


Multiplane Diffractive Acoustic Networks

Athanasios G. Athanassiadis¹, Lennart Schlieder, Kai Melde², Valentin Volchkov³,
Bernhard Schölkopf⁴, and Peer Fischer

Abstract—Acoustic holograms are able to control pressure fields with high spatial resolution, enabling complex fields to be projected with minimal hardware. This capability has made holograms attractive tools for applications, including manipulation, fabrication, cellular assembly, and ultrasound therapy. However, the performance benefits of acoustic holograms have traditionally come at the cost of temporal control. Once a hologram is fabricated, the field it produces is static and cannot be reconfigured. Here, we introduce a technique to project time-dynamic pressure fields by combining an input transducer array with a multiplane hologram, which is represented computationally as a diffractive acoustic network (DAN). By exciting different input elements in the array, we can project distinct and spatially complex amplitude fields to an output plane. We numerically show that the multiplane DAN outperforms a single-plane hologram, while using fewer total pixels. More generally, we show that adding more planes can increase the output quality of the DAN for a fixed number of degrees of freedom (DoFs; pixels). Finally, we leverage the pixel efficiency of the DAN to introduce a combinatorial projector that can project more output fields than there are transducer inputs. We experimentally demonstrate that a multiplane DAN could be used to realize such a projector.

Index Terms—Acoustic holograms, stochastic optimization, ultrasonic transducers, ultrasound.



I. INTRODUCTION

ACOUSTIC holograms provide a high-resolution, low-cost, and low-complexity alternative to other field-shaping methods, such as phased arrays [1], [2], [3]. A single transducer with a hologram can be used to project a complex spatial field with higher resolution than state-of-the-art ultrasonic phased arrays [4]. The ability of holograms to shape such a variety of fields has led to their adoption for diverse applications, including particle trapping and manipulation [1], cellular assembly [5], acoustic fabrication [6], imaging [7], [8],

contactless power transfer [9], and transcranial ultrasonic therapy [10], [11]. However, the advantages of holograms generally come at the cost of flexibility: once a hologram is fabricated, the field that it can produce is fixed and cannot be modulated. Therefore, the existing acoustic holograms are not well suited for applications that require real-time dynamic control over the projected pressure field, as is desirable in emerging areas, such as acoustic microrobotics, acoustic fabrication, compressed sensing, and spatially resolved sonochemistry.

Different techniques have been devised to introduce dynamic effects in ultrasound using a static hologram. Fields with traveling phase fronts, such as rotating (vortex) beams, can be used to apply forces that move particles dynamically. Melde et al. [1] showed that this concept can be generalized to create a “phase surfer”—a particle that follows complex paths defined by a static hologram. However, these techniques fall short of the needs for dynamic control of a pressure field, because neither the shape of the particle trajectory nor the distribution of ultrasonic energy can be changed once the hologram is fabricated.

A promising approach, the hologram-enhanced array, was introduced recently by Cox et al. [4]. By combining a two-element array with a hologram, they showed that two different output fields could be encoded in different halves of the hologram, and the output field could be adjusted dynamically by tuning the relative amplitude of the two elements in the array. Using a hologram-enhanced array, Cox et al. [4] were able to capture and translate a particle arbitrarily in 1-D. The

Manuscript received 4 January 2023; accepted 8 March 2023. Date of publication 10 March 2023; date of current version 26 April 2023. This work was supported by the Max Planck Society through the Max-Planck-Gesellschaft–Fraunhofer Gesellschaft (MPG-FhG) Project Akustogramme. The work of Lennart Schlieder was supported by the International Max Planck Research School for Intelligent Systems (IMPRS-IS). (Corresponding author: Athanasios G. Athanassiadis.)

Athanasios G. Athanassiadis, Kai Melde, and Peer Fischer were with the Max Planck Institute for Intelligent Systems, 70569 Stuttgart, Germany. They are now with the Institute for Molecular Systems Engineering and Advanced Materials, Heidelberg University, 69120 Heidelberg, Germany, and also with the Max Planck Institute for Medical Research, 69120 Heidelberg, Germany (e-mail: thanasi@uni-heidelberg.de; kai.melde@mr.mpg.de; peer.fischer@mr.mpg.de).

Lennart Schlieder, Valentin Volchkov, and Bernhard Schölkopf are with the Max Planck Institute for Intelligent Systems, 72076 Tübingen, Germany (e-mail: lennart.schlieder@tuebingen.mpg.de; valentin.volchkov@tuebingen.mpg.de; bernhard.schoelkopf@tuebingen.mpg.de).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TUFFC.2023.3255992>, provided by the authors. Digital Object Identifier 10.1109/TUFFC.2023.3255992

Highlights

- **Acoustic holograms are powerful tools to shape sound fields, but cannot be changed dynamically. We introduce a technique to project multiple fields using a multi-layer hologram and a transducer array.**
- **Our multiplane hologram projects multiple pressure fields with higher quality than a single-plane hologram. We extend this to make the first projector that projects more fields than input transducers.**
- **Our method extends the benefits of holograms (e.g. low cost and complexity) to applications that require dynamic pressure fields for real-time control (e.g. particle manipulation or stimulation).**

hologram-enhanced array provides a useful design paradigm to combine the simplicity and high resolution of holograms with the dynamic tunability provided by phased arrays. However, this technique has only been shown to interpolate between two simple point foci. For many more complex applications, it is necessary to project and switch between arbitrarily shaped pressure fields.

Such applications would, therefore, benefit from a technique that can multiplex multiple complex holograms into a single structure. While many multiplexing techniques have been developed in optical holography for information processing and information storage [12], [13], [14], related techniques in acoustics have been much slower to develop. Besides the space-division approach described in [4], only wavelength multiplexing and depth stacking have been demonstrated with ultrasonic fields [1], [15]. These techniques offer alternative ways to project multiple planar fields using a single hologram, but are limited in their ability to provide dynamic control of the output pressure fields in a single plane. While space division is, therefore, better suited to provide dynamic control, its performance is limited when trying to combine many fields (see Section II).

To overcome these limitations, we draw inspiration from recent developments in optical holograms. Multiplane holograms have become a popular platform for multiplexing in optical mode conversion, information processing, and computation. Fontaine et al. [16], [17] showed that a seven-plane phase hologram can be used to bidirectionally transform between 200 spatially separated inputs (single-mode fibers) and 200 orthogonal outputs (Laguerre–Gaussian modes of an optical fiber).

Multiplane holograms have also led to improvements in all-optical computation and information processing. Lin et al. [18] introduced a multiplane phase hologram that was designed as a diffractive-neural network, with the phase shift at each pixel representing a neuron that could be trained using a stochastic optimization algorithm. Using this representation of the hologram, they were able to calculate a five-plane hologram that could classify thousands of untrained input fields (MNIST digits [19], [20]) into ten different categories. Finally, Wetzstein et al. [21] combined a multilayer amplitude mask with temporal modulation to create a compressive 3-D optical display.

In acoustics, multiplane holograms have been used previously to overcome certain challenges with static phase holograms. Brown [22] demonstrated that two sequential phase plates could be used to control both amplitude and phase

in the output plane. In addition, Brown et al. [23] showed that by using two holograms with restricted phase modulation depths, a stackable hologram could be created, projecting two independent patterns that could be moved relative to one another by shifting the phase plates. However, neither of these techniques is capable of dynamically switching between multiple distinct pressure patterns in real time.

Here, we adapt the idea of multiplane diffractive neural networks for acoustic holography, introducing a new acoustic hologram architecture—the multiplane diffractive acoustic network (DAN). We build a device that can project multiple distinct and spatially complex pressure fields using a static two-plane hologram and an input array with a few elements (here, we use 10). This device can dynamically switch between the different predefined output fields by sequentially activating the different inputs. In order to combine all the different input–output transformations into a static hologram, the hologram phases are calculated using optimization techniques designed for neural networks.

We show that this multiplane architecture improves the output performance for multifield projection. As a comparison, we consider a single-plane hologram that encodes different outputs in different physical locations of the hologram (space-division multiplexing). We numerically show that such techniques suffer performance limitations when projecting a larger number of output fields. By contrast, we predict and experimentally verify that a multiplane DAN can efficiently multiplex the different input–output transformations, allowing us to encode more information in a fixed number of hologram pixels. Finally, we exploit the increased multiplexing capacity of the multiplane DANs to propose a compressive projection technique that could scale up the number of projected fields while keeping the number of input transducers fixed. Our experimental proof-of-concept projector demonstrates that such a projector is possible, while revealing how we could improve its performance in future realizations.

II. MULTIPLANE DANs

A. Hologram Structure and Calculation

The basic architecture of a space-division hologram is shown in Fig. 1(a). An array of independently controlled transducers is placed a distance dz behind a phase hologram, emitting plane waves that are to be shaped into desired 2-D output fields at a distance dz beyond the hologram.

We extend the space-division hologram described in [4] to a multiplane phase hologram that consists of two $N \times N$ pixel

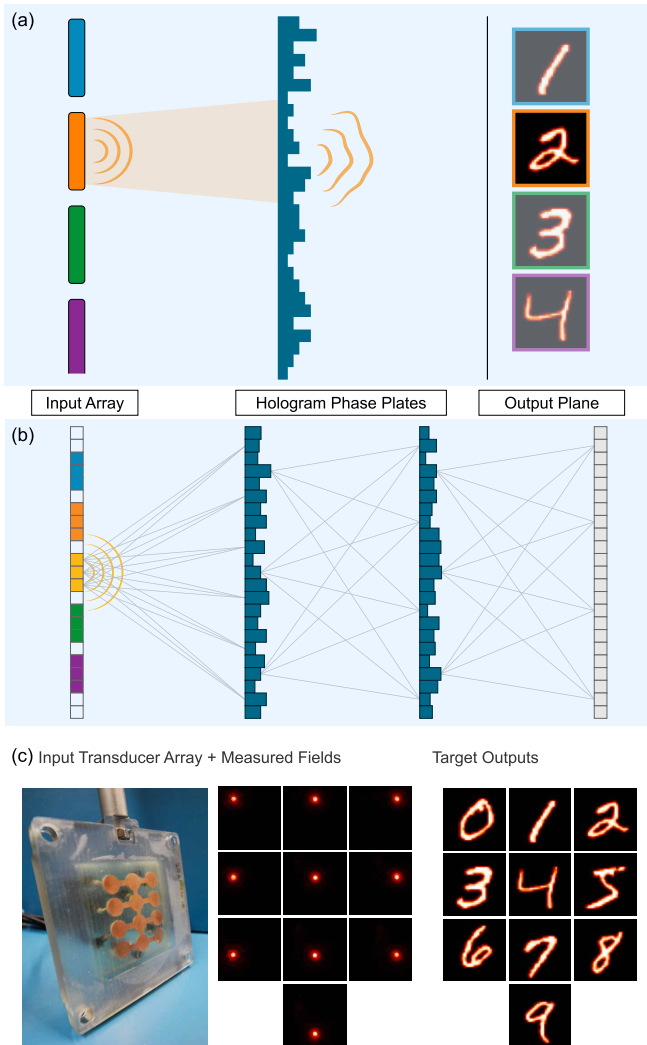


Fig. 1. (a) Space-division multiplexed hologram. Each transducer in the input array is used to excite a different output field in the output plane, thereby using a single hologram to encode multiple input–output transformations. (b) Two-plane DAN model. The system consists of the input array, two independent phase plates, and the output plane. Each plane consists of discrete nodes, which are connected between planes using the free-space Green’s function between the nodes. The phase plates can apply a variable phase shift $\Delta\phi$. (c) Custom-built array of ten circular piezoelectric transducers as the input array. The fields produced by each transducer were measured and are used as the inputs for the numerical predictions. As output fields, we consider ten digits from the MNIST dataset.

phase plates, separated by a distance d_z along the propagation axis, as shown in Fig. 1(b). The input and output planes are separated from the hologram by a distance d_z , so that the total distance between the input array and output plane is $L_z = 3 d_z$. This holographic acoustic system (input, phase plates, and output) is similar to optical diffractive neural networks [12], [18], so we refer to it as a multiplane DAN. The input and output planes are discretized pixels, and the pixels in each plane are analogous to nodes in one layer of the neural network. The weights to be calculated are the phase shifts in the two physical hologram plates, which can generally have a different number of pixels than the input and output planes. The discretized pixels on the phase plates and the input and output planes are connected by free-space wave propagation as

in standard hologram calculations. The propagation operator can be described by a matrix multiplication of the input plane with a complex propagation matrix, which is derived from convolutions with the free-space Green’s function matrix [24]. In practice, we implement the propagation operator using the angular spectrum method following [25]. Since this algorithm relies on multiplications in Fourier space, it can be efficiently implemented in modern backpropagation libraries (e.g., Tensorflow [26] and Pytorch [27]). The hologram phase shifts are then calculated using a stochastic optimization algorithm. We implement the DAN in Tensorflow and train the weights with the ADAM optimizer [28], using the mean-squared error between the calculated pressure field amplitude and the desired output amplitudes as a loss function (see Appendix C for details).

The input and output fields used for the hologram calculations are shown in Fig. 1(c). To facilitate accurate comparisons with experiments, the input fields are measured from a homebuilt transducer array. The array consists of ten circular transducers (diameter 8 mm), driven at $f = 1$ MHz (see Appendix B for more details). The target output fields are taken from the MNIST dataset of handwritten digits [19], [20]. We consider square fields, with length $L = 50$ mm per side, discretized to a 180×180 pixel grid for the propagation calculation. Each phase plate consists of $N \times N = 30 \times 30$ adjustable pixels, which are subsampled using nearest-neighbor interpolation to match the input–output grid size. The input, hologram, and output planes are separated by $d_z = 30$ mm.

B. DAN and Multiplexed Hologram Performance

After the hologram planes in the DAN are calculated, the hologram performance is numerically evaluated by propagating each input field through the hologram planes to the output plane using the angular spectrum method [24], [25]. As a reference, we also calculate the output fields produced by ten individual $N \times N$ holograms designed specifically for each input–output pair. The performance of these independent projectors serves as an upper bound of the multiplexed projector’s performance, given the system parameters (frequency, hologram size, pixel size, and output distance).

In Fig. 2, we compare the output of the two-plane, $N = 30$ pixel DAN with outputs from the individual holograms as well as with outputs from single-plane multiplexed holograms with $N = 30$ and $N = 60$. The single plane holograms are calculated using the same optimization algorithms as the DAN, since we observed that the optimization algorithm always significantly outperformed conventional iterative angular spectrum methods when designing the holograms. While a full hologram (individual) is capable of projecting each digit, the one-plane $N = 30$ multiplexed hologram fails at reproducing all ten output fields. Only certain structures of the desired fields are visible in the multiplexed outputs, and these are largely overshadowed by higher noise levels and interference artifacts. The performance is significantly improved by increasing the total number of pixels in a single plane. The one-plane $N = 60$ hologram better reproduces the

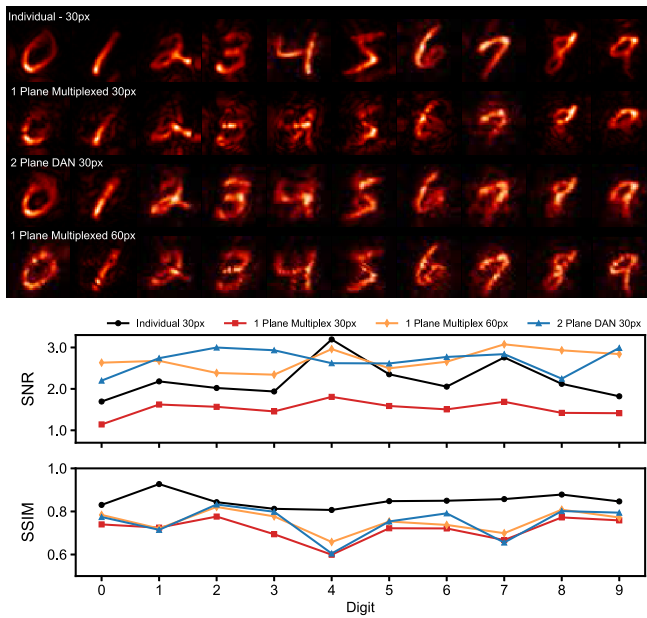


Fig. 2. Comparison of output fields for different hologram configurations. Compared with individually calculated holograms, multiplexing ten fields into one hologram introduces artifacts that result in a lower output quality. The output quality improves significantly compared with the multiplexed hologram when using a two-plane 30×30 pixel DAN or a higher resolution (60×60 pixel) single-plane hologram. The improvements are comparable for both, but the DAN uses half as many total pixels. The improved output quality is reflected primarily in a significant increase in the SNR and to a lesser extent in a small increase of the SSIM for some digits.

target fields, albeit with more artifacts than in the individual outputs. In comparison, the two-plane $N = 30$ DAN also performs very well, visually reproducing the target fields well and with less noise than the one-plane $N = 60$ multiplexed hologram, despite having fewer total degrees of freedom (DoFs; one plane: $\text{DoF} = 60^2 = 3600$ pixels and two planes: $\text{DoF} = 2 \times 30^2 = 1800$ pixels).

The output performance is quantified in Fig. 2 using two standard image quality metrics, the signal-to-noise ratio (SNR) and the structural similarity (SSIM), which describes how closely the output field resembles the target field. These metrics are described in more detail in the Supplemental Information [27] and [28]. While the SSIM is comparable for most of the different hologram architectures—reflecting the general presence of the target digit shape in the output—there is an average improvement for the holograms with more pixels. The quality improvement is more strongly reflected in the SNR, which increases significantly for the $N = 60$ one-plane hologram and $N = 30$ two-plane DAN compared with the $N = 30$ one-plane hologram.

One possible concern when comparing the different hologram architectures is that the number of total DoFs also changes when switching from one to two planes for fixed N . Therefore, we explored how the DAN performance scales with N as well as with the number of planes, providing us different ways to adjust the total number of DoF. The results are plotted in Fig. 3 for one-, two-, and four-plane DANs with N between 30 and 180. In general, there are clear

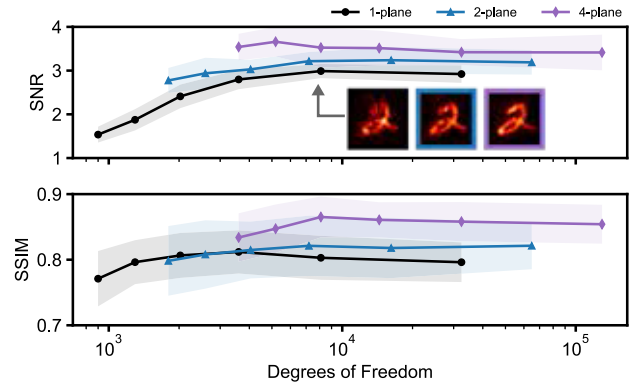


Fig. 3. Performance metrics (SNR and SSIM) for different DAN configurations as a function of total DoFs (number of pixels). For each configuration, data points and error bars represent the mean performance and standard deviation across all ten digits. As the number of DoFs increases, the performance generally increases. For a fixed number of DoFs, using more planes in the DAN also improves performance. Inset: the digit-2 output for different configurations with around 8000 DoFs.

performance improvements when adding more DoF and when adding more planes. This trend is most pronounced for the SNR. Above around 300 pixels per plane, the SNR does not improve when adding more pixels, but increases meaningfully when increasing the number of planes. A comparison of the outputs for digit “2” is shown in the inset for the three different plane counts, with comparable number of DoFs. A similar trend is observable in the SSIM, although the differences are much smaller. Further comparison of the performance of different DAN architectures is provided in the Supplemental Information [27] and [28].

The improved performance of a multilayer network over that of a single-layer network may initially seem counterintuitive for linear systems. In general, one would expect that multiple linear operations (free-space propagation and phase shifts) could all be combined into a single transformation that could be performed in a single hologram plane. However, the physics of the propagation process introduce important constraints that make this infeasible experimentally. By combining free-space propagation and phase shifts, the resulting transfer function would generally involve both phase shifts and amplitude changes, which are not realizable with a phase hologram. Therefore, by splitting the transformation up over multiple planes, a wider range of outputs can be realized without needing amplitude control. Such an approach was leveraged by Brown [22] for accurate amplitude and phase modulation using a two-plane phase hologram.

The performance scaling presented in Fig. 3 demonstrates a clear benefit from using the multiplane architecture. For a fixed number of total pixels, using more planes tends to produce lower noise results. This means that for a fixed output quality, larger pixels can be used, or for fixed pixel sizes, higher quality can be achieved by using multiple planes with a fixed number of total pixels. The multiplane DANs, therefore, utilize the available pixels more efficiently than a single-plane hologram.

To validate the multiplate hologram performance, we measure the output from our two-plane hologram in a water tank. The optimized phase plates are 3-D printed out of a rigid

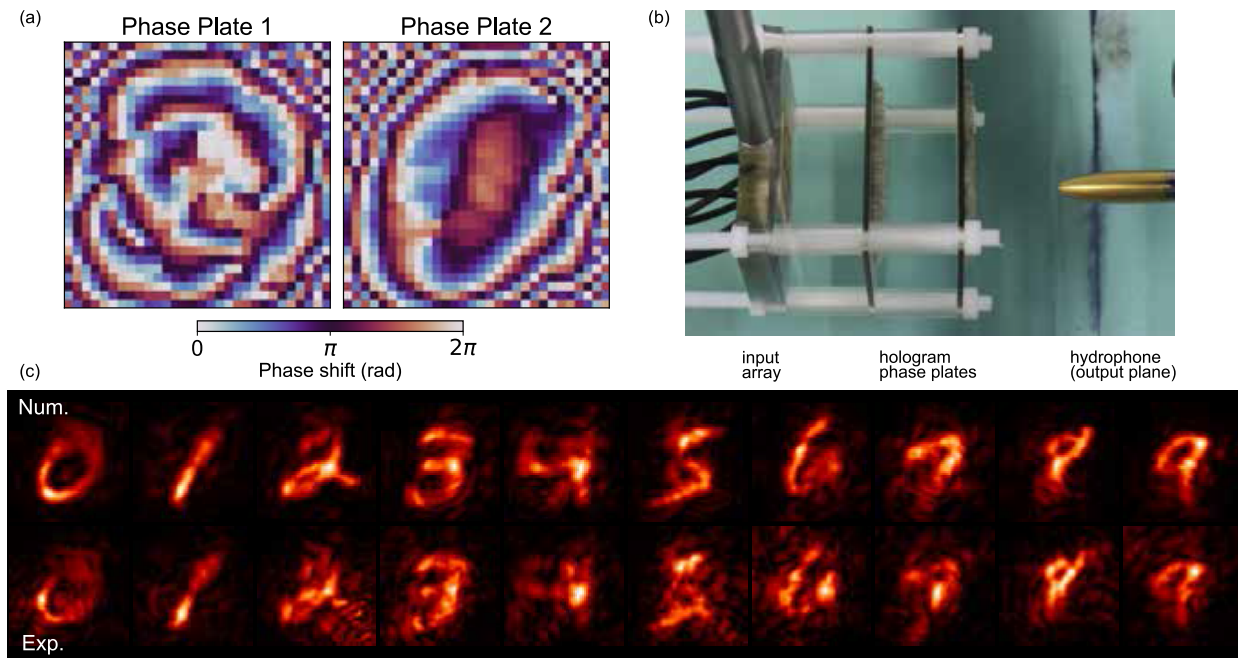


Fig. 4. Experimental outputs from a two-plane DAN match the predicted outputs. (a) Optimized phase plates for the two-plane DAN. (b) Phase plates are 3-D printed and rigidly assembled with the input array. The entire structure is submerged in a water tank, and the output field is scanned with a hydrophone for each input transducer. (c) Experimentally measured output fields (bottom) agree with the numerical predictions for the DAN outputs (top). Although the experiments are generally noisier, the structure of each digit is clearly observable, and features, such as pressure hot spots, match between the predicted and measured fields.

PMMA-like plastic (Objet VeroClear, see Appendix A for properties). The phase shifts imparted by the phase plates are shown in Fig. 4(a), and the assembled multiplane hologram in Fig. 4(b). The printed material has a sound speed higher than that of water, so that the thickness of each pixel on the phase plate determines the phase shift acquired by the incident wave [1]. The two-phase plates are assembled into a rigid structure along with the input array, as shown in Fig. 4(b). Rigid spacers are used to maintain the necessary distances between the planes, and the structure is fastened with threaded rods that provide lateral alignment. When each input of the transducer is sequentially activated, the pressure in the output plane is scanned using a 0.4-mm PVDF hydrophone (HGL0400, Onda Corporation) mounted on a 3-D motorized translation stage. The measured pressure amplitude fields are shown in Fig. 4(c).

The projected fields measured in experiments agree well with the numerically predicted outputs. Although the experimental measurements contain more noise, the digits are clear. In general, the experimentally projected digits are slightly less uniform than numerically predicted. This may be partly explained by small misalignments of the plates relative to each other, or a tilt of the measurement plane relative to the output plane. Additional discrepancies between the numerical results and the experimental measurements could arise from reflections [27] and [28], nonplanar incidence of the pressure waves on the holograms, or interactions with guided acoustic or elastic modes within the hologram plates [22]. However, most of the observed artifacts, such as the locations of hotspots and interference artifacts (see, e.g., amplitude distribution in the digits “0” or “1”), are carried over from the numerical

predictions, indicating that they could be reduced by improving the hologram calculation step.

One challenge that arises when physically implementing the multiplane DANs is the relatively low energy transmission through the multiplane structure. The measured acoustic energy in the output plane is typically around 20% of the energy measured at the input plane. Up to 35% of the losses are expected because of energy diffracting beyond the edges of the measured field, which is captured by the numerical model. Other significant sources of losses are reflections and attenuation, which are not captured by the angular spectrum propagation. Because of the higher sound speed and density of the 3-D printed material, the acoustic impedance of the phase plates is higher than that of water, leading to partial reflections at the water-phase plate interface. Moreover, there is a small amount of ultrasonic attenuation within the phase plates that further reduces the energy transmission. Accounting for these two factors, we estimate an energy transmission coefficient around 30% for transmission through the two-plane holograms. Combining the diffraction and transmission losses, we estimate that 23% of the energy will propagate to the output plane, in good agreement with the experiments. Further losses are possible due to excitation of acoustic or elastic wave modes within the phase plates, which would require more complex 3-D simulations to accurately quantify. To mitigate transmission losses in the future, different phase plate materials could be explored, impedance matching layers could be added to the interfaces, or the two-plane holograms could be made out of a single-body block, as discussed in [22].

An additional experimental challenge is properly aligning the input array to the hologram planes and aligning the

different planes within the multiplane DAN. Small alignment errors can lead to large unanticipated artifacts and errors in the output fields. Such errors were explored by Brown [22], who used a two-plane acoustic phase hologram to control the amplitude and phase of a projected field. They observed that, although the output was robust against small perturbations to the spacing between the planes, small subpixel in-plane misalignment led to significant reduction in the output field quality. The sensitivity of multiplane holograms to alignment issues is a major challenge to implement multiplexed holograms in real-world settings. One benefit, therefore, of using a multiplane DAN, rather than a high-resolution single-plane hologram, is the ability to spread a fixed pixel count over multiple planes. For an equivalent number of total pixels (DoFs), a two-plane hologram can use larger pixels than a one-plane hologram, making it more robust against misalignments.

III. COMPRESSIVE PROJECTION USING DANs

We have shown that DANs can more efficiently use the pixels in each hologram, allowing them to multiplex fields using a two-plane hologram more effectively than can be done with a single-plane hologram. Here, we build on this idea to introduce a compressive projector: a DAN that exploits this pixel efficiency to project more output fields than there are inputs in the array. Such a projector would be useful, for instance, to move an acoustically controlled system between a large number of states using only a small number of input transducers. Related ideas in optics have leveraged multiple amplitude-attenuating layers and temporal multiplexing to create compressive 3-D displays [21].

The basic premise of the compressive acoustic projector is shown in Fig. 5(a) for an idealized four-transducer input array. In order to increase the number of realizable output fields, we excite two transducers at a time, creating six different input combinations from the four transducers. In general, for P elements chosen from an N -transducer array, this approach could provide combinatorially more output fields than input transducers, with the maximum number of output fields scaling as

$${}_N C_P = \frac{N!}{P!(N-P)!}.$$

To explore the limitations and performance in a proof-of-principle device, we consider the $N = 4$ and $P = 2$ configuration, which defines six input fields.

The primary challenge with mapping each of these six input fields to distinct output fields is the number of simultaneous constraints for each given input transducer. For instance, in the four to six (input–output) mapping, each transducer appears in three different input fields. In practice, no existing direct methods exist to design holograms that satisfy these constraints. Therefore, we use our DAN optimization approach to calculate a two-plane hologram that can project unique outputs for different combinations of two inputs.

We evaluate the combinatorial DAN performance both numerically and experimentally, measuring the projected pressure field as described above. Fig. 5 shows the performance

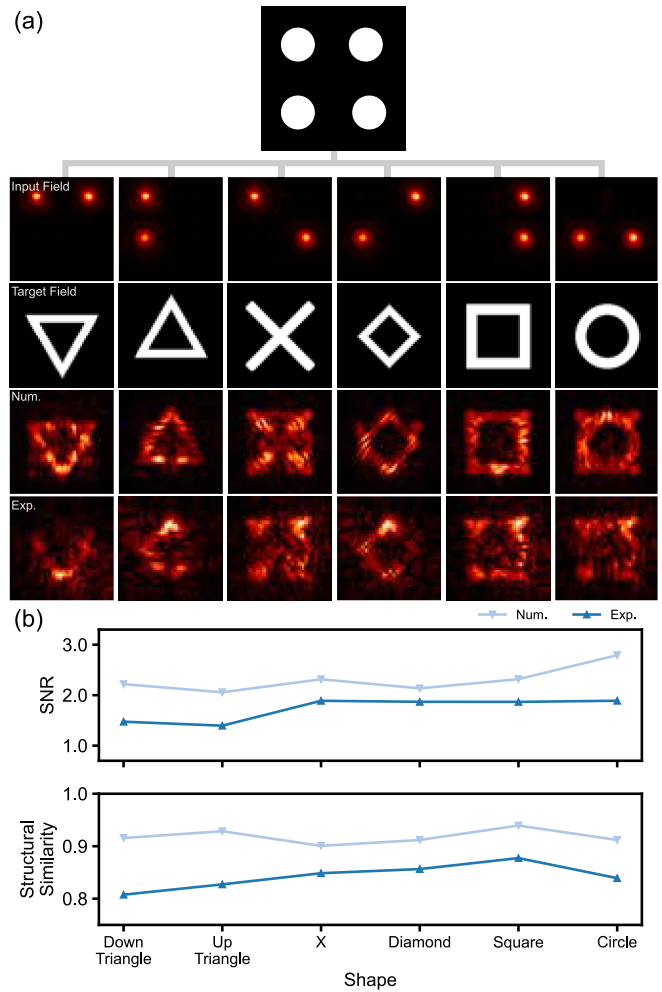


Fig. 5. Pixel efficiency of the DAN can be used for compressive projection: projecting more output fields than there are transducers in the input array. (a) Using combinations of two inputs, six distinct input fields can be generated by a four-transducer array. These inputs can then be treated independently to produce six outputs using a DAN. The numerically predicted and experimentally measured fields are shown in the bottom two rows. While there is some crosstalk between different fields, the different output shapes are clearly identifiable. (b) Performance metrics for predicted and measured fields show similar trends in quality for measured fields as for predicted ones.

of our optimized DAN combinatorial projector for a four-input/six-output mapping. The associated target fields are distinct shapes: a downward-pointing triangle, an upward-pointing triangle, an X, a diamond, a square, and a circle. The fields and holograms are discretized to a 180×180 pixel grid for propagation and evaluation, while the holograms contain 60×60 tunable pixels. As shown in the numerical predictions [third row of Fig. 5(a)], the two-plane DAN outputs can reproduce each shape, demonstrating that the principle of a combinatorial projector is realizable, albeit with crosstalk and interference artifacts. The experimental measurements reveal similar features—amplitude hotspots and crosstalk artifacts prevent more accurate reproduction of the target fields. The quantitative metrics shown in Fig. 5(b) further reflect the observed numerical and experimental performance. There is some variation in the quality for the different outputs, and

the quality in the experimental measurements is, in general, slightly lower than that predicted numerically.

While our results demonstrate that a combinatorial projector can be realized, further work is needed to improve the output quality and reduce crosstalk. As the ratio of combinatorial outputs to inputs scales up for larger input arrays (e.g., 10:5 for a five-element array), we expect crosstalk artifacts will likely become worse, requiring creative solutions to further improve performance. Since the source size generally plays a large role in the hologram output quality, rearranging the input elements with larger spacing and larger apertures could be one concrete step to improve the quality. By integrating more sophisticated numerical simulations into the design step, multiple-reflection and full-wave effects could also be accounted for in the optimization step to produce cleaner outputs with better agreement between the numerical and experimental results. Finally, a much richer set of outputs may be possible if nonlinear layers can be introduced in the hologram itself. Internal nonlinear layers are a common feature of computational neural networks, and they are necessary for such networks to compute arbitrary functions [31], such as classifying inputs (see [32, Ch. 5.1]). In optical neural networks, using even a single nonlinear layer has been shown to increase the performance of various image processing and classification tasks [33]. If appropriate nonlinearities were identified and realized as layers in a DAN, the DAN architecture could provide more diverse and unique capabilities.

IV. CONCLUSION

Acoustic holograms are rapidly growing as a tool to project finely structured ultrasound fields. Their ease of fabrication and low complexity make them attractive devices for applications, including medical therapy, imaging, particle and cellular assembly, and fabrication. However, holograms are currently limited in their ability to project time-dynamic pressure fields. One way around this limitation is to couple a hologram to a multiple-transducer array, so that each element can be used to project a different output field. In this way, temporal control is provided by multiplexing the holograms for different input–output transformations and sequentially switching which input transducer is activated.

Here, we have described a technique for multiplexing holograms using a multiplane DAN. We use an optimization algorithm to design a two-plane DAN that maps ten input pressure fields into ten distinct output fields. We show that this DAN performs comparably to a single-plane hologram with twice as many total pixels. A two-plane DAN is 3-D-printed, and experimental measurements of the output fields agree well with the numerically predicted performance. Finally, we build on the pixel efficiency of DANs and introduce a compressive projector, which can output more fields than there are input transducers. While our proof-of-concept results indicate such a projector is possible, improvements to the performance will require creative solutions. Inspired by recent developments in optical diffractive neural networks, we suggest that developing appropriate nonlinear acoustic materials could be a fruitful path toward realizing the full potential of DAN projectors.

In practice, the ability to rapidly change between structured pressure fields will allow for real-time sensing and control of acoustically responsive systems. Moreover, leveraging the pixel efficiency and the DAN optimization architecture, it should be possible to encode additional features to the holograms by designing an appropriate loss function. For instance, to create holograms that perform well in inhomogeneous media, the DANs could be designed with a stochastic optimization function that accounts for random phase delays in the numerical propagation path. The eventual realization of a combinatorially excited DAN could provide a low-complexity solution to generate a larger number of fields for real-time sensing and control in industrial, clinical, and scientific applications.

APPENDIX

A. Material Properties

In our modeling and experiments, the primary propagation medium is water (sound speed $c_w = 1484$ m/s and density $\rho_w = 1000$ kg/m³ [34], [35]), and the holograms are made out of a UV-cured polymer (Stratasys Vero Clear, sound speed $c_h = 2424$ m/s, and density $\rho_h = 1180$ kg/m³). The specific acoustic impedance of the water and holograms are, therefore, $Z_w = 1.484$ MRayl and $Z_h = 2.86$ MRayl, respectively. In addition, the hologram material absorbs ultrasound, with an attenuation coefficient of $\alpha = 2.58$ dB/cm at 1 MHz.

B. Hologram Input and Output Fields

The input transducer array consists of 10×8 -mm diameter piezoceramic disks arranged on a grid. The elements are resonant at 1 MHz. The elements are driven independently of one another, with a 20-cycle sinusoidal pulse, amplified to $20 V_{pp}$. The pressure amplitude and phase output by the transducers are scanned in a water tank using a 0.4-mm-diameter hydrophone (HGL0400, Onda Corporation) and a lock-in amplifier (Zurich Instruments UHFLI). To extract the equivalent CW pressure field, a Fourier transform is taken of the received pulse, and the pressure amplitude and phase are extracted at 1 MHz.

The calculated and measured fields are square with side length $L = 50$ mm per side, and that are discretized to $N = 180$ pixels per side. The pixel size in the input and output planes is $\delta = 0.28$ mm. The planes are separated by a distance $dz = 30$ mm. Since the phase plates typically have a smaller number of pixels per side (N), for propagation, the hologram planes are resampled to 180 pixels per side using nearest-neighbor sampling.

C. Optimization Loss Function

Single and multiplane holograms were calculated using optimization methods as described in the main text. The loss functions that were minimized during optimization were based on the mean square error of the pressure amplitude in the output plane

$$L_{\text{mse}} = \frac{1}{N^2} \sum_{n=1}^{N^2} (P_t(x_n, y_n) - P_o(x_n, y_n))^2. \quad (1)$$

Here, $P_t = |p_t|$ is the target pressure amplitude, and $P_o = |p_o|$ is the pressure amplitude produced in the output plane of the DAN. The sum is taken over all $N \times N$ pixels in the output plane.

For multiplexed holograms, this loss is, furthermore, averaged over one training batch, which consists of all output images that the DAN is designed to project. The standard MSE without normalization was used for all experiments, since it retained the most energy, which was an important consideration for experimental validation. Normalized loss functions can also be used, although these typically produced phase plates that defocused the waves, producing a slightly improved output shape and better output uniformity at the expense of energy in the output plane.

ACKNOWLEDGMENT

See Supplemental Material for details of calculations and experimental parameters, which includes Refs. [27] and [28].

REFERENCES

- [1] K. Melde, A. G. Mark, T. Qiu, and P. Fischer, "Holograms for acoustics," *Nature* vol. 537, no. 7621, pp. 518–522, Sep. 2016.
- [2] A. Marzo, S. A. Seah, B. W. Drinkwater, D. R. Sahoo, B. Long, and S. Subramanian, "Holographic acoustic elements for manipulation of levitated objects," *Nature Commun.*, vol. 6, no. 1, p. 8661, Oct. 2015.
- [3] R. Morales, I. Ezcurdia, J. Irisarri, M. A. B. Andrade, and A. Marzo, "Generating airborne ultrasonic amplitude patterns using an open hardware phased array," *Appl. Sci.*, vol. 11, no. 7, p. 2981, Mar. 2021.
- [4] L. Cox, K. Melde, A. Croxford, P. Fischer, and B. W. Drinkwater, "Acoustic hologram enhanced phased arrays for ultrasonic particle manipulation," *Phys. Rev. Appl.*, vol. 12, no. 6, Dec. 2019, Art. no. 064055.
- [5] Z. Ma et al., "Acoustic holographic cell patterning in a biocompatible hydrogel," *Adv. Mater.*, vol. 32, no. 4, Jan. 2020, Art. no. 1904181.
- [6] K. Melde, E. Choi, Z. Wu, S. Palagi, T. Qiu, and P. Fischer, "Acoustic fabrication via the assembly and fusion of particles," *Adv. Mater.*, vol. 30, no. 3, Jan. 2018, Art. no. 1704507.
- [7] P. Kruizinga et al., "Compressive 3D ultrasound imaging using a single sensor," *Sci. Adv.*, vol. 3, no. 12, Dec. 2017, Art. no. e1701423.
- [8] J. Janjic et al., "Structured ultrasound microscopy," *Appl. Phys. Lett.*, vol. 112, no. 25, Jun. 2018, Art. no. 251901.
- [9] M. Bakhtiari-Nejad, A. Elnahas, M. R. Hajj, and S. Shahab, "Acoustic holograms in contactless ultrasonic power transfer systems: Modeling and experiment," *J. Appl. Phys.*, vol. 124, no. 24, Dec. 2018, Art. no. 244901.
- [10] S. Jiménez-Gambín, N. Jiménez, J. M. Benlloch, and F. Camarena, "Holograms to focus arbitrary ultrasonic fields through the skull," *Phys. Rev. Appl.*, vol. 12, no. 1, Jul. 2019, Art. no. 014016.
- [11] S. Jiménez-Gambín, N. Jiménez, J. M. Benlloch, and F. Camarena, "Generating Bessel beams with broad depth-of-field by using phase-only acoustic holograms," *Sci. Rep.*, vol. 9, no. 1, p. 20104, Dec. 2019.
- [12] J. W. Goodman, *Introduction to Fourier Optics*, 3rd ed. Englewood, CO, USA: Roberts & Company, 2005.
- [13] X. An and D. Psaltis, "Experimental characterization of an angle-multiplexed holographic memory," *Opt. Lett.*, vol. 20, no. 18, pp. 1913–1915, Sep. 1995.
- [14] G. Barbastathis and D. Psaltis, "Volume holographic multiplexing methods," in *Holographic Data Storage* (Springer Series in Optical Sciences), H. J. Coufal, D. Psaltis, and G. Sincerbox, Eds. Heidelberg, Germany: Springer, vol. 76, 2000, pp. 21–62.
- [15] M. D. Brown, B. T. Cox, and B. E. Treeby, "Design of multi-frequency acoustic kinoforms," *Appl. Phys. Lett.*, vol. 111, no. 24, Dec. 2017, Art. no. 244101.
- [16] N. K. Fontaine, R. Ryf, H. Chen, D. Neilson, and J. Carpenter, "Design of high order mode-multiplexers using multiplane light conversion," in *Proc. Eur. Conf. Opt. Commun. (ECOC)*, Sep. 2017, pp. 1–3.
- [17] N. K. Fontaine, R. Ryf, H. Chen, D. T. Neilson, K. Kim, and J. Carpenter, "Laguerre-Gaussian mode sorter," *Nature Commun.*, vol. 10, no. 1, p. 1865, Apr. 2019.
- [18] X. Lin et al., "All-optical machine learning using diffractive deep neural networks," *Science*, vol. 361, no. 6406, pp. 1004–1008, Sep. 2018.
- [19] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998. [Online]. Available: <http://ieeexplore.ieee.org/document/726791/>
- [20] Y. LeCun, C. Cortes, and C. Burges. (2010). *MNIST Handwritten Digit Database*. [Online]. Available: <http://yann.lecun.com/exdb/mnist/>
- [21] G. Wetzstein, D. Lanman, M. Hirsch, and R. Raskar, "Tensor displays: Compressive light field synthesis using multilayer displays with directional backlighting," *ACM Trans. Graph.*, vol. 31, no. 4, p. 80, Jul. 2012.
- [22] M. D. Brown, "Phase and amplitude modulation with acoustic holograms," *Appl. Phys. Lett.*, vol. 115, no. 5, Jul. 2019, Art. no. 053701.
- [23] M. D. Brown, B. T. Cox, and B. E. Treeby, "Stackable acoustic holograms," *Appl. Phys. Lett.*, vol. 116, no. 26, Jun. 2020, Art. no. 261901.
- [24] D.-L. Liu and R. C. Waag, "Propagation and backpropagation for ultrasonic wavefront design," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 44, no. 1, pp. 1–13, Jan. 1997.
- [25] X. Zeng and R. McGough, "Evaluation of the angular spectrum approach for simulations of near-field pressures," *J. Acoust. Soc. Amer.*, vol. 123, no. 1, pp. 68–76, Jan. 2008.
- [26] M. Abadi et al., (2015). *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. [Online]. Available: <https://www.tensorflow.org/>
- [27] A. Paszke et al., "Pytorch: An imperative style, high-performance deep learning library," in *Proc. Adv. Neural Inf. Process. Syst.*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds. Red Hook, NY, USA: Curran Associates, 2019, pp. 8024–8035. [Online]. Available: <http://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [28] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
- [29] Z. Wang and A. C. Bovik, "Mean squared error: Love it or leave it? A new look at signal fidelity measures," *IEEE Signal Process. Mag.*, vol. 26, no. 1, pp. 98–117, Jan. 2009.
- [30] S. van der Walt et al., "Scikit-image: Image processing in Python," *PeerJ*, vol. 2, p. e453, Jun. 2014.
- [31] G. Cybenko, "Approximation by superpositions of a sigmoidal function," *Math. Control Signals Syst.*, vol. 2, no. 4, pp. 303–314, Dec. 1989.
- [32] C. Bishop, *Pattern Recognition and Machine Learning* (Information Science and Statistics). New York, NY, USA: Springer-Verlag, 2006.
- [33] T. Yan et al., "Fourier-space diffractive deep neural network," *Phys. Rev. Lett.*, vol. 123, no. 2, Jul. 2019, Art. no. 023901.
- [34] H. Medwin, "Speed of sound in water: A simple equation for realistic parameters," *J. Acoust. Soc. Amer.*, vol. 58, no. 6, pp. 1318–1319, 1975.
- [35] L. E. Kinsler, A. R. Frey, A. B. Coppens, and J. V. Sanders, *Fundamentals of Acoustics*, 4th ed. New York, NY, USA: Wiley, 2000.