



OPEN

DATA DESCRIPTOR

An in-air synthetic aperture sonar dataset of target scattering in environments of varying complexity

Thomas E. Blanford¹✉, David P. Williams², J. Daniel Park², Brian T. Reinhardt², Kyle S. Dalton³, Shawn F. Johnson² & Daniel C. Brown³

This paper describes a synthetic aperture sonar (SAS) dataset collected in-air consisting of four types of targets in four environments of different complexity. The in-air laboratory based experiments produced data with a level of fidelity and ground truth accuracy that is not easily attainable in data collected underwater. The range of complexity, high level of data fidelity, and accurate ground truth provides a rich dataset with acoustic features on multiple scales. It can be used to develop new signal-processing and image reconstruction algorithms, as well as machine learning models for object detection and classification. It may also find application in model verification and validation for acoustic simulators. The dataset consists of raw acoustic time series returns, associated environmental conditions, hardware configuration, array motion, as well as the reconstructed imagery.

Background & Summary

Synthetic aperture sonar (SAS) is a coherent acoustic remote-sensing technique that is typically used to produce high resolution images of objects in underwater environments¹. It is conceptually similar to synthetic aperture radar (SAR)², computed tomography (CT)³, magnetic resonance imaging (MRI)⁴, and seismic migration imaging⁵ techniques used in other domains. SAS arrays are mounted to a moving platform, transmit pulses at regular intervals, and record the backscattered echoes on one or more receivers. Using the estimated motion of the platform and an estimate of the local sound speed, these echoes are reconstructed into imagery that is a spatial map of the acoustic reflectivity of the scene. Although photograph like in appearance, SAS imagery is generally collected with a low grazing angle geometry and the data contain significant physical complexity. Targets have complicated scattering behavior that depends on their shape and material composition. They often lie proud or partially buried on seafloors with multiple scales of roughness and inhomogeneous acoustic properties. Both the raw echoes and the reconstructed imagery contain information about the targets, the environment, and how the two interact with acoustic ensonification. Other sources of high resolution underwater imagery, such as those collected from multibeam echosounders or airborne laser sources, are typically collected from normal incidence to the seafloor and have resolution that depends on range^{6–8}.

Underwater SAS data typically lack the accurate ground truth necessary to isolate features related to specific acoustic phenomena and may be confounded by deleterious factors such as uncompensated motion and noise. Supervised machine learning algorithms such as convolutional neural networks (CNNs) are increasingly being employed for automated detection and classification of targets in SAS imagery^{9,10}. Their development is hindered, however, by small datasets¹¹, lack of ground truth, and class imbalance¹². Phenomena of interest often occur in the tails of the distribution of large underwater datasets, but there are often insufficient samples with high enough fidelity to train networks to exploit these specific features^{13,14}. The motivating goal of this experiment was to produce tightly controlled SAS data with multiple types of acoustic interactions and degrees of complexity to use in a study of CNN performance.

¹Center for Acoustics Research and Education, University of New Hampshire, Durham, NH, 03824, USA.

²Pennsylvania State University, University Park, PA, 16804, USA. ³Graduate Program in Acoustics, Pennsylvania State University, University Park, PA, 16804, USA. ✉e-mail: thomas.blanford@unh.edu

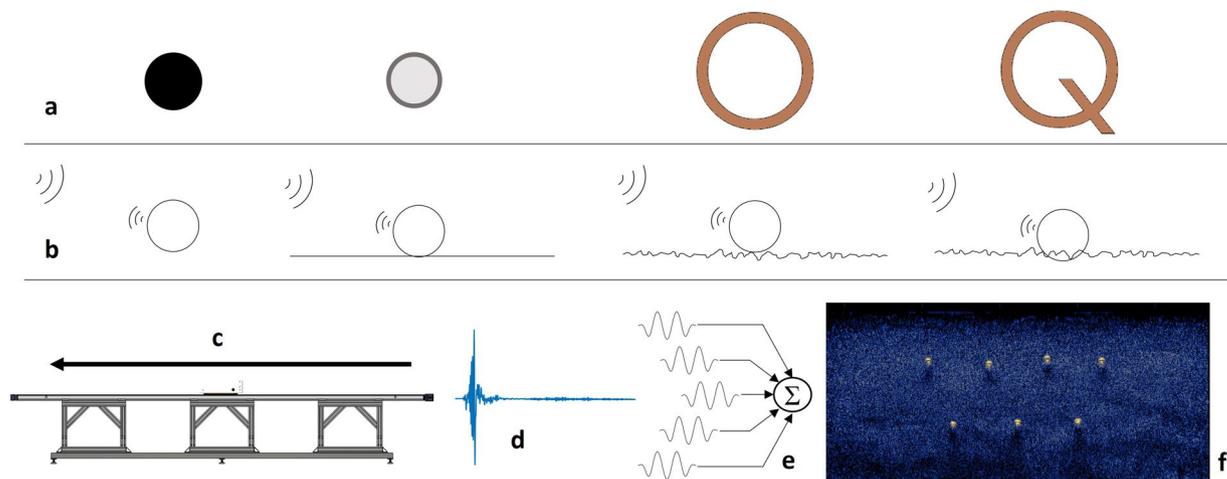


Fig. 1 The experiment collected tightly controlled acoustic scattering data, suitable for SAS reconstruction, from four target classes (a) in four environments (b). The targets were designed to have multiple discriminating features between classes, and the environments were designed to incrementally increase in complexity. An array of microphones and a loudspeaker were moved relative to the scene on a linear actuator (c), repeatedly transmitting a pulse and measuring the backscattered pressure (d). The resulting data were coherently reconstructed (e) into complex-valued imagery (f) that is a spatial representation of the estimated acoustic reflectivity of the scene.

This dataset consists of in-air SAS data of multiple types of targets and backgrounds collected in a controlled, quiet, indoor laboratory environment. It contains the raw acoustic signals and the associated non-acoustic data needed for image reconstruction, as well as the complex-valued SAS imagery. Natural underwater environments contain many degrees of complexity and can be difficult or impossible to control and accurately characterize. Laboratory experiments can be designed to invoke specific physical phenomena and allow for much greater control than field experiments. Furthermore, conducting controlled acoustic experiments is both substantially easier and less expensive in-air than underwater. The advantages afforded by in-air experimentation were intended to allow for accurate modeling and quantification of uncertainty in the data that would be infeasible underwater.

The experiment comprises four target classes (solid sphere, hollow sphere, block letter O, and block letter Q) in four classes of environment that increase in physical complexity (free-field, proud on a planar interface, proud on a rough interface, and partially buried in a rough interface). A loudspeaker and an array of four microphones were incrementally moved relative to the scene on a linear actuator. At each position along the actuator, the loudspeaker transmitted a short pulse and the backscattered echoes on each microphone were simultaneously sampled. The array motion and environmental properties were measured on a per-ping basis. The raw acoustic data, the non-acoustic data, and the reconstructed SAS imagery are all provided for each image collection. Additionally, regular acoustic characterization of the background environment and noise, as well as non-acoustic characterization of the targets and the instrumentation were conducted in order to aid data reuse. Figure 1 depicts a schematic overview of the experiment.

This dataset has potential reuse in the development and testing of SAS reconstruction algorithms (including interferometric processing)^{15,16}, automated detectors¹⁷, and classifiers^{18,19}. Additionally it may find use in experimental validation of acoustic models for both target²⁰ and rough interface scattering²¹.

Methods

The experiment was conducted from March 15, 2023 to April 28, 2023 in an indoor auditorium space. Acoustic scattering data was collected from four classes of targets in four types of background environments and was designed to allow synthetic aperture image reconstruction with the data. The data acquisition system consisted of commercial off-the-shelf (COTS) hardware with software automation.

Experimental design. Acoustic data were collected from an array consisting of a Peerless OX20SC00-04 loudspeaker (<https://products.peerless-audio.com/transducer/108>) with a 1.91 cm (0.75”) diameter diaphragm and four GRAS46AM microphones (<https://www.grasacoustics.com/products/measurement-microphone-sets/product/551-46am>). The array was mounted to the carriage of a Parker HPLA-080 5m’ linear actuator (<https://ph.parker.com/us/en/product-list/hpla-080-belt-driven-roller-wheel-rodless-linear-actuator>). Errors in estimating the phase of the recorded acoustic signals deteriorate the quality of the reconstructed image. Unwanted platform motion is the primary source of phase errors in SAS²². Moving the array with a precisely controlled linear actuator minimizes uncertainty in the sensor position that will degrade the image quality and eliminates the need for data-driven motion estimation that is commonly used in underwater SAS²³. The linear actuator was installed on the stage floor of the auditorium and its position within the space was unmoved during the course of the experiment. The actuator was leveled with a Bosch GLL2 laser line level to within 1.6 mm over the 5 m length. Rockwool covered in cotton fabric was placed over the supports of the linear actuator in order to minimize acoustic scattering from the support structure. Acoustically absorbent open-cell foam was also placed behind

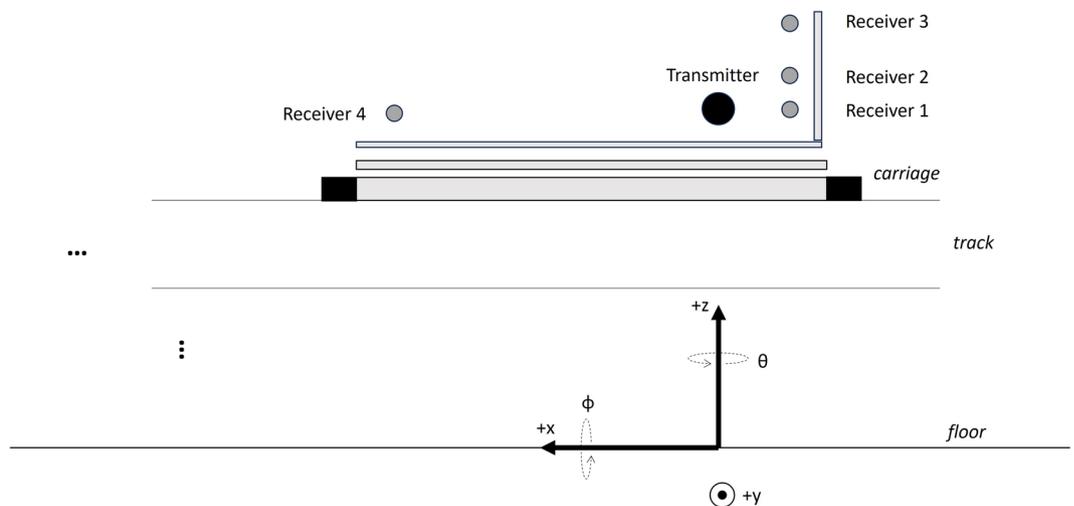


Fig. 2 The data collection geometry was defined in a right-handed Cartesian coordinate system with the origin on the floor directly below the loudspeaker when the array is in the home position. SAS data is typically referenced in a sensor-centric coordinate system. In this experiment, the x -axis corresponds to the along-track direction and the y -axis corresponds to the cross-track direction. The choice of origin in this coordinate system was for reliable placement and location of targets and backgrounds relative to the array.

the array and on the toe clamps holding the actuator to the supports to minimize unwanted scattering from the experimental infrastructure. The scenes for measurement consisted of targets and the background environments placed in front of the actuator.

The positions of the sensors and targets were defined in the Cartesian coordinate system illustrated in Fig. 2. An electromagnetic sensor was mounted to the actuator frame and used return the array to a fixed starting position (home). The weight of the actuator's support system ensured that the home position could not move relative to the floor during the course of the experiment. The origin of the coordinate system is the point on the stage floor directly below the center of the speaker with the array in the home position on the actuator. This choice of definition for the origin allowed for reliable placement of the targets and backgrounds throughout the experiment. Synthetic aperture sonar data is commonly referenced in a sensor-centric coordinate system of along-track (the principal direction of motion) and cross-track (range perpendicular to the array). In this experiment, the positive x -axis is the along-track direction and the positive y -axis is the cross-track direction.

The array geometry was designed to facilitate multiple types of processing and analysis with the data. Three microphones were arranged adjacent to the loudspeaker in a vertical line to allow for interferometric processing. The fourth microphone was located 45 cm away from the loudspeaker, forming a bistatic scattering geometry where targets in the scene are not in the far-field of this transmitter-receiver pair. This near-field sensing condition occurs in some types of underwater SAS systems²⁴. The coordinates of the transducers in the acoustic array were verified with a laser range finder and are reported in the data record. The microphones were oriented such that their faces were normal to the y -axis. The loudspeaker was mounted in a fixture with a depression angle $\phi = \pm 25^\circ$. The depression angle, ϕ , is defined as a counterclockwise rotation of the loudspeaker about the x -axis. This angle ensured that the entire scene was ensonified by the main beam of the loudspeaker. The microphones' beam patterns are substantially less directive, and all signals backscattered from the scene would arrive inside their main beams without any rotation. Figure 3 shows the transducer array installed on the carriage of the linear actuator.

Figure 4 describes the hardware configuration in the data acquisition system. The microphones were connected with BNC cables to a GRAS 12AX four-channel signal conditioner (<https://www.grasacoustics.com/products/power-module/product/690-12ax>) which both powered the microphones and applied +20 dB gain to the signals. The outputs of the signal conditioner were connected to an NI USB-4431 data acquisition device (<https://www.ni.com/en-us/shop/model/usb-4431.html>) which simultaneously sampled the signals. The USB-4431 also generated a pulsed signal that was transmitted to the loudspeaker through a QSC RMX4050a amplifier (<https://www.qsc.com/solutions-products/power-amplifiers/portable/2-channel/rmxa-series/rmx-4050a/>). Two temperature probes with K-type thermocouples sampled the air temperature at two different locations in the scene. The first probe, τ_1 , was located at the edge of the linear actuator at approximately $\xi_1 = [2.5, 0, 1]$ m and was digitized by an NI TC-01 analog to digital converter (<https://www.ni.com/en-us/shop/model/usb-tc01.html>). The second probe, τ_2 , was located at the edge of the scene near the data acquisition electronics at approximately $\xi_2 = [-0.5, 1, 0.05]$ m. A relative-humidity probe was co-located with τ_2 . A Thorlabs TSP01 (<https://www.thorlabs.com/thorproduct.cfm?partnumber=TSP01>) digitized both the temperature and humidity data at ξ_2 . Data acquisition was automated in LabVIEW. Software timing in LabVIEW ensured synchronization between the transmitted and received signals and simultaneous sampling of the signals on all four microphones. Motion of the linear actuator was commanded through LabVIEW to a Parker IPA04 motion controller (<https://ph.parker.com/us/en/product-list/ipa04-hc-single-axis-servo-drive-controller-3-0a-1-100-240vac-1-1kva>), which also provided feedback about the actuator's position. The range resolution of reconstructed SAS imagery



Fig. 3 The transducer array consisted of a loudspeaker and four microphones. The array is shown mounted to the carriage of the linear actuator, with the loudspeaker pointed upward to $\phi = +25^\circ$ for the free-field environment. Acoustically absorptive foam behind the transducers attenuates sound transmitted and received from the rear that could interfere with the desired signals scattered by the scene.

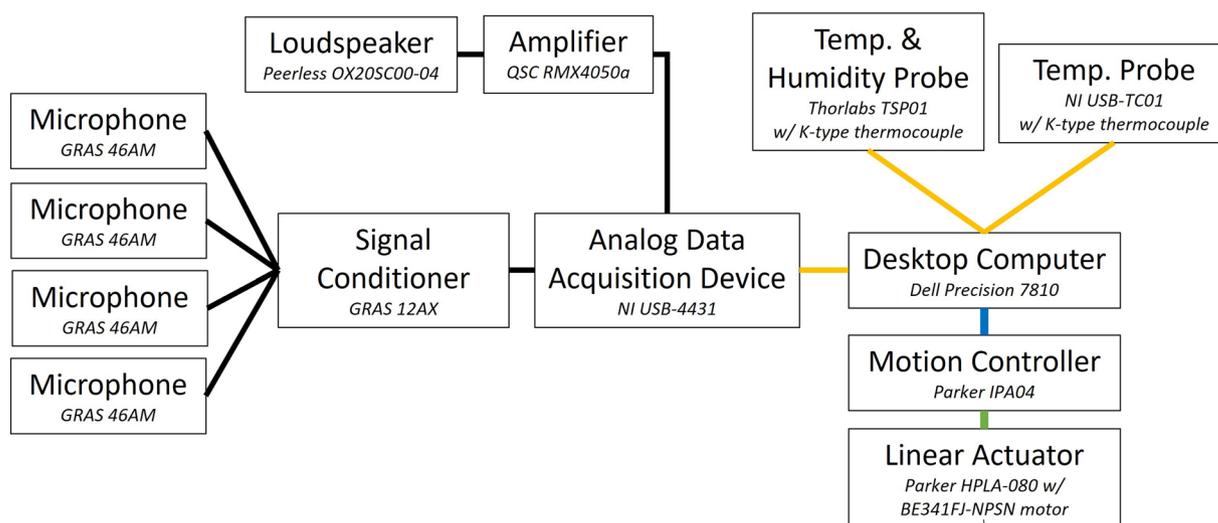


Fig. 4 The experimental instrumentation is described in a block diagram. COTS hardware was integrated together in the data acquisition system. Black lines indicate analog signals, yellow lines indicate digital communication via USB, the blue line indicates digital communication via Ethernet, and the green line indicates digital communication between the motor and the controller. A Dell Precision 7810 computer running LabVIEW software automated the data collection.

is determined by the pulse bandwidth, while the along-track resolution is determined by the sensor directivity and spatial sampling pattern. By transmitting a pulse with 20 kHz bandwidth from a speaker with a 1.91 cm aperture, the system was designed to produce SAS imagery with approximately $0.9 \text{ cm} \times 0.9 \text{ cm}$ pixel resolution²⁵.

Four classes of targets, shown in Fig. 5 were employed: solid 10.2 cm (4") diameter polyurethane spheres (McMaster-Carr, USA <https://www.mcmaster.com/6490K27/>), hollow 10.2 cm (4") diameter aluminum spherical shells with 1.5 mm wall thickness (Custom Ornamental Iron Works Ltd, USA https://customironworks.com/metal-balls-c-1/aluminum-hollow-balls-c-1_142/aluminum-hollow-balls-30815-p-1716.html), 20.3 cm (8") diameter block letter Os, and 20.3 cm (8") diameter block letter Qs. Both the O and Q targets were fabricated from 1.91 cm (0.75") thick medium density fiberboard (MDF). These targets were designed to have distinguishing physical features that could be clearly resolved in the SAS imagery, but also have one or more acoustic effects that might be discriminatory. For example, an incident waveform on the hollow sphere would be expected to excite structural modes that would differentiate the response from that of the solid sphere of the same size²⁶. That is, elastic scattering from resonant targets can produce discriminating features in SAS imagery^{27,28}. Additionally, the Q target has a tail that differentiates its shape from the O target. The exterior of the tail also forms a corner reflector with the body of the Q, which can cause an enhancement at this point in the SAS imagery²⁹. Seven replicas of each class of target were procured and labeled in marker with an integer from 1 to 7 indicating that target's position within the scene. The mass of each target was measured using a digital balance scale to partially characterize variation in physical properties among the targets.



Fig. 5 Four classes of targets were employed in the experiment: (a) solid polyurethane spheres, (b) hollow aluminum spherical shell, (c) block letter O's, and (d) block letter Q's. The targets have shapes that can be clearly resolved in the SAS imagery and one or more acoustic affects that might be discriminating features.

SAS imagery of underwater objects typically contains features that relate to the target, features that relate to the background, and features that relate to their interaction. Although the experimental configuration is similar in some ways to underwater SAS systems and scenes, it is not intended to replicate an underwater environment in air. Instead, targets and backgrounds were developed with features that scale in complexity, could be reliably procured or manufactured, and could be carefully characterized. Four environment classes were designed that increased in complexity by incrementally introducing these features to the data. The simplest background is a free-field environment. In this case, the targets were suspended by thin wires far from any surfaces. This approximates, to the greatest extent possible in this experimental configuration, the absence of a background. Next, the targets were placed on a smooth flat, planar interface. Scattering from the smooth interface is principally in the forward direction and there is minimal energy backscattered to the microphone array. The reflection of sound, however, does introduce local multipath between the interface and the target³⁰. This interaction adds a phenomenon to the acoustic data that is not present in the free-field environment, but the lack of backscattering from the interface still minimizes phenomena relating to the background alone. A rough interface background was created from plastic pellets, and the targets were placed proud on this interface. The rough interface produces diffuse backscattering from the background, and occlusion (shadowing) of portions of the background by the proud target adds an additional interaction. Finally, in the most complex environment, targets were partially buried amid the rough interface. This adds further interactions between the target and the background because each may occlude the other. Figure 6 shows photographs of the experiment with the different background environments installed. The four environments were each characterized with measurements of acoustic scattering without targets present. Additional characterization of these environments (such as roughness estimation) was not made as part of this data set.

SAS Data collection. For a given background, scene data was collected sequentially for each class of targets placed in the environment. Then a new background was introduced and the process was repeated. To prevent data labeling errors, each scene consisted of only a single background and target class. The procedure for a SAS scan involved first preparing the scene and then collecting the acoustic data. The targets were placed in the same nominal positions within the scene, defined by the bounding boxes in Fig. 7. The centers of the seven boxes were arranged in a row of four and a row of three, separated by 0.75 m in each dimension. The targets were assigned a position in the scene that matched the number written on the target. This ensured that the same target was in each position across each background in order to reduce uncertainty in the data related to variation in the target properties.

Each target was manually picked up and replaced between collections in order to introduce natural target-position variability into the data. The targets were allowed to randomly vary in position as long as upon replacement the targets visually appeared to be within the bounding box at each location. The block letter O and Q targets were placed with the faces nominally parallel to the floor. The orientation of the block letter Q targets was allowed to vary randomly such that the legs were within a 10.2 cm line centered on the near-range edge of the box. The other three target classes are rotationally symmetric so their orientation was not considered upon replacement. The details of preparing the scenes for each background environment are described in the following subsections.

Once the scene was prepared, the acoustic data was collected with automated software in LabVIEW. First, the linear actuator initiated a routine to place the carriage in the home position as indicated by the electromagnetic sensor. Once the home position was established, the system repeated a sequence of transmitting a pulse, recording the signals on the microphones, recording the temperature and humidity, then advancing the carriage 5 mm. The transmitted pulse, $u(t)$, was a 500 μ s linear frequency modulated (LFM) downchirp from 30 to 10 kHz with a 10% Tukey window,

$$u(t) = w(t) \sin \left(2\pi \left(f_1 + \frac{f_2 - f_1}{2t_p} t \right) t \right) \quad (1)$$

where t is time in seconds, $f_1 = 30$ kHz, $f_2 = 10$ kHz, $t_p = 500$ μ s. The window function, $w(t)$, is defined by

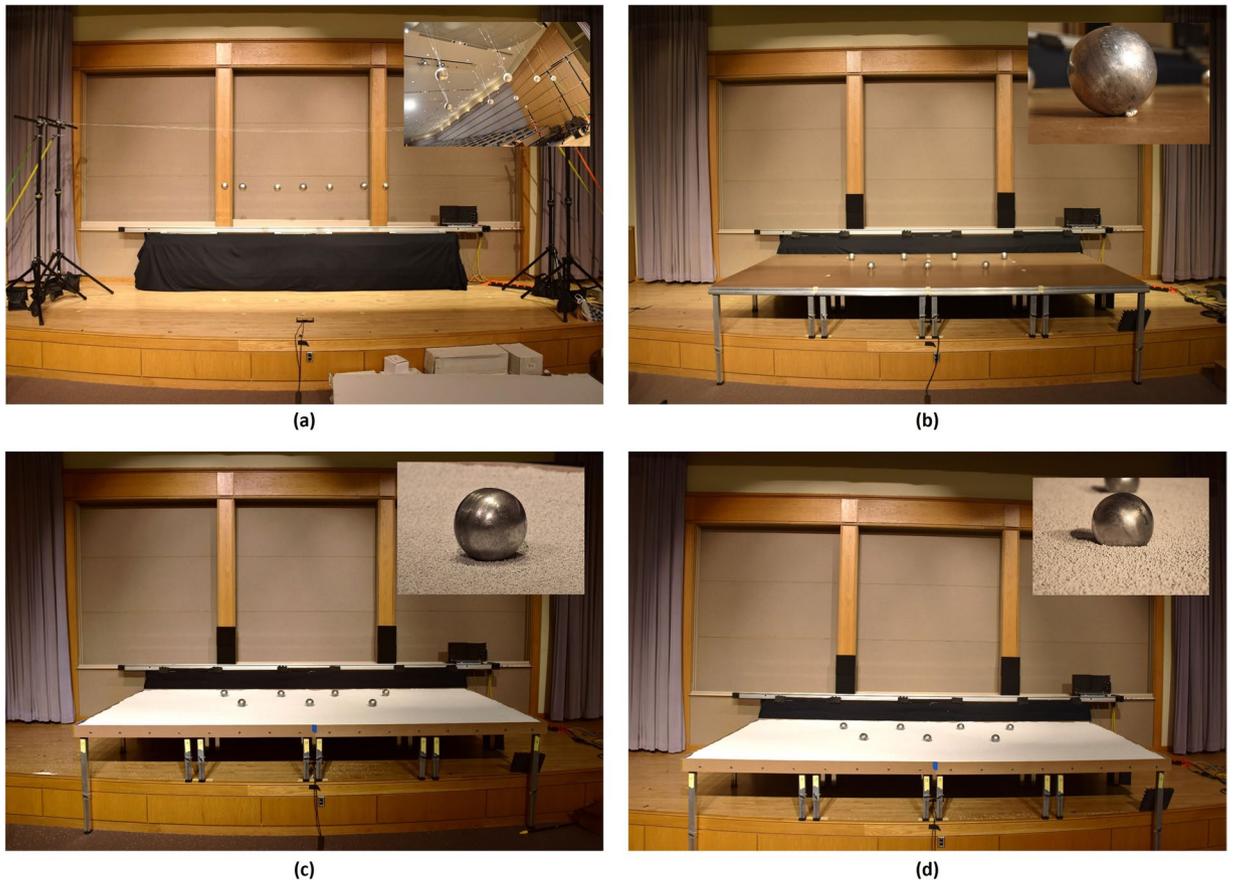


Fig. 6 Four environment classes that increase in complexity were designed for the experiment. First targets were suspended from fishing line in a free-field environment (a). Then the targets were placed upon a flat planar interface (b). Next, the targets were placed proud on a rough interface composed of HDPE pellets (c). Finally, targets were partially buried in the rough interface (d).

$$w(t) = \begin{cases} \frac{1}{2} \left(1 - \cos \left(\frac{2\pi t}{rT} \right) \right) & 0 \leq t < \frac{rT}{2} \\ 1 & \frac{rT}{2} \leq t \leq \left(1 - \frac{r}{2} \right) T \\ \frac{1}{2} \left(1 - \cos \left(\frac{2\pi t}{rT} \right) \right) & \left(1 - \frac{r}{2} \right) T < t \leq T \end{cases} \quad (2)$$

where the Tukey window fraction $r = 0.1$. SAS range resolution is inversely proportional to the bandwidth of the transmitted signal. The choice of frequencies was intended to maximize the image resolution with this particular model of speaker. The Tukey window applied to the chirp reduces the amplitude at the beginning and end of the waveform in order to minimize the transient response of the speaker and shorten the temporal ambiguity function of the transmitted signal. This sequence was repeated 1001 times, advancing the carriage a total of 5 m from the home position. The nominal 5 mm advance per ping is less than half a wavelength at 30 kHz (the highest frequency in the band), which sufficiently samples the synthetic aperture so that an image can be reconstructed from the signals on any single microphone without aliasing²⁵.

The actual advance per ping varies stochastically by a small amount compared to the nominal advance. The mean of the measured advance was 5.004 mm and the standard deviation was 65.42 μm . The position of the carriage at each ping was monitored by a feedback encoder in the actuator's motion controller. The time series of acoustic data, temperature, humidity, and along-track position of the transmitter (as reported from the motion controller) on each ping are recorded in hierarchical data format (.h5) files (<https://www.hdfgroup.org/solutions/hdf5/>) corresponding to each collection as described in Table 1. Any unexpected behavior or events that occurred during a collection were noted by the system operator and qualitatively described in the data record. Each configuration of a given target and environment was scanned at least 31 times. Some configurations were measured more times as time allowed in the experimental schedule. Scattering measurements of each background (with no targets present) were also collected in order to characterize that portion of the scene. The background noise measured by the system was also collected periodically throughout the experiment.

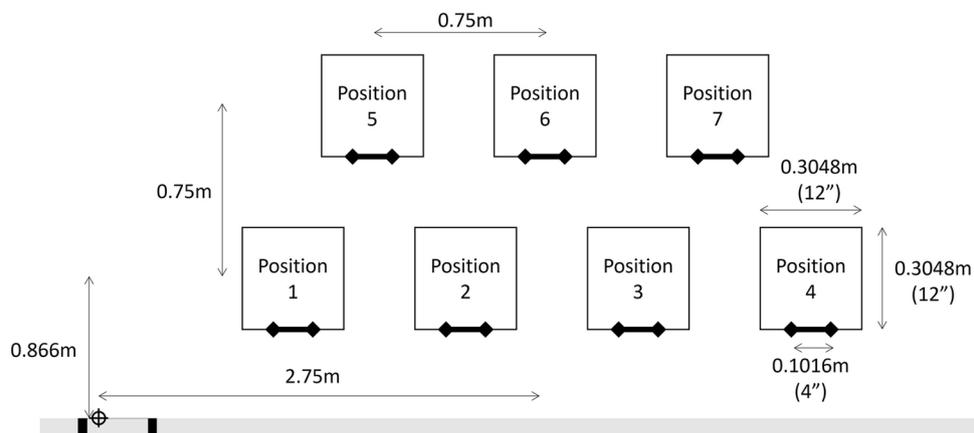


Fig. 7 Targets were arranged in two rows, with four objects in the front row and three objects in the back row, to ensure that each target was ensounded over the same range of angles. Between scans the targets were manually picked up and replaced within a box centered around each position in order to randomly perturb their positions. The block letter “Q” targets were oriented such that the leg was within 10.2 cm of the center of the box.

Filename	Description	Number of Collections
t0e1_xx.h5	No target present, free-field	8
t1e1_xx.h5	Solid sphere, free-field	42
t2e1_xx.h5	Hollow sphere, free-field	32
t3e1_xx.h5	Letter O, free-field	33
t4e1_xx.h5	Letter Q, free-field	31
t0e2_xx.h5	No target present, flat interface	14
t1e2_xx.h5	Solid sphere, flat interface	34
t2e2_xx.h5	Hollow sphere, flat interface	34
t3e2_xx.h5	Letter O, flat interface	63
t4e2_xx.h5	Letter Q, flat interface	53
t0e3_xx.h5	No target present, rough interface	32
t2e3_xx.h5	Hollow sphere, proud on rough interface	32
t3e3_xx.h5	Letter O, proud on rough interface	32
t4e3_xx.h5	Letter Q, proud on rough interface	31
t1e4_xx.h5	Solid sphere, partially buried in rough interface	32
t2e4_xx.h5	Hollow sphere, partially buried in rough interface	32
t3e4_xx.h5	Letter O, partially buried in rough interface	36
t4e4_xx.h5	Letter Q, partially buried in rough interface	35
noise_xx.h5	Noise recording (no waveform transmission)	5

Table 1. The acoustic data from each collection are stored in .h5 files. The filenames indicate the type of data present in the file. Each file is a unique data collection of that particular configuration.

Free-field environment. In the free-field environment, targets were suspended in front of the linear actuator with fishing line. Steel aircraft cable 1.5 mm in diameter was stretched between aluminum tripod stands placed outside of the imaging scene. Ratchet straps were connected between the tripods and 5 gallon buckets filled with concrete in order to tension the lines. Each target was hung from the lines with four 10 pound test fishing lines. Snap swivels connected the fishing lines to the steel cables. For the spherical targets, the four lines were connected through a small aluminum loop attached to the spheres with cyanoacrylate adhesive. For the O and Q targets, the fishing lines were connected with staples at 90° spacing. The length of fishing line used to suspend each target was measured from the point of attachment at the target to the end of the swivel with a tape measure. These values were recorded in a spreadsheet described in Table 3.

Masking tape was placed on the stage floor below the center of each target. These markings identified the center and edges of the bounding boxes and were used to visually position and align targets. After repositioning of the targets between scans, the targets would tend to swing like pendula from the overhead lines. Any gross oscillations were manually damped by gently touching the target with a cloth. After the manual damping, the targets were left to settle for at least 5 minutes before starting the acoustic collection. Upon setting a scene of each class of targets for the first time and letting them settle, the elevation of the targets was measured using a laser range finder. These elevations are recorded in the spreadsheet described in Table 3. The elevations of the spherical targets was measured to the lowest point on the sphere. The elevation of the O and Q targets were

Data Path	Description	Units
/position	$n_{pings} \times 1$ vector of along-track position at each ping	m
/temperature	$n_{pings} \times 2$ vector of temperature at $[\xi_1, \xi_2]$ for each ping	°C
/humidity	$n_{pings} \times 1$ vector of relative humidity at each ping	%
/xVec	vector of x-coordinates for pixel centers in reconstructed imagery	m
/yVec	vector of y-coordinates for pixel centers in reconstructed imagery	m
/zPlane	z-coordinate for the plane of the imagery	m
/chN/ts	$n_{samples} \times n_{pings}$ matrix of time series measured on channel N	V
/chN/img_re	Real component of complex-valued imagery reconstructed from data on channel N	V
/chN/img_im	Imaginary component of complex-valued imagery reconstructed from data on channel N	V

Table 2. The variables within the .h5 files containing data from each collection are organized into a group for each of the receivers with acoustic time series and imagery, and a fifth group containing non-acoustic data about the collection.

filename	Description	Units
lineLength.csv	Measurements of the lengths of fishing line used to suspend targets in the free-field environment	m
targetElevation.csv	Measurements of the elevation of the targets in the free-field environment	m
targetMass.csv	Measurements of mass of the targets	g
O.pdf	Dimensioned drawing for the block letter O target	inch
Q.pdf	Dimensioned drawing for the block letter Q target	inch
notes.csv	Qualitative notes about any anomalies that occurred during each data collection, along with the time stamp of the start of the collection	
acquisitionParams.csv	Properties of the data acquisition system during the acoustic collection	various
sensorCoordinates.csv	Measured coordinates of the speaker and microphones in the home position	m
waveformCapture.csv	Recording of the transmitted waveform, sampled at 100 kHz	V
speakerPattern.csv	Measured directivity pattern of the speaker at 10 kHz, 20 kHz, and 30 kHz	dB re: maximum
micPattern.csv	Measured directivity pattern of the channel 1 microphone at 10 kHz, 20 kHz, and 30 kHz	dB re: maximum
micCalibration.csv	Manufacturer provided electroacoustic calibration (with uncertainty) of each microphone used in the experiment at 250 Hz	V/Pa
micFrequencyResponse.csv	Manufacturer provided frequency response of each microphone used in the experiment relative to the response at 250 Hz	dB re: V/pa at 250 Hz
preampCalibration.csv	Manufacturer provided gain calibration of each microphone preamplifier channel	dB re: input

Table 3. Data from the development, calibration, and characterization of the hardware used in the experiment are stored in the “characterization data” folder.

measured to the object bottoms at four points: the negative cross-track end, the positive cross-track end, the negative along-track end, and the positive along-track end.

While there is always a background in experimental data, the free-field environment was designed to minimize the interaction of the environment with the targets. The targets were hung above the array, at a nominal elevation $z = 1.6$ m to prevent multipath interference with the floor from appearing in imagery. The rigging hardware was chosen to minimize the cross-sectional area so that it would not strongly scatter the incident acoustic signals. The maximum response axis of the speaker was rotated to an angle $\phi = +25^\circ$ from the positive y -axis so that the main beam of the projector pointed toward the targets hanging above the array. This configuration also reduced the amount of transmitted acoustic energy that was incident upon the stage floor. Upon completion of the free-field environment testing, the attachment points (staples and aluminum loops) were removed from the targets.

Flat interface environment. In the flat interface environment, a set of four 1.22 m \times 2.44 m ($4' \times 8'$) platforms covered with a 4.76 mm ($3/16''$) sheet of tempered hardboard (Eucaboard) were installed in front of the linear actuator. Both the platforms and the hardboard were aligned so that the 2.44 m dimension was parallel to the y -axis of the experiment's coordinate system. This ensured that imperfections in the joints between the platforms and hardboard would not create a discontinuity that appreciably scatters sound toward the array. The tempered hardboard has a hard finish that is smooth with surface roughness much smaller than an acoustic wavelength so that it will reflect sound. The platforms were set so that the top of the hardboard was at $z = 0.60325 \pm 0.0015$ m, as verified with a laser range finder. This elevation placed the targets at a symmetric grazing angle relative to the free-field experiments. With the targets located below the array in elevation, the maximum response axis of the speaker was rotated to an angle $\phi = -25^\circ$ from the positive y -axis so it was again pointing toward the targets. By symmetry in the experimental design, the nominal incident grazing and scattering angles between the targets and the array were unchanged between the free-field and flat interface cases.

The bounding boxes around each target position in Fig. 7 were drawn onto the hardboard with permanent marker. To set each scene, the targets were manually picked up and replaced within the boxes drawn onto the

hardboard. The spherical targets were held in place with a small piece of putty placed between hardboard and the positive y side of the target.

Proud on rough interface environment. A “sandbox” was created to build the rough interface environment on top of the same platforms and hardboard that were used in the flat interface environment. Sides made from 19.05 mm (0.75”) thick MDF were attached to the outside of the platforms. The tops of these sides were set at $z = 0.635$ m and the interior edge (visible from the perspective of the array) had a 19.05 mm (0.75”) radius applied to the corner during fabrication using a router. The radiused (rounded) corner reduces the cross-section of the rail facing the array and minimizes the amplitude of acoustic scattering from this edge. The “sandbox” was filled with high-density polyethylene (HDPE) pellets in a layer approximately 2 cm thick. The shape of the pellets was nominally spherical with 3.1 mm diameter and irregular dimples on the surface. The x - and y -positions corresponding to the centers of each target position and the bounding boxes were projected onto the rails and marked with a permanent marker. These indicators were used to visually place and align targets. To set each scene, all of the targets were first removed from the platforms. Then a push broom with a 61 cm \times 8.9 cm (24” \times 3.5”) head was used to sweep the top of the layer of pellets. This perturbed the positions of the pellets to produce a new, random realization of the rough interface. Sweeping was done in strokes parallel to the y -axis to prevent formation of ripples in the interface parallel to the x -axis that could strongly scatter sound. Finally, the targets were gently placed on top of the interface so that they sat as proud as possible above the pellets. The bearing capacity of the HDPE pellets was insufficient to support the solid sphere targets and they immediately sank to the level of the hardboard upon placement. Scans of solid spheres proud on the rough interface were therefore not possible to collect and are accordingly missing from the list of configurations in Table 1. Like the flat interface environment, the speaker was pointed downwards at an angle $\phi = -25^\circ$.

Partially buried in rough interface environment. Partially buried targets were set using the same experimental configuration as the rough interface environment. For each new scene, the HDPE pellets were swept, targets were placed in their designated positions, and then pressed into the HDPE beads. For the spherical targets, they were first placed onto the interface and then pushed down until the bottom of the sphere hit the hardboard beneath the pellets. For the O and Q targets, they were first placed proud on top of the pellets. Next, an edge of the target was pushed down until approximately 25% of the target was submerged beneath the pellets. The portion of the target that was buried was allowed to vary randomly by target position and by scene.

Noise. The background noise of the system and the environment was periodically characterized by setting the amplitude of the transmitted waveform to 0 V and collecting a scan of the scene. As no waveform was transmitted by the speaker, the signals recorded from each microphone are from the ambient acoustic noise in the space and the electronic noise of the data acquisition system.

Characterization and Calibration. Several additional experiments were performed to characterize the acoustical response of the array and the electrical response of the data acquisition system. The group delay of the data acquisition system (which accounts for the electronic delay of converters and filters in the data acquisition) was characterized by electrically connecting the output of the USB-4431 to the input, transmitting a broadband pulse, and estimating the delay between transmission and reception using cross-correlation. The directivity of the speaker and microphone were characterized using standard electroacoustic calibration procedures³¹ at 10 kHz, 20 kHz, and 30 kHz. The directivity was measured by rotating the transducer relative to a reference transducer in 1° increments. It was calculated as the ratio of the root mean square pressure at each angle relative to the maximum root mean square pressure across all angles. Only one microphone (receiver position 1) was characterized as the microphones are assumed to all be well-matched in the factory. Polar plots of the measured directivity are shown in Fig. 8. The transmitted waveform was captured by a GRAS46AM reference microphone (separate from the four used in the array) aligned with the maximum response axis of the speaker at a range of 1.002 m. The reference microphone was connected to preamplifier channel 4 and digitized with the same system described in Fig. 4. The electroacoustic response of all the microphones and preamplifier were calibrated in the factory prior to delivery. These values were tabulated and provided in the characterization data.

SAS Image reconstruction. The acoustic returns recorded on each microphone were reconstructed into complex-valued imagery using the signal processing flow described in Fig. 9. This pre-processing sequence follows commonly used SAS reconstruction techniques²⁵. First, the processing flow removed non-acoustic artifacts (DC bias and group delay) that are introduced by the data acquisition electronics. Next, a “transmit blank” algorithm set the portions of the recorded waveforms corresponding to the direct path transmission from the speaker to the microphone equal to zero. The data was high-pass filtered using a finite impulse response filter with a 5 kHz -3 dB corner to suppress out-of-band noise. Finally, the real-valued time series data was converted to a complex-valued representation using the Hilbert transform and replica correlated with the transmitted pulse described in Eq. (1). The replica correlation pulse-compresses the broadband pulse in order to improve the range resolution of the imagery. A local estimate of the sound speed c , in units of m/s was obtained for each ping using³²

$$c = 331.6 + 0.61\tau, \quad (3)$$

where $\tau = \frac{\tau_1 + \tau_2}{2}$, is the average temperature in Celsius measured by the sensor at the center of rail (τ_1) and the sensor at the edge of the scene (τ_2). Finally, complex-valued imagery was reconstructed from the N pings using delay and sum reconstruction,

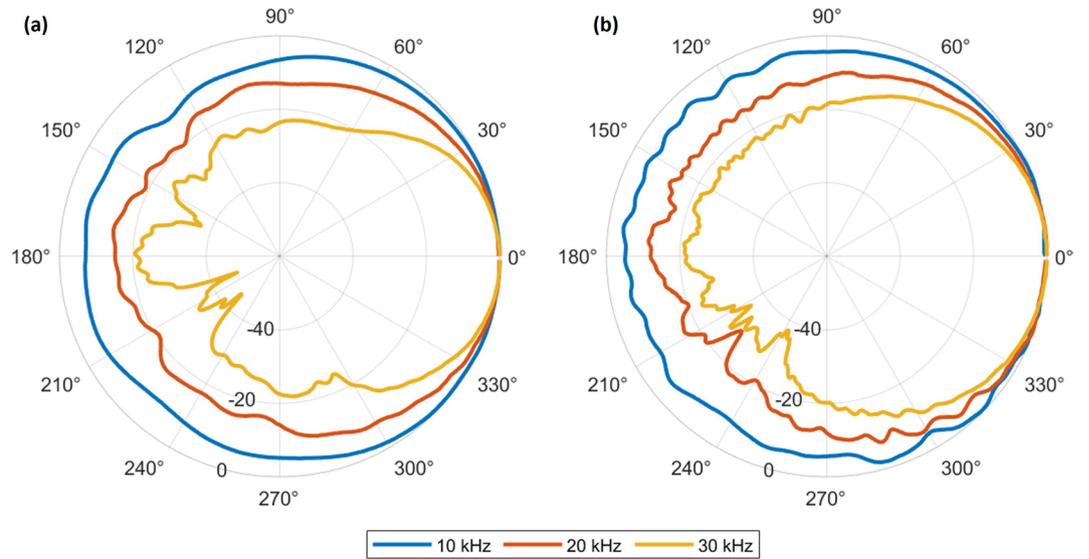


Fig. 8 The measured directivity patterns of (a) the loudspeaker and (b) the microphone at 10, 20, and 30 kHz are shown as polar plots. The nulls around 240° are due to diffraction around transducer mount and cabling.

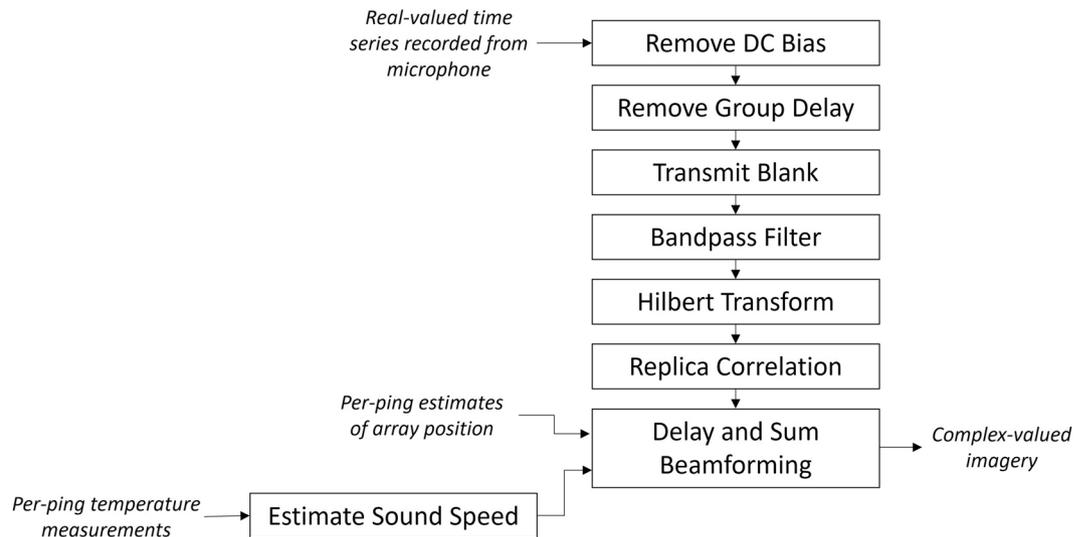


Fig. 9 The raw acoustic data were processed through a sequence of steps to minimize non-acoustic artifacts. Using the per-ping estimates of position and sound speed, the data were coherently combined to form complex-valued imagery of the scene.

$$f(\bar{\xi}) = \sum_{n=1}^N p_n \left(\frac{1}{c} (|\bar{\xi} - \bar{\xi}_R| + |\bar{\xi} - \bar{\xi}_T|) \right) u(\bar{\xi}, \bar{\xi}_T) \tag{4}$$

where $p_n(t, \bar{\xi}_T, \bar{\xi}_R)$ is the replica correlated pressure time series measured by the array on ping n , c is the speed of sound, $\bar{\xi}_T$ is the position of the transmitter on ping n , and $\bar{\xi}_R$ is the position of the receiver in the synthetic array on ping n . $u(\bar{\xi}, \bar{\xi}_T)$ is a windowing function that limits the reconstruction to a fixed azimuthal field of view

$$u(t) = \begin{cases} 1, & \left| \arctan \frac{(\bar{\xi}_{T,x} - \bar{\xi}_x)}{(\bar{\xi}_{T,y} - \bar{\xi}_y)} \right| \leq \frac{\phi_f}{2} \\ 0, & \text{otherwise} \end{cases} \tag{5}$$

where $\bar{\xi}_x$ and $\bar{\xi}_y$ and the x and y components of $\bar{\xi}$.

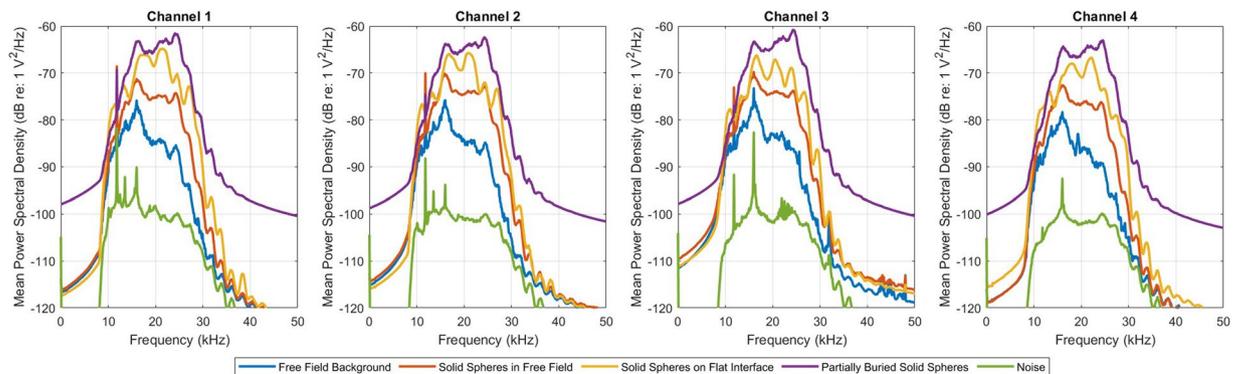


Fig. 10 Typical power spectral density plots of the acoustic signals (averaged across all pings) collected with the solid sphere targets show that acoustic signals are primarily in the 10–30 kHz band of the transmitted waveform. This power is substantially higher than the background noise, and is consistent across all four channels. This demonstrates that the acoustic signals scattered from the background and targets have a high SNR.

Equation (4) was implemented numerically for pixels within a $\phi_f = 120^\circ$ azimuthal field of view from the transmitter on each ping. This field of view encompasses more than 80% of the energy transmitted by the speaker. Energy from outside these angles is predominantly noise and is excluded from the image formation to improve image quality. Because the time delays, $\frac{1}{c}(|\xi - \xi_R| + |\xi - \xi_T|)$, often correspond to times between integer samples of the time series, nearest-neighbor interpolation after upsampling the complex-valued data by a factor of ten was used to estimate the complex-value for each pixel.

Data Records

The dataset is available on figshare³³. It is organized into two folders: “scenes” and “characterization data.” Within the “scenes” folder, the acoustic and non-acoustic data from the collection of each scene, along with the reconstructed imagery of the scene, are saved in .h5 files. Each .h5 file contains data from a unique collection of a particular target and background configuration. A four character prefix in the filename identifies the configuration of the targets and the background, and a two-digit suffix indicates the unique number of the collection. This suffix ranges from 01 to the total number of collections in that configuration. Noise recordings, where no waveform was transmitted in order to characterize the background noise in the experiment, are saved as .h5 files in the same manner using the prefix “noise.” Table 1 summarizes the naming and quantity of the .h5 files in the dataset.

The variables within the .h5 files are organized into five groups. One group contains the non-acoustic data for each collection. Four other groups contain the acoustic time series data recorded from each receiver channel and the complex-valued imagery reconstructed from that time series. Table 2 describes the organization of the data within each .h5 file. Each data component in the .h5 is annotated with their units, which are also denoted in Table 2 for convenience.

Additional data from the calibration and characterization of the experiment are stored in the “characterization data” folder. The contents of this folder are summarized in Table 3. This portion of the dataset includes the electroacoustic calibration of the receiver electronics, the directivity measurements of the transducers, measurements of the target coordinates and support lines in the free-field environment, and qualitative notes about collection anomalies.

Technical Validation

The dataset was validated through spectral analysis of the time series data, human inspection of the SAS imagery, and automated target detection of the imagery. First, the pre-processed time series signals were analyzed to estimate the signal to noise ratio (SNR) in the various configurations of the data. Figure 10 shows the typical power spectral density of the data corresponding to scattering from within the scene for each of the four microphone channels. In all cases the signals in the 10–30 kHz band are stronger than the background noise by at least 10 dB. The addition of targets and the rough interface further increases the signal level. This is expected and consistent with the scattering strength increasing in the scene from the addition of these elements. The levels are also consistent across all four microphone channels, which indicates there are no issues in the analog data acquisition hardware.

Next, the reconstructed SAS imagery was inspected for obvious defects. Errors in the along-track sampling pattern, estimation of the sensor position, or estimation of the local sound speed can introduce artifacts to the imagery such as defocusing (blur) and aliased copies of targets²². Human subject-matter experts screened the imagery for these errors and found none. This review also inspected the data for labeling errors. As illustrated in Fig. 11, close visual agreement between the known physical configuration of the scene and the appearance of the SAS image confirms that the data are free of labeling errors. The quality of the image reconstruction indicates that the timing, synchronization, and motion estimation in the data acquisition were free of significant errors.

Finally, a constant false alarm rate (CFAR) automated detector (<https://www.mathworks.com/help/phased/ref/2dcfarDetector.html>) was applied to the reconstructed imagery to estimate the target locations and dimensions. This detector is used to find pixels in imagery that contain targets. A detection is registered for a pixel

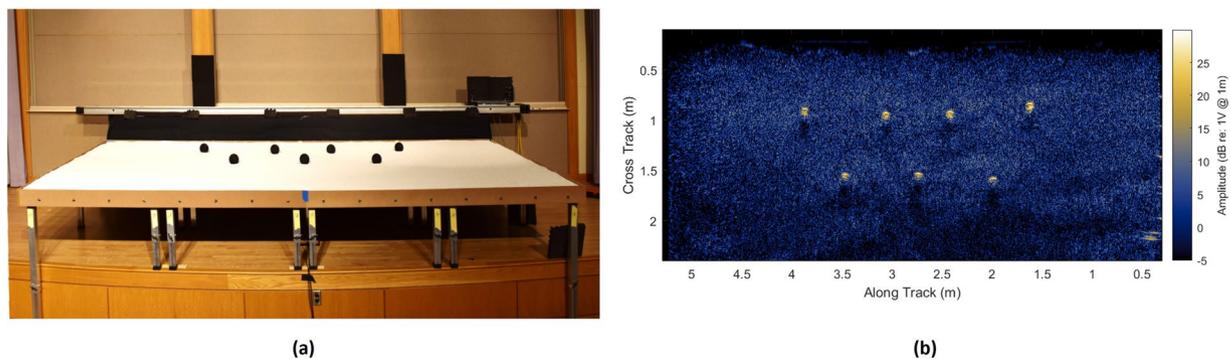


Fig. 11 The backscattered acoustic signals were reconstructed into complex-valued imagery that is an estimate of the acoustic reflectivity at each position of the scene. These images were reviewed by subject-matter experts to confirm the data is free of obvious reconstruction errors. Close visual agreement between the known physical configuration of (a) the scene and (b) the SAS image confirmed that the data is free of labeling errors.

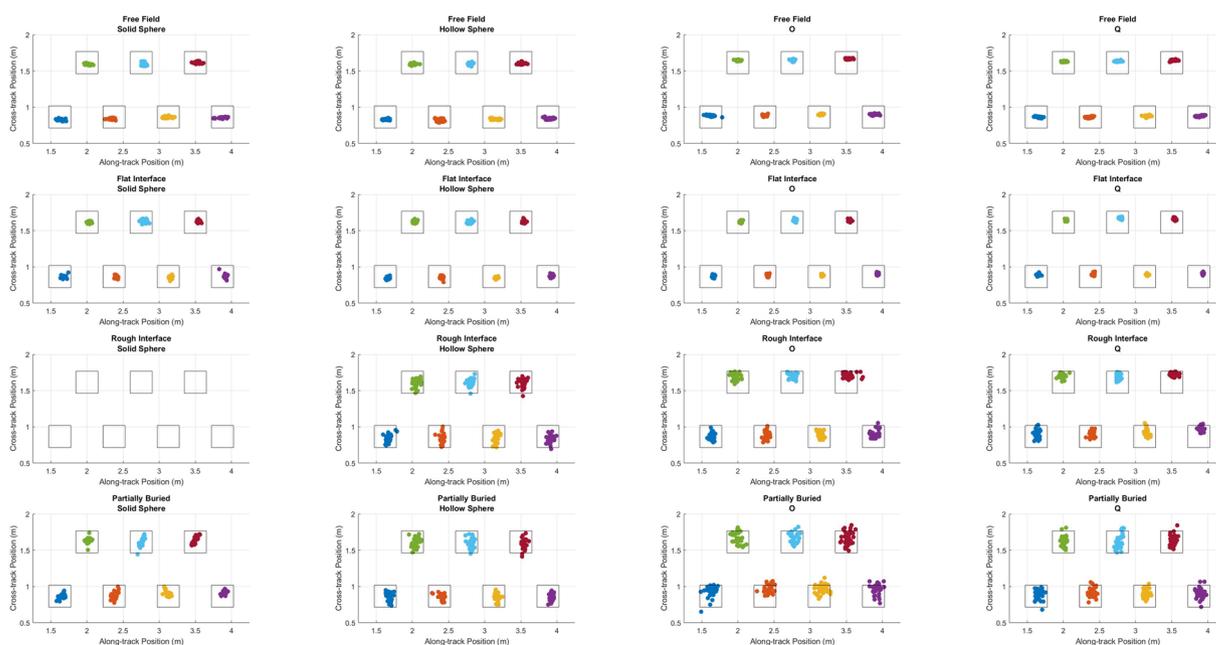


Fig. 12 Reconstructed imagery was analyzed with a CFAR automated target detection algorithm. The estimated target positions are plotted for each combination of background and target. Each point represents a position estimate from a single scene collection. Close agreement is observed between the estimated target location and the placement of the target in the local coordinate system.

if its value exceeds the noise power in the image, which is estimated from neighboring cells. The imagery was first decimated by a factor of 3 to critically sample it in each dimension. The detector was parameterized with a probability of false alarm of $1e - 7$, 6 guard band cells per side, and 10 training cells per side. The estimated centroid of each target was estimated by the mean of the coordinates of all detections in a 0.75×0.75 m box centered on the nominal position of each target as described in Fig. 7. The dimensions of each target were estimated as twice the standard deviation of the coordinates of each detection within the same region. Figure 12 shows a scatter plot of the estimated centroid of each target overlaid on a map of the bounding boxes described in Fig. 7 for each combination of target and background. The estimated centroid of the targets is generally tightly clustered within each bounding box. In the free-field case, the estimated centroids have greater variability in the along-track direction than the cross-track direction. This is because the targets were suspended from lines which highly constrained the amount of variation in the cross-track direction. The greatest variability in estimated position occurs for the rough interface and partially buried background cases. This is likely caused by greater variability in accurately placing the targets because the nearest placement markings were at the edges of the platforms. The close agreement between the estimated and intended positions of the targets in the imagery indicate that the reconstructed imagery is in focus and accurately registered in the experimental coordinate system.

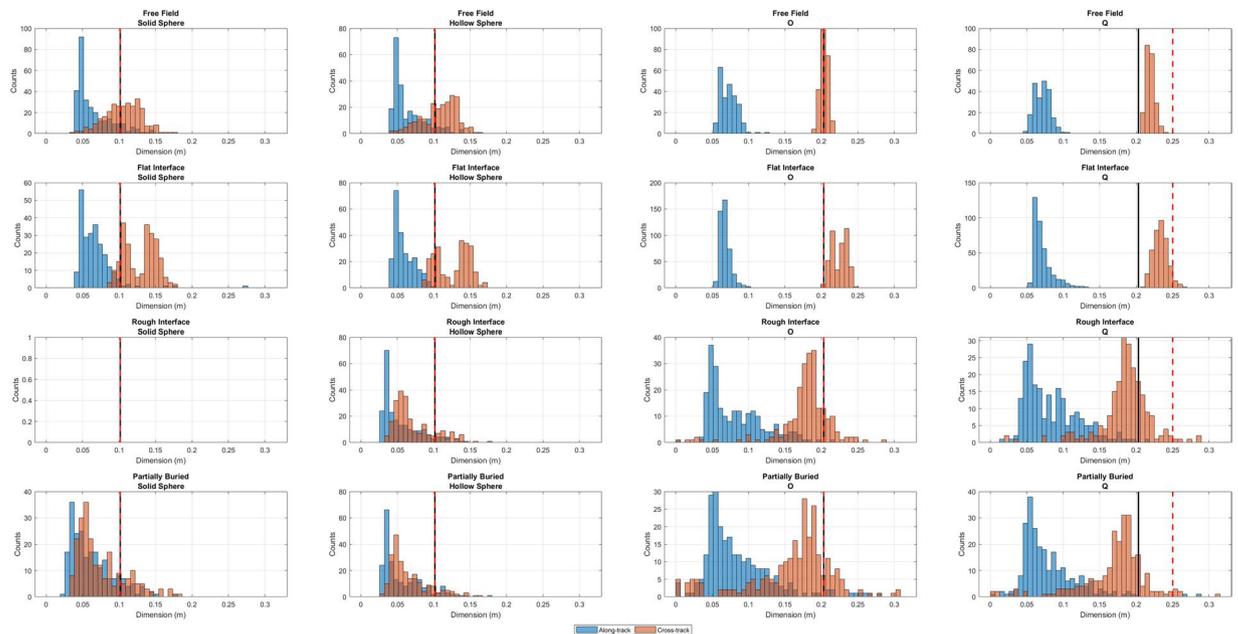


Fig. 13 The CFAR detector was used to estimate the dimensions of each target based on the detections in a $0.75\text{ m} \times 0.75\text{ m}$ box centered on the nominal position of each target. Histograms of the estimated dimensions of the bounding boxes in the along-track and cross-track dimension are plotted for each combination of target and background. The combination of solid spheres proud on the rough interface was omitted because the spheres sank into the HDPE pellets upon placement.

The estimated target dimensions from the detections are plotted by target type and by background type in the histograms of Fig. 13. The solid black line indicates the nominal dimension of the target in the along-track direction and the dashed red line indicates the nominal dimension of the target in the cross-track direction, as determined from the target construction. These lines overlap due to the symmetry of the target in all cases except for Q, where the leg increases the nominal cross-track dimension. Acoustic imaging effects will cause these estimated dimensions to diverge from the nominal, physical dimensions of the target. For example, the finite beamwidth of the transmitter does not permit observation of the target from all azimuthal angles. This will cause the estimated along-track dimension of the target to be less than the nominal dimension. On the other hand, the point spread function of the imaging system as well as acoustic interactions with the environment like target-local multipath will cause the estimated dimensions (especially in the cross-track direction) to be larger than the nominal dimensions. Nonetheless, the dimensions estimated from the detector output provide a measure of focus quality and data variability. The generally close agreement between the estimated and nominal dimensions, and the close clustering of the dimensions, indicate that the reconstructed imagery is in focus. The greatest variability in the estimated dimensions occurs for the partially buried O and Q targets. This is because in some instances the buried portions of the target are occluded by the rough interface, reducing the apparent size of the target to the detector.

Code availability

The source code used to process the acoustic data and reconstruct the imagery are available on GitHub (<https://github.com/tblanford/airsas>). The scripts, as configured in the repository, will reproduce Figs. 8, 10, 11b, 12, and 13 in this paper. Small changes should be made to the scripts to specify the path to the dataset prior to execution. This code was developed using MATLAB™ version 2022b and uses the Signal Processing and Phase Array Toolboxes.

Received: 25 April 2024; Accepted: 28 October 2024;

Published online: 05 November 2024

References

1. Bellettini, A. & Pinto, M. Design and experimental results of a 300-khz synthetic aperture sonar optimized for shallow-water operations. *IEEE Journal of Oceanic Engineering* **34**, 285–293, <https://doi.org/10.1109/JOE.2007.907933> (2009).
2. Moreira, A. *et al.* A tutorial on synthetic aperture radar. *IEEE Geoscience and Remote Sensing Magazine* **1**, 6–43, <https://doi.org/10.1109/MGRS.2013.2248301> (2013).
3. Hsieh, J. & Flohr, T. Computed tomography recent history and future perspectives. *Journal of Medical Imaging* **8**, 052109, <https://doi.org/10.1117/1.JMI.8.5.052109> (2021).
4. Wright, G. Magnetic resonance imaging. *IEEE Signal Processing Magazine* **14**, 56–66, <https://doi.org/10.1109/79.560324> (1997).
5. Gazdag, J. & Sguazzero, P. Migration of seismic data. *Proceedings of the IEEE* **72**, 1302–1315, <https://doi.org/10.1109/PROC.1984.13019> (1984).

6. Janowski, L. *et al.* High resolution optical and acoustic remote sensing datasets of the puck lagoon. *Scientific Data* **11**, <https://doi.org/10.1038/s41597-024-03199-y> (2024).
7. Misiuk, B. & Brown, C. J. Benthic habitat mapping: A review of three decades of mapping biological patterns on the seafloor. *Estuarine, Coastal and Shelf Science* **296**, 108599, <https://doi.org/10.1016/j.ecss.2023.108599> (2024).
8. Menandro, P. S., Bastos, A. C., Misiuk, B. & Brown, C. J. Applying a multi-method framework to analyze the multispectral acoustic response of the seafloor. *Frontiers in Remote Sensing* **3**, <https://doi.org/10.3389/frsen.2022.860282> (2022).
9. Williams, D. P. On the use of tiny convolutional neural networks for human-expert-level classification performance in sonar imagery. *IEEE Journal of Oceanic Engineering* **46**, 236–260, <https://doi.org/10.1109/JOE.2019.2963041> (2021).
10. Emigh, M., Marchand, B., Cook, M. & Prater, J. Supervised deep learning classification for multi-band synthetic aperture sonar. *Proceedings of the Institute of Acoustics* **40** (2018).
11. Cobb, T. J. Synthetic aperture sonar seabed environment dataset (sassed). *Mendeley Data*, V4, <https://doi.org/10.17632/s5j5g2r2vc.4> (2022).
12. Neupane, D. & Seok, J. A review on deep learning-based approaches for automatic sonar target recognition. *Electronics* **9**, <https://doi.org/10.3390/electronics9111972> (2020).
13. Chen, J. L. & Summers, J. E. Deep neural networks for learning classification features and generative models from synthetic aperture sonar big data. *Proceedings of Meetings on Acoustics* **29**, 032001, <https://doi.org/10.1121/2.0000458> (2017).
14. Huo, G., Wu, Z. & Li, J. Underwater object classification in sidescan sonar images using deep transfer learning and semisynthetic training data. *IEEE Access* **8**, 47407–47418, <https://doi.org/10.1109/ACCESS.2020.2978880> (2020).
15. Reed, A. *et al.* Neural volumetric reconstruction for coherent synthetic aperture sonar. *ACM Trans. Graph.* **42**, <https://doi.org/10.1145/3592141> (2023).
16. Sæbo, T. O., Hansen, R. E. & Synnes, S. A. V. Phase ambiguities in synthetic aperture sonar interferometry. In *OCEANS 2022, Hampton Roads*, 1–10, <https://doi.org/10.1109/OCEANS47191.2022.9977342> (2022).
17. Williams, D. P. The Mondrian detection algorithm for sonar imagery. *IEEE Transactions on Geoscience and Remote Sensing* **56**, 1091–1102, <https://doi.org/10.1109/TGRS.2017.2758808> (2018).
18. Williams, D. P. Underwater target classification in synthetic aperture sonar imagery using deep convolutional neural networks. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, 2497–2502, <https://doi.org/10.1109/ICPR.2016.7900011> (2016).
19. Bryan, O., Haines, T. S. F., Hunter, A., Hansen, R. E. & Warakagoda, N. Automatic recognition of underwater munitions from multi-view sonar surveys using semi supervised machine learning: a simulation study. *Proceedings of Meetings on Acoustics* **47**, 070018, <https://doi.org/10.1121/2.0001632> (2022).
20. Abawi, A. T. Kirchhoff scattering from non-penetrable targets modeled as an assembly of triangular facets. *J. Acoust. Soc. Am.* **140**, 1878–1886, <https://doi.org/10.1121/1.4962735> (2016).
21. Lyons, A. P., Olson, D. R. & Hansen, R. E. Modeling the effect of random roughness on synthetic aperture sonar image statistics. *J. Acoust. Soc. Am.* **152**, 1363–1374, <https://doi.org/10.1121/10.0013837> (2022).
22. Cook, D. A. & Brown, D. C. Analysis of phase error effects on stripmap SAS. *IEEE J. of Oceanic Engineering* **34**, 250–261, <https://doi.org/10.1109/JOE.2007.907935> (2009).
23. Bellettini, A. & Pinto, M. A. Theoretical accuracy of synthetic aperture sonar micronavigation using a displaced phase centre antenna. *IEEE J. Oceanic Eng.* **27**, 780–789 (2002).
24. Brown, D. C., Johnson, S. F., Gerg, I. D. & Brownstead, C. F. Simulation and testing results for a sub-bottom imaging sonar. *Proceedings of Meetings on Acoustics* **36**, 070001, <https://doi.org/10.1121/2.0001012> (2019).
25. Hawkins, D. W. *Synthetic Aperture Imaging Algorithms: with application to wide bandwidth sonar*. Ph.D. thesis, University of Canterbury (1996).
26. Junger, M. C. Sound Scattering by Thin Elastic Shells. *J. Acoust. Soc. Am.* **24**, 366–373, <https://doi.org/10.1121/1.1906905> (1952).
27. Hoang, T. *et al.* Resonant scattering-inspired deep networks for munition detection in 3d sonar imagery. *IEEE Transactions on Geoscience and Remote Sensing* **61**, 1–17, <https://doi.org/10.1109/TGRS.2023.3324223> (2023).
28. Marston, T. M., Marston, P. L. & Williams, K. L. Scattering resonances, filtering with reversible SAS processing, and applications of quantitative ray theory. In *OCEANS 2010 MTS/IEEE SEATTLE*, 1–9, <https://doi.org/10.1109/OCEANS.2010.5664606> (2010).
29. Blanford, T. E., McKay, J. D., Brown, D. C., Park, J. D. & Johnson, S. F. Development of an in-air circular synthetic aperture sonar system as an educational tool. *Proceedings of Meetings on Acoustics* **36**, 070002, <https://doi.org/10.1121/2.0001025> (2019).
30. Kargl, S. G., España, A. L., Williams, K. L., Kennedy, J. L. & Lopes, J. L. Scattering from objects at a water-sediment interface: Experiment, high-speed and high-fidelity models, and physical insight. *IEEE Journal of Oceanic Engineering* **40**, 632–642, <https://doi.org/10.1109/JOE.2014.2356934> (2015).
31. Bobber, R. J. *Underwater Electroacoustic Measurements* (Peninsula Publishing, 1988).
32. Beranek, L. L. & Mellow, T. J. *Acoustics: Sound Fields and Transducers* (Elsevier Science and Technology, 2012).
33. Blanford, T. E. *et al.* <https://doi.org/10.6084/m9.figshare.26961892> (2024).

Acknowledgements

The research presented in this work was funded by Office of Naval Research Grants N00014-22-1-2607, N00014-23-1-2846, N00014-21-1-2438, N00014-19-1-2221, and N00014-22-1-2627. The authors wish to thank the Applied Research Laboratory at Penn State University for providing the auditorium space and the linear actuator, Hannah Kurdilla, Jonah Warner, and Steven Todd for their assistance with installing and executing the experiment, and Isaac Gerg for supplying the colormap used for the display of SAS imagery.

Author contributions

T.B., D.W., J.P., and B.R. conceived the experiments, T.B., J.P., K.D., B.R., S.J., and D.B. conducted the experiments. All authors analyzed the results and reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to T.E.B.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2024