).

Besides the time-reversal work, our technique also shares some common ground with the 2D localization work by Clement et al. (35–37). They successfully showed that frequency-dependent wave field patterns can be used to localize two point scatterers in a 2D plane using only one sensor in receive. Instead of solving a linear system as we propose here, the authors used a dictionary containing measured impulse responses and a cross-correlation technique to find the two point scatterers. By contrast, we propose to induce signal diversity using local delays in both transmit and receive. This type of "wave field coding" can be applied to many other imaging techniques and seems to offer a much higher dynamic range than the slowly varying frequency-dependent wave fields. As a result, we are able to move beyond the localization of isolated point scatterers but perform actual 3D imaging as we show in this paper.

The findings reported in this paper demonstrate that compressive imaging in ultrasound using a coded aperture is possible and practically feasible. Our device has disadvantages and limitations and does not deliver similar functionality as existing 3D ultrasound arrays. For example, we cannot focus the transmission beam, making techniques such as tissue harmonic imaging that require high acoustic power less accessible (37). In addition, the mechanical rotation of the mask is expected to introduce a fair amount of error in the prediction of the ultrasound field and limits the usability when it comes down to imaging moving tissue or the application of ultrasound Doppler. Besides, a rotating wave field becomes less effective at pixels closer to the center of rotation—assuming a spatially uniform generation of spatial frequencies. Future systems should consider the incorporation of other techniques to introduce signal variation, such as controlled linearized motion, or a mask that can be controlled electronically, similar to the digital mirror devices used in optical compressive imaging (13). Furthermore, the angular spectrum approach method that predicts the ultrasound field inside the medium is not flawless in cases where the medium contains strong variations in sound speed and density. These issues may possibly be solved, however, by incorporating more advanced ultrasound field prediction tools that can deal with these kinds of variations (38). This will increase the computational burden of the image reconstruction even further but may potentially lead to better results than what we have shown here.

One of the fundamental ideas of this paper is to prove that spatial wave compression allows for compressive imaging. Hence, this technique is applicable not only to ultrasound but also to many other imaging

and localization techniques that normally rely on spatial sampling of coherent wave fields. With respect to clinical ultrasound, we believe that the proposed technique may prove to be useful in cases where signal reduction is beneficial, for example, in minimally invasive imaging catheters. Furthermore, we think that a coded aperture mask could potentially be used in conjunction with conventional 1D ultrasound arrays to better estimate the out-of-plane signals and possibly extend the normal 2D imaging capabilities to 3D.

Alternative coded aperture implementations for ultrasound compressive imaging are certainly possible. However, the range of the physical parameters that we can vary to create a useful coded aperture for ultrasound seems less diverse compared to optics. The properties of light are such that it can easily be blocked very locally (for example, by absorption), which is hardly possible with ultrasound waves because of their mechanical nature and larger wavelengths. The idea of applying local delays in the ultrasound field to uniquely address every pixel therefore seems more practical. Ultimately, it will be these kinds of physical parameters that will dictate whether and to what extent true compressive imaging is possible. Restricting ourselves to local time delays, we have successfully shown in this paper that the physics of ultrasound allow for the ultimate compression of a 3D volume in a few measurements with a single sensor and a simple coding mask.

MATERIALS AND METHODS

Mask manufacturing

Our aim was to make a spatiotemporally varying ultrasound field by locally introducing phase shifts in the wave transmitted by the sensor. We did so by introducing a mask that delays the wave locally, similar to the function of a phase spatial light modulator in optics (39). The aberration mask that we put in front of the sensor was made from an 11-mm-thin circular plastic layer (Perspex), in which we had drilled holes 1 mm in diameter (Fig. 5C). The depth and positions of these holes were randomly chosen by an algorithm, allowing for a depth ranging between 0.1 and 1 mm and for a maximum of 0.5-mm overlap between the holes. The holes were drilled using a computer numeric control machine (P60 HSC, Fehlmann). The choice for 1-mm holes was made because of convenience of manufacturing. At room temperature, the speed of sound propagation in water is approximately 1480 m/s, whereas the speed of sound in the Perspex material is 2750 m/s. Perspex has good properties in terms of processing, high sound speed, and low acoustic absorption.

Acoustic setup

For the transmission and reception of ultrasonic waves, we used one large sensor of piezocomposite material that had an active diameter of 12.7 mm and a center frequency of 5 MHz (C309-SU, Olympus Panametrics NDT Inc.). The sensor was coupled to a pulser/receiver box (5077PR, Olympus Panametrics NDT Inc.) that can generate a short −100 V transmit spike and amplifies signals in receive up to 59dB. No filtering was applied in receive. The sensor with the mask was mounted in a tank filled with water measuring 200 × 300 × 200 mm. In the water, we placed two plastic letters that had been 3D-printed (Objet30, Stratasys Ltd.). To enhance the ultrasound reflection from these letters, we sprayed only the letters with silicon carbide powder (size, k800). The two letters "E" and "D" refer to the names of our academic institutions, Erasmus Medical Center and Delft University of Technology.

The thin plastic coding mask was part of a larger round casing that fitted snugly around the sensor (Fig. 6E). To this casing, we fitted a cogwheel that allows the casing to be rotated in front of the sensor using a small rubber band and an electronic rotary stage (T-RS60A, Zaber Technologies). By rotating the aberration mask, we were essentially able to rotate the interference pattern inside the stationary medium, allowing for multiple observations of the same object. We allowed a thin water film between the plastic mask and the sensor surface to ensure optimal coupling between the piezo material and the plastic mask. With this setup, it took about 1 min to acquire the signals needed to reconstruct the 3D volume shown in Fig. 6.

Acoustic calibration procedure

Before we could use our device for imaging, we needed to know the signature of our interference pattern. That is, we needed to know what the transmission and receive wave fields looked like in both space and time, so that we could populate our system matrix $\mathbf{H}$. Note that for every spatial point, we had a unique temporal ultrasound signal. This calibration step is an essential component in this kind of compressive imaging (40). The better $\mathbf{H}$ is known, the better the reconstruction of the image. For our calibration procedure, we recorded the transmission wave field in a plane perpendicular to the ultrasound propagation axis using a 3D translation stage (BiSlide MN10-0100-M02-21, Velmex) and a broadband needle hydrophone (0.5 to 20 MHz, 0.2-mm diameter; Precision Acoustics) coupled to a programmable analog-to-digital converter (DP310, Acqiris USA) using 12 bits per sample and a 200-MHz sampling rate.

For the results presented in this paper, we recorded the wave field in a 30 × 30 mm plane located 2.5 mm in front of the aberration mask using a step size of 0.12 mm (less than two times the highest spatial frequency detected) in both the x and y directions (Figs. 3 and 4). The signal of every pixel was obtained using the angular spectrum approach to address pixels at greater depth, using autoconvolution to account for the pulse-echo acquisition, and using linear interpolation to account for additional rotation measurements; all of which were implemented in MATLAB software, release 2015b (The MathWorks). We noticed that our sensor was not completely axisymmetric, meaning that the mask wave interaction changes slightly upon mask rotation. This issue was solved by measuring four rotated planes (0°, 90°, 180°, and 270°). We generated all intermediate planes by linear interpolation between each pair of neighboring planes. The calibration procedure took around 1.5 hours, after which the complete spatiotemporal impulse response of the sensor was known. Note that this calibration is sensor- and mask-specific and, in principle, only needs to be done once.

Approximate model and CRLB

To analyze the CRLB for the estimation of a single scatterer from one or more pulse-echo measurements, we derived the expressions for computing the Fisher information matrix, the inverse of which was computed numerically to obtain the CRLB covariance matrix. The following assumptions and conventions were used:

(i) Any noise or other errors can be modeled as zero mean white Gaussian additive noise, with variance σ^2. Hence, this analysis is only valid if noise sources, such as electronic, quantization, or jitter noise, can be modeled as Gaussian random variables, and their total effect is compounded into the signal-to-noise figure.

(ii) The ultrasound field due to the coding mask can be approximated as a collection of point sources, where the point sources are located at the surface of the mask. The point in time at which each virtual point source is excited is determined by the time of arrival of the ultrasound field from the true sensor location through the mask to the end of the mask.

(iii) The impulse response for the main sensor is taken as being equal to a Gaussian pulse: $h(t) = \exp[-\frac{t^2}{2l^2}]$, where l determines the pulse width or frequency bandwidth.

(iv) The origin of the spatial axes corresponds to the center of the main sensor surface, with the sensor lying in the xy plane and the z axis perpendicular to the sensor surface.

(v) The sensor is excited by a single Dirac delta pulse.

(vi) The position of the virtual source j is denoted as $\mathbf{s}_j \in \mathbb{R}^3$, and the set of all $\mathbf{s}_j$ is denoted as S. Vector variables are indicated by bold typesetting.

(vii) The ith element of a vector $\mathbf{k}$ is denoted as $[\mathbf{k}]_i$.

Under the assumptions above, the forward pressure wave field measured at position $\boldsymbol{\theta} = [x, y, z]^T$ can be written as

$$p(t, \boldsymbol{\theta}) = \sum_{j=0}^{|S|-1} h(t - \tau(\boldsymbol{\theta}, \mathbf{s}_j))$$

where $\tau(\boldsymbol{\theta}, \mathbf{s}_j) = \frac{1}{c_1}[\mathbf{s}_j]_z + \frac{1}{c_0}||\mathbf{s}_j - \boldsymbol{\theta}||_2$. This can be regarded as the convolution of the sensor impulse response with the mask impulse response: $p(t, \boldsymbol{\theta}) = h(t) * m(t, \boldsymbol{\theta}, S)$. The echo signal can then be expressed as $a(t, \boldsymbol{\theta}) = h(t) * m(t, \boldsymbol{\theta}, S) * m(t, \boldsymbol{\theta}, S) * h(t)$, or simply $a(t, \boldsymbol{\theta}) = p(t, \boldsymbol{\theta}) * p(t, \boldsymbol{\theta})$, which is equal to the autoconvolution of the forward pressure wave field at $\boldsymbol{\theta}$. Under the white Gaussian noise assumption, it can be shown that the (i, j)-th entry of the Fisher information matrix $I(\boldsymbol{\theta}) \in \mathbb{R}^{3x3}$ is equal to

$$I(\boldsymbol{\theta})_{i,j} = \frac{1}{\sigma^2} \frac{\partial \mathbf{a}(\boldsymbol{\theta})^T}{\partial [\boldsymbol{\theta}]_i} \frac{\partial a(\boldsymbol{\theta})}{\partial [\boldsymbol{\theta}]_j}$$

where $\mathbf{a}(\boldsymbol{\theta})$ is a vector representing the discretized signal $a(t, \boldsymbol{\theta})$. Next, we derived the derivative of $\mathbf{a}(\boldsymbol{\theta})$. We started by computing the derivative of $a(t, \boldsymbol{\theta})$

$$\frac{d}{d[\boldsymbol{\theta}]_i} a(t, \boldsymbol{\theta}) = \frac{d}{d[\boldsymbol{\theta}]_i} p(t, \boldsymbol{\theta}) * p(t, \boldsymbol{\theta}) = \frac{d}{d[\boldsymbol{\theta}]_i} \int_{-\infty}^{\infty} p(t - T, \boldsymbol{\theta}) p(T, \boldsymbol{\theta}) dT$$

$$= \int_{-\infty}^{\infty} \frac{d}{d[\boldsymbol{\theta}]_i} p(t - T, \boldsymbol{\theta}) p(T, \boldsymbol{\theta}) dT = 2 p(t, \boldsymbol{\theta}) * \frac{d}{d[\boldsymbol{\theta}]_i} p(t, \boldsymbol{\theta})$$

Using the definition of $p(t, \boldsymbol{\theta})$ and $\tau(\boldsymbol{\theta}, \mathbf{s}_j)$, we find that

$$\frac{d}{d[\boldsymbol{\theta}]_i} p(t, \boldsymbol{\theta}) = \sum_{j=0}^{|S|-1} \frac{[\boldsymbol{\theta}]_i - [\mathbf{s}_j]_i}{||\boldsymbol{\theta} - \mathbf{s}_j||_2^2} h(t - \tau(\boldsymbol{\theta}, \mathbf{s}_j))$$

$$\times \left[\frac{1}{c_0 l^2} (t - \tau(\boldsymbol{\theta}, \mathbf{s}_j)) - \frac{1}{||\boldsymbol{\theta} - \mathbf{s}_j||_2} \right]$$

which can be discretized to obtain $d\mathbf{a}(\boldsymbol{\theta})/d[\boldsymbol{\theta}]_i$ for a single measurement. From $I(\boldsymbol{\theta})$, the CRLB is found as $I(\boldsymbol{\theta})^{-1}$.

Image reconstruction

Here, we formalized our image reconstruction problem in a set of linear equations $\mathbf{u} = \mathbf{Hv} + \mathbf{n}$. This formalization allowed for a multitude of solvers to be applied to this problem. We applied two inversion methods to find $\mathbf{v}$ from $\mathbf{u}$. The first was the least squares estimate,

which finds the image $\mathbf{v}$ that minimizes the squared error between the model $\mathbf{Hv}$ and observation $\mathbf{u}$ as follows: $\hat{\mathbf{v}} = \mathrm{argmin}_{\mathbf{v}'} ||\mathbf{u} - \mathbf{Hv}'||_2^2$. This estimate is known to be the minimum variance unbiased estimator if the linear model is valid, and the noise $\mathbf{n}$ is the white Gaussian. To implement this estimator, we used the LSQR algorithm (25), which is especially appropriate for large-scale sparse systems. By limiting the number of iterations, one can regularize the reconstruction problem [see, for example, the study by Hansen (26)]; this also allows for balancing of the resolution, the noise robustness, or the contrast.

For the reconstructed letters shown in Fig. 6 (B and D), we used about 15 iterations, which took several seconds to compute. The results did not seem to vary much between 7 and 40 iterations. Above 50 iterations, we observed a serious loss in contrast as the algorithm starts to fit the solution to the noise and errors present in our $\mathbf{H}$ matrix. For the complete reconstruction of the 3D volume shown in Fig. 6F, we applied a BPDN algorithm (27), exploiting the sparsity of the image because we were only interested in two reflecting letters in an otherwise empty medium (water). For this purpose, we chose to use the iterative two-step shrinkage/thresholding algorithm, named TwIST, using l_1-norm regularization (41) for implementation, although other algorithms that are capable of l_1-based optimization are also viable. By this approach, we minimized the BPDN cost-function $\hat{\mathbf{v}} = \mathrm{argmin}_{\mathbf{v}'} \frac{1}{2} ||\mathbf{u} - \mathbf{Hv}'||_2^2 + \lambda ||\mathbf{v}'||_1$, where λ is a dimensionless scalar that weights the amount of regularization. The BPDN algorithm regularization parameter λ was chosen by defining a region of interest (ROI) containing the reconstructed letters and an ROI outside the letters and computed the contrast ratio as $20 \log_{10}(I_{\mathrm{letter}}/I_{\mathrm{background}})$ for a range of λ values, where I_{letter} or $I_{\mathrm{background}}$ are the average values of $|\mathbf{v}|$ over the pixels in the corresponding ROI. We then chose the λ that offered the highest contrast ratio. Using this procedure, we observed a contrast of 9 dB for regularized least squares and 29 dB for the l_1 regularization. For this proof-of-principle study, it shows that selecting a suitable λ is possible in a similar way as in real-life clinical ultrasound imaging, where the operator tunes some critical image reconstruction parameters (such as speed of sound) by visual feedback on image quality. In further research, it should be investigated if λ should be tuned for image content or a globally optimal λ can be established. If the latter is the case, an optimal λ should be established over large sets of training images, and its generality tested on independent test images. The 3D rendering of the letters shown in Fig. 6F were performed using the open-source visualization tool ParaView (www.paraview.org) (42).

SUPPLEMENTARY MATERIALS

REFERENCES AND NOTES

1. K. K. Shung, M. Zippuro, Ultrasonic transducers and arrays. *IEEE Eng. Med. Biol. Mag.* **15**, 20–30 (1996).

2. K. E. Thomenius, Evolution of ultrasound beamformers. in *1996 IEEE Ultrasonics Symposium Proceedings* (IEEE, 1996), pp. 1615–1622.

3. A. Fenster, D. B. Downey, H. N. Cardinal, Three-dimensional ultrasound imaging. *Phys. Med. Biol.* **46**, R67 (2001).

4. E. Merz, 25 Years of 3D ultrasound in prenatal diagnosis (1989–2014). *Ultraschall. Med.* **36**, 3–8 (2015).

5. E. J. Candès, J. Romberg, T. Tao, Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theory* **52**, 489–509 (2006).

6. D. L. Donoho, Compressed sensing. *IEEE Trans. Inf. Theory* **52**, 1289–1306 (2006).

7. C. E. Shannon, Communication in the presence of noise. *PIRE* **37**, 10–21 (1949).

8. R. G. Baraniuk, Compressive sensing [Lecture Notes]. *IEEE Signal Process. Mag.* **24**, 118–121 (2007).

9. N. Wagner, Y. C. Eldar, Z. Friedman, Compressed beamforming in ultrasound imaging. *IEEE Trans. Signal Process.* **60**, 4643–4657 (2012).

10. H. Liebgott, A. Basarab, D. Kouame, O. Bernard, D. Friboulet, Compressive sensing in medical ultrasound, in *2012 IEEE International Ultrasonics Symposium,* (IEEE, 2012), pp. 1–6.

11. C. Quinsac, A. Basarab, D. Kouamé, Frequency domain compressive sampling for ultrasound imaging. *Adv. Acoust. Vib.* **2012**, 231317 (2012).

12. M. F. Schiffner, G. Schmitz, Fast pulse-echo ultrasound imaging employing compressive sensing in *Proceedings of the IEEE International Ultrasonics Symposium (IUS)* (IEEE, 2011), pp 688–691.

13. M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, R. G. Baraniuk, Single-pixel imaging via compressive sampling. *IEEE Signal Process. Mag.* **25**, 83–91 (2008).

14. L. Wang, L. Li, Y. Li, H. C. Zhang, T. J. Cui, Single-shot and single-sensor high/super-resolution microwave imaging based on metasurface. *Sci. Rep.* **6**, 26959 (2016).

15. D. Shrekenhamer, C. M. Watts, W. J. Padilla, Terahertz single pixel imaging with an optically controlled dynamic spatial light modulator. *Opt. Express* **21**, 12507–12518 (2013).

16. N Huynh, E. Zhang, M. Betcke, S. R. Arridge, P. Beard, B. Cox, A real-time ultrasonic field mapping system using a Fabry Pérot single pixel camera for 3D photoacoustic imaging, in *Proceedings of SPIE Volume 9323,* (International Society for Optics and Photonics, 2015), pp. 93231O.

17. E. Caroli, J. B. Stephen, G. Di Cocco, L. Natalucci, A. Spizzichino Coded aperture imaging in X-and gamma-ray astronomy. *Space Sci. Rev.* **45**, 349–403 (1987).

18. R. F. Marcia, Z. T. Harmany, R. M. Willett, Compressive coded aperture imaging, in *Proceedings of SPIE Volume 7246,* (International Society for Optics and Photonics, 2009), pp. 72460G.

19. X. Peng, G. J. Ruane, G. A. Swartzlander Jr., Randomized aperture imaging. https://arxiv.org/abs/1601.00033 (2016).

20. K. Melde, A. G. Mark, T. Qiu, P. Fischer, Holograms for acoustics. *Nature* **537**, 518–522 (2016).

21. M. D. Brown, D. I. Nikitichev, B. E. Treeby, B. T. Cox, Generating arbitrary ultrasound fields with tailored optoacoustic surface profiles. *Appl. Phys. Lett.* **110**, 094102 (2017).

22. S. M. Kay, *Fundamentals of statistical signal processing: Estimation Theory*, (PTR Prentice-Hall, 1993).

23. R. Lavarello, F. Kamalabadi, W. D. O'Brien, A regularized inverse approach to ultrasonic pulse-echo imaging. *IEEE Trans. Med. Imaging* **25**, 712–722 (2006).

24. R. Stoughton, S. Strait, Source imaging with minimum mean-squared error. *J. Acoust. Soc. Am.* **94**, 827–834 (1993).

25. C. C. Paige, M. A. Saunders, LSQR: An algorithm for sparse linear equations and sparse least squares. *ACM Trans. Math. Softw.* **8**, 43–71 (1982).

26. P. C. Hansen, *Discrete Inverse Problems: Insight and Algorithms* (SIAM, 2010).

27. S. S. Chen, D. L. Donoho, M. A. Saunders, Atomic decomposition by basis pursuit. *SIAM Rev. Soc. Ind. Appl. Math.* **43**, 129–159 (2001).

28. Y. Du, H. Jensen, J. A. Jensen, Investigation of an angular spectrum approach for pulsed ultrasound fields. *Ultrasonics* **53**, 1185–1191 (2013).

29. P. R. Stepanishen, K. C. Benjamin, Forward and backward projection of acoustic fields using FFT methods. *J. Acoust. Soc. Am.* **71**, 803–812 (1982).

30. M. Fink, Time reversal of ultrasonic fields. I. Basic principles. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **39**, 555–566 (1992).

31. G. Montaldo, D. Palacio, M. Tanter, M. Fink, Building three-dimensional images using a time-reversal chaotic cavity. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **52**, 1489–1497 (2005).

32. C. Draeger, J.-C. Aime, M. Fink, One-channel time-reversal in chaotic cavities: Experimental results. *J. Acoust. Soc. Am.* **105**, 618–625 (1999).

33. N. Etaix, M. Fink, R. K. Ing, Acoustic imaging device with one transducer. *J. Acoust. Soc. Am.* **131**, EL395–EL399 (2012).

34. G. Montaldo, D. Palacio, M. Tanter, M. Fink, Time reversal kaleidoscope: A smart transducer for three-dimensional ultrasonic imaging. *Appl. Phys. Lett.* **84**, 3879–3881 (2004).

35. P. J. White, G. T. Clement, Two-dimensional localization with a single diffuse ultrasound field excitation. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **54**, 2309–2317 (2007).

36. F. C. Meral, M. A. Jafferji, P. J. White, G. T. Clement, Two-dimensional image reconstruction with spectrally-randomized ultrasound signals. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **60**, 2501–2510 (2013).

37. G. T. Clement, Two-dimensional ultrasound detection with unfocused frequency-randomized signals. *J. Acoust. Soc. Am.* **121**, 636–647 (2007).

38. B. E. Treeby, B. T. Cox, k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave fields. *J. Biomed. Opt.* **15**, 021314 (2010).

39. N. Savage, Digital spatial light modulators. *Nat. Photonics* **3**, 170–172 (2009).

40. A. Liutkus, D. Martina, S. Popoff, G. Chardon, O. Katz, G. Lerosey, S. Gigan, L. Daudet, I. Carron, Imaging with nature: Compressive imaging using a multiply scattering medium. *Sci. Rep.* **4**, 5552 (2014).

41. J. M. Bioucas-Dias, M. A. T. Figueiredo, A new TwIST: two-step iterative shrinkage/thresholding algorithms for image restoration. *IEEE Trans. Image Process.* **16**, 2992–3004 (2007).

42. J. Ahrens, B. Geveci, C. Law, C. Hansen, C. Johnson, 36-ParaView: An End-User Tool for Large-Data Visualization, in *The Visualization Handbook*, (Elsevier, 2005), pp. 717–731.

43. E. J. Candes, T. Tao, Decoding by linear programming. *IEEE Trans. Inf. Theory* **51**, 4203–4215 (2005).