

Acoustic Wave Manipulation Through Sparse Robotic Actuation

Tristan Shah^{1*}, Noam Smilovich^{1*}, Feruza Amirkulova², Samer Gerges², Stas Tiomkin^{1†}

Abstract—Recent advancements in robotics, control, and machine learning have facilitated progress in the challenging area of object manipulation. These advancements include, among others, the use of deep neural networks to represent dynamics that are partially observed by robot sensors, as well as effective control using sparse control signals. In this work, we explore a more general problem: the manipulation of acoustic waves, which are partially observed by a robot capable of influencing the waves through spatially sparse actuators. This problem holds great potential for the design of new artificial materials, ultrasonic cutting tools, energy harvesting, and other applications. We develop an efficient data-driven method for robot learning that is applicable to either focusing scattered acoustic energy in a designated region or suppressing it, depending on the desired task. The proposed method is better in terms of a solution quality and computational complexity as compared to a state-of-the-art learning based method for manipulation of dynamical systems governed by partial differential equations. Furthermore our proposed method is competitive with a classical semi-analytical method in acoustics research on the demonstrated tasks. We have made the project code publicly available, along with a web page featuring video demonstrations: <https://gladisor.github.io/waves/>.

I. INTRODUCTION

Manipulation has long been a central challenge in robotics [1], [2]. Robots are frequently tasked with physically interacting with their environment, be it through picking, placing, or altering objects in a variety of applications, from industrial manufacturing [3], [4] to healthcare [5], [6]. In this work, we explore robotic manipulation on a class of vastly more complex phenomena: acoustic waves.

Manipulation of acoustic waves introduces unique challenges on several fronts. Unlike traditional object manipulation, where the robot operates in direct contact with the object, wave manipulation involves indirect influence through intermediary tools such as leveraging the interaction between the scatterers and the incident field [7]–[9]. Further complicating the task, acoustic waves propagate at an extremely rapid speed [10] compared to robotic actuation. Lastly, designing an optimal controller for wave manipulation requires knowledge about the state of wave, which is a function of time and space, and usually can be only partially observed.

Establishing robust control of acoustic waves has wide reaching applications. Robotic systems capable of manipulating waves can lead to groundbreaking technologies, such as seismic wave mitigation [11], [12] and super-focusing




Fig. 1. Schematic representation of interaction between an agent and acoustic wave. The agent observes wave through its ‘Sensor’ readings, and affects it back by choosing the optimal locations and radii of the cylindrical scatterers (gray blobs) according to its ‘Controller’ (policy). Wave propagates with the speed of sound according to the Wave PDE, while robot dynamics (actuation of the cylinders) is an order of magnitude slower time scale. ‘Excitation’ represents an exogenous source of acoustic energy (e.g., speaker). The ‘Dissipation’ layer prevents wave reflections from the boundaries and allows for the efficient simulation of open space (infinite) environments. Robot’s policy is trained to achieve a desired wave configuration such as focusing energy in a particular location or minimizing total scattered energy.

devices for high-resolution ultrasound imaging and surgery [13], [14]. Mastering wave manipulation opens the door to these and many other downstream technologies.

Aside from numerical methods like finite difference [15]–[17] and discontinuous Galerkin [18], [19], which rely on computationally intensive simulations, as well as the Reduced Order Model approach [20], [21], typically suited for low-frequency wave components, there are two main approaches for wave manipulation. The first one assumes an analytical solution to the wave equation, which is a partial differential equation (PDE) [22]. If such an analytical solution can be obtained, it is used in conjunction with gradient based optimization (GBO) techniques to perform optimization over a system parameters [8], [23]–[25], towards a desired control objective. However, in general, an analytical solution to a PDE with arbitrary initial and boundary conditions is unknown, which leads to simplifying and limiting assumptions of the problem [8].

Another approach is based on machine learning and data-driven methods. The existing data-driven methods for PDE control suggest to learn wave dynamics from a large number of samples in order to predict its response to a control signal [26], [27]. Subsequently, this model can be used by a predictive controller to select optimal actions. Usually the learning of PDE dynamics is done by way of a black-box neural network model such as a (long short term memory) LSTM model [26], [27], Latent Evolution of PDE (LE-PDE)

* Equal contribution. Listing order is random.

¹ Computer Science Dept., Whitacre CoE, Texas Tech University; {trishshah, noam.smilovich, stas.tiomkin}@ttu.edu.

² Mechanical Engineering Dept., Davidson CoE, San Jose State University; {samer.gerges, feruza.amirkulova}@sjsu.edu.

[†] Corresponding Author.

[28], or, more advanced methods such as a Neural Ordinary Differential Equation (NODE) [29].

The main disadvantage of these black-box models is that there is no guarantee that a model will learn essential properties of the underlying dynamics of a PDE such as forces, energy dissipation, or viscosity. Consequently, due to the uninterpretable nature of black-box models, there is no clear way to represent physical properties of a wave within the model's architecture. That hampers the applicability of black-box models to important real-life applications such as wave propagation in open space, which requires physically meaningful dissipation of energy.

To address the gaps in the existing methods we design an interpretable framework for realizable control of acoustic waves. Our framework is specifically designed for adaptation to real-world environments, effectively addressing key challenges in wave manipulation. Namely: considering a realizable actuation mechanism which affects the wave in a sparse manner by inducing scattering is employed, actions are selected by an agent on a slower timescale than wave propagation, and observations of the wave state which could be feasibly collected by sensors are provided as input to our model. Furthermore, our framework addresses the lack of interpretability [30] of previous-ML based methods by explicitly constraining a learnable latent representation to abide by a one-dimensional (1D) wave equation. By construction, physical meaning is assigned to quantities within the latent space in the form of parameterized initial conditions, boundary conditions, and constraints (explained in Section V). Moreover, our framework allows for wave control in open spaces, which we achieve by the design of a trainable data-driven dissipation layer. To our best knowledge, the latter has not been demonstrated before, while it is an essential component in acoustics research, known as *perfectly matched layer* (PML) [31], [32].

We demonstrate the application of our framework to two important tasks in acoustics: sound suppression (minimization of scattered energy in the domain) and sound focusing (maximization of sound energy in a desired location). In our experiments we show that our method is superior to a standard ML dynamics model in these tasks, and competitive with a semi-analytical GBO method. Our work is pioneering research in dynamical control of acoustic waves. While this work focuses on establishing control of a particular type of PDE for wave phenomenon, it could be applied to controlling arbitrary systems governed by PDE such as wildfires, pandemics, and eventually climate [33]–[35].

II. PRIOR WORK

There have been several attempts to apply machine learning to the manipulation of complex phenomena governed by differential equations. Overall there has been less focus on designing systems that could be realized by real robots.

One approach used a surrogate dynamics model to enhance reinforcement learning for PDE-governed systems [26]. The

authors did not prioritize realizability in this work, assuming their agent could apply forces across the entire spatial domain of the 1D Kuramoto-Sivashinsky (KS) system—a non-realistic scenario, as fluid manipulation typically involves sparse control through rotating cylinders [27] or airfoils [36]. Additionally, their reliance on full state observability, with an LSTM model trained on high-resolution simulation data, limits its applicability to real-world systems, where only sensor data is available.

A realizable approach to using ML for fluid dynamics manipulation was demonstrated through MPC of a Navier-Stokes governed system [27]. The authors made specific choices to promote the realizability of their system, such as using sparse control through rotating cylindrical devices as well as assuming partial observability through sensors. Similar to the example of KS manipulation [26], this work also employed an LSTM architecture [27]. While the authors demonstrated the success of their method, the sample efficiency of such an architecture may not be optimal for modeling PDE dynamics. Furthermore, the dynamics of an LSTM are not interpretable and do not provide a straightforward pathway to incorporating domain knowledge from fluid dynamics.

Recent works have explored the incorporation of physics knowledge in ML dynamics models [37], [38]. These works and others [39] have shown that incorporating such knowledge into the training of neural networks leads to better generalization and sample efficiency. In their work, the authors utilized Physics-Informed Neural Networks (PINNs) [39] to model the dynamics of simple systems governed by ODE in response to control. They demonstrated that their PINN based dynamics models could effectively control a mechanical multi-link manipulator [37], the Van der Pol Oscillator [38], and the classic Four Tank System [38]. While these works demonstrated that it is possible to incorporate physics information into ML dynamics models, they are restricted to Ordinary Differential Equation (ODE) systems which are less general than PDEs are.

In the following sections, we present our approach which incorporates prior knowledge of PDEs into learnt dynamics models for the purpose of manipulating acoustic waves. In our work, we prioritize realizability, with the goal of enabling wave manipulation with real robots.

III. PRELIMINARIES

In this section, we introduce notations and provide an overview of the problem. The formal problem definition and its solution are provided in Sections IV and V.

A. Notations

We denote $t \in \mathbb{R}^+$, $x \in \mathbb{R}^{d_x}$ and $a(t) \in \mathbb{R}^{d_a}$ as the time coordinate, the space coordinate, and the agent's action, respectively. The partial derivative of a function w.r.t time is denoted by ∂_t . A differential operator, acting on a function $\zeta(x, t)$ in a domain $x \in \Omega$, is denoted by $\mathcal{N}(\zeta; d(t), \ell(x), t)$, where $d(t)$ and $\ell(x)$ represent functions controllable and uncontrollable by the agent, respectively. Both $d(t)$ and

$\ell(x)$ can affect the solution to the PDE, $\partial_t \zeta(x, t) = \mathcal{N}(\zeta(x, t); d(t), \ell(x), t)$. We represent the state of a robot observed at a particular time t_i as $d(t_i)$ whereas $a(t)$ represents actions applied over a time span $t \in [t_i, t_{i+1}]$.

B. Wave-Robot Coupled Environment

The simulated environment we consider in this work is defined by a PDE $\mathcal{N}$ for environment dynamics and an ODE F for the robot dynamics. The agent indirectly interacts with the environment through its actions, $a(t)$, which influence the robot, $d(t)$. The latter represents a dynamically changing boundary condition within the environmental domain, which evolves according to a dynamical control system, F :

$$\partial_t \zeta(x, t) = \mathcal{N}(\zeta(x, t); d(t), \ell(x), t) \quad \text{PDE} \quad (1)$$

$$a(t) = \pi(\zeta(x, t), d(t)) \quad \text{Policy} \quad (2)$$

$$\dot{d}(t) = F(d(t), a(t)) \quad \text{ODE} \quad (3)$$

where π is an agent's policy. In this setting the agent controls a low dimensional state of the robot, which in turn affects the function $\zeta(x, t)$ in the domain Ω . In the sequel we formally define an optimization problem for sample-based derivation of an optimal π from a partially observable function $\zeta(x, t)$. In the interaction model, the ODE controls the PDE creating nested dynamics with two different time scales.

C. Sparse Robotic Actuation

The separation between the agent actions, $a(t)$, and the environment ζ , through the robot, $d(t)$, allows for a physical realization of control of PDEs with a spatially sparse control. This is because it is not usually possible to construct a controller which is capable of directly acting on a PDE-governed environment at every point in space.

In this work, we consider a robotic controller which is capable of manipulating acoustic waves by inducing scattering. Specifically, the robot consists of a number of cylindrical scattering devices that the agent can control through adjustment of their position, radii, or both. In turn, the scattered wave energy produced by reflection of incident waves upon the scatterers can produce complex wave interference patterns, resulting in energy focusing or suppression [8].

D. Wave Sensing

In practice, a sensor which is capable of fully observing wave functions is not possible. Instead, real time total displacement and pressure fields can be observed on discretized grids through sensors such as microphone arrays and acoustic cameras [40]. In our simulated environment, the pixels of the sensor images are obtained by observing the displacement of the wave function at points on a uniformly spaced grid over Ω . We assume the agent can only partially observe the wave function, $\zeta(x, t)$, via its sensors, which we denote as images: $\text{Sensor}(\zeta(x, t)) = X(t) \in \mathbb{R}^{d_1 \times d_2}$.

In the next section, we define the control problem, and then propose a tractable method for its solution.

IV. PROBLEM DEFINITION

In the fully observable case, we can formally define the control problem as minimization of a cost function in Eq. (4) for $t \in [t_i, t_f]$, subject to the dynamics $\mathcal{N}$ of ζ and dynamics F of d . A shorthand notation is used $\zeta = \zeta(x, t)$:

$$\min_{a(\cdot)} \int_{t_i}^{t_f} \|\mathcal{T}(\zeta(x, t')) - \sigma^*(t')\|^2 dt', \quad (4)$$

$$\text{s.t.} \quad \begin{cases} \partial_t \zeta = \mathcal{N}(\zeta; d(t), \ell(x), t), \\ \dot{d}(t) = F(d(t), a(t)). \end{cases} \quad (5)$$

where $\mathcal{T}(\zeta(x, t)) = \sigma(t)$ is a function that computes a signal we wish to predict and control, such as scattered energy, at a particular location or in the entire domain, Ω . The goal of this control problem is to select actions which influence ζ to possess properties defined by a reference signal $\sigma^*(t)$.

This problem is challenging, because of the coupling between ODE and PDE in the constraint in Eq. (5), and the complexity of control of an entire function, ζ . Moreover, wave propagates at a much higher speed in comparison to the responsiveness of a physically realizable robot. That creates two different time scales in the above-mentioned coupling, which requires a particular consideration.

We assume the robot knows the dynamics, F , of its scatterers which define the time-varying boundary conditions for the wave, ζ . A direct solution to this problem is intractable, so we propose to learn a low-dimensional dynamical control system of ζ , which captures the essential properties of ζ , and to utilize this low dimensional dynamics for the design of control law.

V. PROPOSED SOLUTION

We propose to solve the problem defined in Section IV in two stages. First, we learn a 1D dynamical control system, $\mathbf{z}(\bar{x}, t)$, which captures the essential properties of full dynamics of wave ζ , such as energy distribution in space, scattering, excitation, and dissipation. Second, we derive optimal actions to control the energy distribution of ζ via its 1D model $\mathbf{z}(\bar{x}, t)$, which is used as a surrogate in the optimization of Eq. (4). A shorthand notation is used $\mathbf{z} = \mathbf{z}(\bar{x}, t)$.

A. Learning the Reduced Dimensional Model

Given an initial sensor observation, $X(t_i)$, robot state $d(t_i)$, and actions $a(t)$, the agent constructs $\mathbf{z}$ for $t \in [t_i, t_f]$, a 1D representation of the true function, ζ , in the reduced spatial coordinates, $\bar{x} \in \Gamma \subset \mathbb{R}$. Construction of $\mathbf{z}$ is accomplished by way of two encoders shown in Figure 2 which output functions defined over Γ : the wave encoder, W , and the robot encoder, D :

$$W_\mu : X(t) \rightarrow [\mathbf{g}(\bar{x}, t), l(\bar{x})] \quad \text{(Wave Encoder)} \quad (6)$$

$$D_\phi : d(t) \rightarrow c(\bar{x}, t) \quad \text{(Robot Encoder)} \quad (7)$$

where $c(\bar{x}, t)$ and $l(\bar{x})$ represent the environment components which are controllable and uncontrollable by the agent's actions, respectively. This way $c(\bar{x}, t)$ and $l(\bar{x})$ correspond to $d(t)$ and $\ell(x)$, respectively, in the original environment,

ζ . This separation is beneficial because it allows for explicit representation of energy excitation and dissipation properties of the original environment (cf., Figure 1) within $l(\bar{x})$. Additionally, W_μ outputs $\mathbf{g}(\bar{x}, t)$, which defines the initial conditions of $\mathbf{z}$. We denote the parameters of W_μ and D_ϕ by $\theta \doteq [\mu, \phi]$.

Given these definitions, we introduce the following optimization problem for $t \in [t_i, t_f]$:

$$\min_{\theta} \left\{ \int_{t_i}^{t_f} \|\mathcal{T}(\mathbf{z}(\bar{x}, t')) - \mathcal{T}(\zeta(x, t'))\|^2 dt' \right. \quad (8)$$

$$+ \int_{t_i}^{t_f} \int_{\Gamma} \|\partial_t \mathbf{z}(\bar{x}, t') - \mathcal{N}^\Gamma(\mathbf{z}(\bar{x}, t'); c_\phi(\bar{x}, t'), l_\mu(\bar{x}), t')\|^2 d\bar{x} dt' \quad (9)$$

$$\left. + \int_{\Gamma} \|\mathbf{z}(\bar{x}, t_i) - \mathbf{g}_\mu(\bar{x}, t_i)\|^2 d\bar{x} \right\} \quad (10)$$

$$\text{s.t. } \dot{d}(t) = F(d(t), a(t)),$$

where $\mathcal{N}^\Gamma$ is a PDE similar to $\mathcal{N}$ that operates on a 1D space Γ instead of Ω . Eq. (8) is the prediction loss for scattered energy. Eq. (9) and Eq. (10) are the PDE consistency and the initial condition losses, respectively. Notably, $\mathbf{z}$ implicitly depends on θ through learnable $\mathbf{g}_\mu$, and c_ϕ and l_μ .

We chose a deep convolutional neural network to represent W_μ , in Figure 2 (left), which maps the pixels of the sensor observations $X(t)$ onto functions over Γ . Similarly, the robot encoder is implemented by an artificial neural network D_ϕ , in Figure 2 (right), maps the robot's trajectory to latent control functions. Since the dynamics of the robot in Eq. (3) are known ahead of time, D_ϕ can construct $c(\bar{x}, t)$ from an initial robot state $d(t_i)$ and $a(t)$.

In our method, $\mathbf{z}$ is constructed using numerical integration of $\mathcal{N}^\Gamma$. Assuming prior knowledge of the governing equation of ζ in lower dimensional space, the encoded initial conditions, boundary conditions, and exogenous forces specify a solution $\mathbf{z}$. Relying on numerical integration to generate $\mathbf{z}$ has the unique benefit of guaranteeing that it is a solution to $\mathcal{N}^\Gamma$, while introducing no trainable parameters aside from the encoders W_μ and D_ϕ . This is in contrast to the approaches that require an additional network (e.g., NODE) to represent the dynamics $\mathcal{N}^\Gamma$ of $\mathbf{z}$, which increases sample and computational complexity.

Our method requires the following data for training:

$$\mathcal{I} = \{(X(t_i), d(t_i), a(t), \sigma(t))\}_{k=1}^N, \quad (11)$$

where $X(t_i)$, $d(t_i)$, $a(t)$, and $\sigma(t)$, are a sensor observation and a robot state observed at time t_i , and actions and corresponding target signal in time period $t \in [t_i, t_f]$. Algorithm 1 summarizes the learning of physics-informed model.

Once the robot learns how its sparse control affect the wave, ζ , through $\mathbf{z}$, optimal actions can be selected according to Model Predictive Control, as explained below.




Fig. 2. Encoding Scheme. Left: Sensor observation $X(t_i)$ of the PDE ζ is encoded by W to the latent initial condition $\mathbf{g}(\bar{x}, t_i)$ and exogenous function $l(\bar{x})$. Right: sequences of robot configurations, $d(t_i), d(t_{i+1}), \dots$, produced by agent's actions, $a_i(t)$ are individually encoded to latent control functions $c(\bar{x}, t_i), c(\bar{x}, t_{i+1}), \dots$.

Algorithm 1 Learning Physics-Informed Model

Require: Training data, $\mathcal{I}$, (cf., Eq. (11)), parameters of W_μ and D_ϕ , denoted by $\theta = (\mu, \phi)$, and design dynamics F .

- 1: Randomly initialize θ
- 2: **repeat**
- 3: **Sample** a data point from $\mathcal{I}$: $\{X(t_i), d(t_i), a(t), \sigma(t) \sim \mathcal{I}\}$ {Eq. (11)}
- 4: **Encode** a Sensor Reading (wave image): $\{\mathbf{g}_\mu(\bar{x}, t_i), l_\mu(\bar{x}) = W_\mu(X(t_i))\}$ {Eq. (6)}
- 5: **Integrate** a Design Trajectory from $d(t_i)$: $d(t) = d(t_i) + \int_{t_i}^t F(d(t'), a(t')) dt'$
- 6: **Encode** a Design Trajectory: $\{c_\phi(\bar{x}, t) = D_\phi(d(t))\}$ {Eq. (7)}
- 7: **Integrate** a Latent Trajectory from $\mathbf{g}(\bar{x}, t_i)$: $\mathbf{z}_\theta(\bar{x}, t) = \mathbf{g}_\mu(\bar{x}, t_i) + \int_{t_i}^t \mathcal{N}^\Gamma(\mathbf{z}(\bar{x}, t'); c_\phi(\bar{x}, t'), l_\mu(\bar{x}), t') dt'$
- 8: **Calculate** the Prediction Loss: $\{\text{loss}(\theta) = \int_{t_i}^{t_f} (\mathcal{T}(\mathbf{z}_\theta(\bar{x}, t')) - \sigma(t'))^2 dt'\}$ {Eq. (8)}
- 9: **Perform Gradient Descent** on $\text{loss}(\theta)$
- 10: **until** Convergence
- 11: **return** W_μ^*, D_ϕ^*

B. Model Predictive Control of Wave Energy

We apply MPC [41]–[43] for control of wave energy over the time interval $[t_i, t_f]$, by dividing a control trajectory, $a(t)$, to a sequence of N_a piece-wise constant actions, $a(t) = \{a_\tau : \text{if } t_\tau \leq t < t_{\tau+1} \forall \tau\}$. This time discretization leads to a series of time intervals $\Delta_\tau = (t_{\tau+1} - t_\tau)$, with boundary conditions $t_{\tau=1} = t_i$ and $t_{\tau=N_a+1} = t_f$, which reflects the slow time scale of robotic actuation in comparison to the high speed of wave propagation. The objective below with the reference control signal $\sigma^*(t)$, allows for the derivation

of the optimal control sequence $\{a_\tau^*\}_{\tau=1}^{N_a}$:

$$\min_{\{a_\tau\}_{\tau=1}^{N_a}} \int_{t_i}^{t_f} \|\mathcal{T}(\mathbf{z}(\bar{x}, t')) - \sigma^*(t')\|^2 dt' + \beta \sum_{\tau=1}^{N_a} \|a_\tau\|^2 \Delta_\tau \quad (12)$$

$$\text{s.t.} \begin{cases} \partial_t \mathbf{z}(\bar{x}, t) = \mathcal{N}^\Gamma(\mathbf{z}(\bar{x}, t); c_{\phi^*}(\bar{x}, t), l_{\mu^*}(\bar{x}, t)) \\ \mathbf{z}(\bar{x}, t_i) = \mathbf{g}_{\mu^*}(\bar{x}, t_i) \end{cases} \quad (13)$$

where the constraints are the latent 1D dynamics with the optimal parameters $\theta^* = [\mu^*, \phi^*]$ in Eq. (9)-(10), in addition to F , which is known to the agent. The agent's power is constrained through hyperparameter β .

To the best of our knowledge, this proposed solution has not been studied prior to our work, while the most relevant works [37], [38] considered ODEs only, which is a restricted class of dynamics. Moreover, there are no efficient data-driven methods for controlling acoustic PDEs in open space with sparse control and partially observable state.

VI. EXPERIMENTS

The goal of our experiments is to validate that the AEM approach can effectively exploit prior knowledge about the physical properties of a system governed by an acoustic PDE. We demonstrate this by first showing that the model can predict scattered energy in the domain for action sequences extending far beyond its training horizon. Next, we show that AEM can be effectively used with MPC on two benchmark tasks in acoustics: energy suppression and focusing. We compare our method to a ML approach (NODE) and a ground truth solution derived by GBO.

A. Setting

The robotic actuator we consider in this environment is a configuration of cylindrical scatterers. We assume that an agent has the ability to manipulate by $a(t)$, the radii and/or position of $d(t)$, which create time-variant boundary conditions for the domain, $\Omega(d(t))$ of wave propagation:

$$\begin{aligned} \partial_t^2 \zeta(x, t) &= c^2 \partial_x^2 \zeta(x, t) + f(x, t) \in \Omega(d(t)) \\ \dot{d}(t) &= F(d(t), a(t)), \end{aligned}$$

where $\zeta(x, t)$ and $d(t)$ evolve in two different timescales: fast and slow for wave and robot, respectively.

We conduct the experiments in a simulated open space with size $x \in [-15.0 \text{ m}, 15.0 \text{ m}] \times [-15.0 \text{ m}, 15.0 \text{ m}]$, where the excitation source creates a spherical wave in the domain (cf., Figure 1). Simulation of open space without wave reflection from the boundaries is achieved through a PML [31], which dissipates wave energy as it leaves Ω . Within $\mathcal{N}^\Gamma$ we include a trainable PML through $l(\bar{x})$ which allows our agent to explicitly model energy dissipation.

B. Robot Actuation

We vary the number of scatterers, denoted as M , and test several configurations for robot actions. We use **R**, **P**, and **F** to represent "Radii Adjustment", "Positional Adjustment", and "Full Adjustment" respectively:

TABLE I

STEADY-STATE SCATTERED ENERGY IN THE SIMULATED ENVIRONMENT. THE STEADY-STATE ENERGY IS MEASURED FROM 0.10 TO 0.20 SECONDS OF TIME IN THE ENVIRONMENT. STATISTICS ARE COMPUTED OVER 12 RUNS WITH RANDOMIZED INITIAL ROBOT CONDITIONS AND FIXED SOURCE LOCATION AT $(-10.0, 0.0)$.

Method	Random	NODE	AEM	GBO
Focusing (Higher is Better)				
M=1 (P)	0.12 ± 0.03	0.17 ± 0.07	2.47 ± 1.91	N/A
M=2 (P)	0.35 ± 0.21	1.22 ± 1.22	7.54 ± 2.24	N/A
M=4 (P)	1.05 ± 0.46	1.45 ± 0.92	5.94 ± 1.90	N/A
M=2 (F)	0.87 ± 0.60	6.19 ± 1.51	8.49 ± 1.37	N/A
Suppression (Lower is Better)				
M=1 (P)	3.00 ± 0.72	2.86 ± 0.53	1.08 ± 0.18	1.01
M=2 (P)	6.87 ± 1.89	4.40 ± 0.84	1.98 ± 0.38	2.13
M=4 (P)	10.04 ± 1.84	10.16 ± 2.11	3.78 ± 0.74	3.53
M=2 (F)	5.41 ± 1.33	3.48 ± 1.01	0.38 ± 0.02	0.38

- 1) **Radii Adjustment (R)** (Ring Configuration): This setup features $M = 19$ scatterers arranged in a ring formation with fixed positions. The actions in this experiment involve solely adjusting the radii of the scatterers.
- 2) **Positional Adjustment (P)**: Here, scatterers have fixed radii, and the agent can only change their positions. We tested this configuration with $M = 1, 2$, and 4 scatterers.
- 3) **Full Adjustment (F)**: In this scenario, the agent can modify both the positions and radii of the scatterers. We demonstrated performance for $M = 2$ in this configuration.

C. Evaluation Tasks

We evaluate our model in both prediction and control.

- **Long-Term Prediction**: We compare predicted scattered energy over a horizon of 200 steps (much beyond the training horizon) between our model, the baseline model, and the ground truth.
- **Focusing**: Concentrate the scattered energy within the upper-right quadrant of the domain by the robots actions.
- **Suppression**: Minimize the total scattered energy within the domain by the robots actions.

These tasks represent standard benchmarks and are frequently used in acoustics research [8], [44]. We use Model Predictive Control in both manipulation tasks.

D. Results

1) **Long-Term Prediction of Scattered Energy**: We assess the robustness of our model by comparing its long-term prediction capabilities to a NODE-based approach, as well as an AEM-based approach that does not incorporate a trainable PML inside $\mathcal{N}^\Gamma$. Figure 3 shows that AEM (PML) follows the ground truth scattered energy compared to NODE and AEM (No PML) which diverge. Our model without a PML experiences rapid energy buildup because




Fig. 3. Long-term prediction of scattered energy over an episode of 0.2 seconds (200 action steps) for the ring configuration (**R**). The training horizon for the models is 20 actions, which demonstrate generalization of the model to unseen in training horizons.

it is unable to dissipate energy in z . The NODE model also diverges from the ground truth signal and provides no pathway to explicitly model energy dissipation. These results highlight the critical role of including a PML in $\mathcal{N}^\Gamma$ when ζ occurs in open space. Table I shows that AEM achieves similar results to the ground truth results by GBO in the 'Suppression' task, when a semi-analytical solution by GBO is known in the literature, otherwise it is marked 'N/A'.

2) *Focusing and Suppression of Scattered Energy*: We utilize MPC through optimization of Eq. (12) for both benchmark tasks. Figure 4 shows the results for the fully adjustable configuration with two scatterers (F, $M=2$). The plot demonstrates the advantages of our approach in both focusing (top) and suppressing (bottom) scattered energy. In terms of steady-state performance, our method achieves higher focused energy and lower suppressed energy. Furthermore, our approach exhibits better transient response time, reaching steady-state faster than the other approaches. Similar trends were observed across other scatterer configurations.

We compare the performance of our model against the NODE-based approach with both models having an equal parameter count of 2.7 million. Table I summarizes the results of our MPC experiments. A random policy, and gradient-based optimization approaches are also included as part of the comparison. Our AEM model achieves good performance on both benchmarks: for focusing, it yields higher focused energy while maintaining comparable error rates; for suppression, it attains lower suppressed energy than NODE and Random control with the added benefit of reduced error compared to other methods. For some of the cases AEM achieves slightly better performance than the steady-state GBO solution.




Fig. 4. Control of scattered energy by focusing (top) and suppression (bottom). Both the positions and radii of two fully adjustable scatterers are varied (F, $M=2$). Source is fixed at location $(-10, 0)$ throughout the simulations. Shading represents ± 0.5 standard deviation from the mean. The mean (solid lines) and variance are calculated with 12 runs with randomized initial conditions. The dashed line is for the ground truth by GBO.

VII. CONCLUSIONS

This work extends classical object manipulation in robotics to the domain of acoustic waves. It opens new research opportunities at the intersection of robotics and acoustics, with potential applications in such areas as ultrasound cutting, energy harvesting, and the design of new acoustic materials. This work provides proof of concept for the implementation of a physically realizable robot for wave manipulation. Importantly, the methodology is not limited to acoustics—it can be adapted to other domains governed by PDEs.

We highlight the unique features of our approach, including sparse robotic actuation under partial observability and a fully interpretable, physics-constrained solution. This allows for a trainable PML embedded within the AEM's latent space ensuring robust and accurate long-term predictions. The proposed AEM allows to find a solution, which can be derived only in particular cases with GBO. That makes AEM a strong candidate in robot learning for PDE manipulation, where interpretable and reliable methods are required.

ACKNOWLEDGEMENTS

The authors thank P. E. Leser from NASA for the valuable suggestions. ST was supported in part by NSF (2246221) and PAZI (195-2020). TS was supported in part by the Koh Family Scholarship. TS and NS acknowledge the RSCA Awards. Part of this work was done when TS and NS were graduate students at SJSU.

REFERENCES

- [1] J. Zhu, A. Cherubini, C. Dune, D. Navarro-Alarcon, F. Alambeigi, D. Berenson, F. Ficuciello, K. Harada, J. Kober, X. Li, J. Pan, W. Yuan, and M. Gienger, "Challenges and outlook in robotic manipulation of deformable objects," *IEEE Robotics & Automation Magazine*, vol. 29, no. 3, pp. 67–77, 2022.
- [2] J. Matas, S. James, and A. J. Davison, "Sim-to-real reinforcement learning for deformable object manipulation," in *Proceedings of The 2nd Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, A. Billard, A. Dragan, J. Peters, and J. Morimoto, Eds., vol. 87. PMLR, 29–31 Oct 2018, pp. 734–743. [Online]. Available: <https://proceedings.mlr.press/v87/matas18a.html>
- [3] J. Sanchez, J.-A. Corrales, B.-C. Bouzgarrou, and Y. Mezouar, "Robotic manipulation and sensing of deformable objects in domestic and industrial applications: a survey," *The International Journal of Robotics Research*, vol. 37, no. 7, pp. 688–716, 2018. [Online]. Available: <https://doi.org/10.1177/0278364918779698>
- [4] M. R. Pedersen, L. Nalpantidis, R. S. Andersen, C. Schou, S. Bøgh, V. Krüger, and O. Madsen, "Robot skills for manufacturing: From concept to industrial deployment," *Robotics and Computer-Integrated Manufacturing*, vol. 37, pp. 282–291, 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0736584515000575>
- [5] A. Hamed, S. C. Tang, H. Ren, A. Squires, C. Payne, K. Masamune, G. Tang, J. Mohammadpour, and Z. T. H. Tse, "Advances in haptics, tactile sensing, and manipulation for robot-assisted minimally invasive surgery, noninvasive surgery, and diagnosis," *Journal of Robotics*, vol. 2012, no. 1, p. 412816, 2012. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1155/2012/412816>
- [6] P. Boonvisut and M. C. Çavuşoğlu, "Estimation of soft tissue mechanical parameters from robotic manipulation data," *IEEE/ASME Transactions on Mechatronics*, vol. 18, no. 5, pp. 1602–1611, 2013.
- [7] P. A. Martin, *Multiple Scattering: Interaction of Time-harmonic Waves with N Obstacles*. Cambridge University Press, New York, 2006.
- [8] F. A. Amirkulova and A. N. Norris, "The gradient of total multiple scattering cross-section and its application to acoustic cloaking," *Journal of Theoretical and Computational Acoustics*, p. 1950016, 03 2020.
- [9] A. N. Norris and J.-M. Conoir, "Multiple scattering by cylinders immersed in fluid: High order approximations for the effective wavenumbers," *The journal of the acoustical society of america*, vol. 129, no. 1, pp. 104–113, 2011.
- [10] V. Del Grosso and C. Mader, "Speed of sound in pure water," *the Journal of the Acoustical Society of America*, vol. 52, no. 5B, pp. 1442–1446, 1972.
- [11] Z. Chen, G. Wang, and C. Lim, "Artificially engineered metaconcrete with wide bandgap for seismic surface wave manipulation," *Engineering Structures*, vol. 276, p. 115375, 2023.
- [12] S. Y. Kim, W. Lee, J. S. Lee, and Y. Y. Kim, "Longitudinal wave steering using beam-type elastic metagratings," *Mechanical Systems and Signal Processing*, vol. 156, p. 107688, 2021.
- [13] T. Ilovitsh, A. Ilovitsh, J. Foiret, B. Z. Fite, and K. W. Ferrara, "Acoustical structured illumination for super-resolution ultrasound imaging," *Communications biology*, vol. 1, no. 1, p. 3, 2018.
- [14] A. Ozelik, J. Rufo, F. Guo, Y. Gu, P. Li, J. Lata, and T. J. Huang, "Acoustic tweezers for the life sciences," *Nature methods*, vol. 15, no. 12, pp. 1021–1028, 2018.
- [15] H.-O. Kreiss, N. A. Petersson, and J. Yström, "Difference approximations for the second order wave equation," *SIAM Journal on Numerical Analysis*, vol. 40, no. 5, pp. 1940–1967, 2002.
- [16] R. J. LeVeque, *Finite difference methods for ordinary and partial differential equations: steady-state and time-dependent problems*. SIAM, 2007.
- [17] J. B. Angel, J. W. Banks, A. M. Carson, and W. D. Henshaw, "Efficient upwind finite-difference schemes for wave equations on overset grids," *SIAM Journal on Scientific Computing*, vol. 45, no. 5, pp. A2703–A2724, 2023.
- [18] B. Cockburn, G. E. Karniadakis, and C.-W. Shu, *Discontinuous Galerkin methods: theory, computation and applications*. Springer Science & Business Media, 2012, vol. 11.
- [19] C.-W. Shu, "High order weno and dg methods for time-dependent convection-dominated pdes: A brief survey of several recent developments," *Journal of Computational Physics*, vol. 316, pp. 598–613, 2016.
- [20] J. Morris, W. Wang, T. Plaisted, C. J. Hansen, and A. V. Amirkhizi, "Optimizing graded metamaterials via genetic algorithm to control energy transmission," *International Journal of Mechanical Sciences*, p. 108775, sep 2023.
- [21] W. Wang, "Reduced order modeling and design of 3d printable mechanical metamaterials," Ph.D. dissertation, University of Massachusetts Lowell, 2023.
- [22] R. Courant, K. Friedrichs, and H. Lewy, "On the partial difference equations of mathematical physics," *IBM Journal of Research and Development*, vol. 11, no. 2, pp. 215–234, 1967.
- [23] F. Amirkulova, S. Gerges, and A. Norris, "Sound localization through multi-scattering and gradient-based optimization," *Mathematics*, vol. 9, no. 22, pp. 2862, 10.3390/math9222862, nov 2021.
- [24] F. A. Amirkulova, S. Gerges, and A. N. Norris, "Broadband acoustic lens design by reciprocity and optimization," *JASA Express Letters*, vol. 2, no. 2, p. 024005, feb 2022.
- [25] S. Gerges and F. Amirkulova, "The gradient of absolute pressure and its application to sound localization," (In Preparation), 2024.
- [26] S. Werner and S. Peitz, "Learning a model is paramount for sample efficiency in reinforcement learning control of pdes," *ArXiv*, vol. abs/2302.07160, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:256846832>
- [27] K. Bieker, S. Peitz, S. L. Brunton, J. N. Kutz, and M. Dellnitz, "Deep model predictive flow control with limited sensor data and online learning," *Theoretical and Computational Fluid Dynamics*, vol. 34, no. 4, pp. 577–591, 03 2020. [Online]. Available: <https://doi.org/10.1007/s00162-020-00520-4>
- [28] T. Wu, T. Maruyama, and J. Leskovec, "Learning to accelerate partial differential equations via latent global evolution," in *Proceedings of the 36th International Conference on Neural Information Processing Systems*, ser. NIPS '22. Red Hook, NY, USA: Curran Associates Inc., 2024.
- [29] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural ordinary differential equations," *Advances in neural information processing systems*, vol. 31, 2018.
- [30] Y. Zhang, P. Tiño, A. Leonardis, and K. Tang, "A survey on neural network interpretability," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 5, no. 5, pp. 726–742, 2021.
- [31] J.-P. Berenger, "A perfectly matched layer for the absorption of electromagnetic waves," *Journal of computational physics*, vol. 114, no. 2, pp. 185–200, 1994.
- [32] S. G. Johnson, "Notes on perfectly matched layers (pmls)," 2021. [Online]. Available: <https://arxiv.org/pdf/2108.05348>
- [33] M. Ángel Javaloyes, E. Pendás-Recondo, and M. Sánchez, "A general model for wildfire propagation with wind and slope," *SIAM Journal on Applied Algebra and Geometry*, vol. 7, no. 2, pp. 414–439, 05 2023. [Online]. Available: <https://doi.org/10.1137%2F22m1477866>
- [34] S. H. Schneider and R. E. Dickinson, "Climate modeling," *Reviews of Geophysics*, vol. 12, no. 3, pp. 447–493, 1974. [Online]. Available: <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/RG012i003p00447>

- [35] F. Majid, A. M. Deshpande, S. Ramakrishnan, S. Ehrlich, and M. Kumar, "Analysis of epidemic spread dynamics using a pde model and covid-19 data from hamilton county oh usa," *IFAC-PapersOnLine*, vol. 54, no. 20, pp. 322–327, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2405896321022382>
- [36] L.-W. Chen, C.-Y. Xu, and X.-Y. Lu, "Numerical investigation of the compressible flow past an aerofoil," *Journal of Fluid Mechanics*, vol. 643, p. 97–126, 2010.
- [37] J. Nicodemus, J. Kneifl, J. Fehr, and B. Unger, "Physics-informed neural networks-based model predictive control for multi-link manipulators," *IFAC-PapersOnLine*, vol. 55, no. 20, pp. 331–336, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2405896322013118>
- [38] E. A. Antonelo, E. Camponogara, L. O. Seman, J. P. Jordanou, E. R. de Souza, and J. F. Hübner, "Physics-informed neural nets for control of dynamical systems," *Neurocomputing*, vol. 579, p. 127419, 2024.
- [39] M. Raissi, P. Perdikaris, and G. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0021999118307125>
- [40] J. A. Bocanegra, D. Borelli, T. Gaggero, E. Rizzuto, and C. Schenone, "A novel approach to port noise characterization using an acoustic camera," *Science of The Total Environment*, vol. 808, p. 151903, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0048969721069795>
- [41] P. Bakaráč and M. Kvasnica, "Fast nonlinear model predictive control of a chemical reactor: a random shooting approach," *Acta Chimica Slovaca*, vol. 11, no. 2, pp. 175–181, 2018.
- [42] J. Richalet, A. Rault, J. Testud, and J. Papon, "Model predictive heuristic control: Applications to industrial processes," *Automatica*, vol. 14, no. 5, pp. 413–428, 1978. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0005109878900018>
- [43] M. Morari, C. E. Garcia, and D. M. Prett, "Model predictive control: Theory and practice," *IFAC Proceedings Volumes*, vol. 21, no. 4, pp. 1–12, 1988, iFAC Workshop on Model Based Process Control, Atlanta, GA, USA, 13–14 June. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/B978008035735500061>
- [44] A. N. Norris, "Acoustic cloaking theory," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 464, no. 2097, pp. 2411–2434, 04 2008. [Online]. Available: <https://doi.org/10.1098%2Frspa.2008.0076>