



StableLev: Data-Driven Stability Enhancement for Multi-Particle Acoustic Levitation

Lei Gao
lei.gao.20@ucl.ac.uk
Department of Computer Science
University College London
London, United Kingdom

Giorgos Christopoulos
georgios.christopoulos.19@ucl.ac.uk
Department of Computer Science
University College London
London, United Kingdom

Prateek Mittal
prateek.mittal@ucl.ac.uk
Department of Computer Science
University College London
London, United Kingdom

Ryuji Hirayama
r.hirayama@ucl.ac.uk
Department of Computer Science
University College London
London, United Kingdom

Sriram Subramanian
s.subramanian@ucl.ac.uk
Department of Computer Science
University College London
London, United Kingdom

ABSTRACT

Acoustic levitation is an emerging technique that has found application in contactless assembly and dynamic displays. It uses precise phase control in an ultrasound transducer array to manage the positions and movements of multiple particles. Yet, maintaining stable mid-air particles is challenging, with unexpected drops disrupting the intended motion and position. Here, we present StableLev, a data-driven pipeline for the detection and amendment of instabilities in multi-particle levitation. We first curate a hybrid levitation dataset, blending optimized simulations with labels based on actual trajectory outcomes. We then design an AutoEncoder to detect anomalies in the simulated data, correlating closely with observed particle drops. Finally, we reconstruct the acoustic field at anomaly regions to improve particle stability and experimentally demonstrate successful dynamic levitation for trajectories within our dataset. Our work provides new insights into multi-particle levitation and enhances its robustness, which will be valuable in a wide range of applications.

CCS CONCEPTS

• **Human-centered computing** → **User interface programming**; • **Computing methodologies** → *Machine learning*.

KEYWORDS

acoustic levitation, stability, levitation dataset, anomaly detection

ACM Reference Format:

Lei Gao, Giorgos Christopoulos, Prateek Mittal, Ryuji Hirayama, and Sriram Subramanian. 2024. StableLev: Data-Driven Stability Enhancement for Multi-Particle Acoustic Levitation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3613904.3642286>



This work is licensed under a Creative Commons Attribution International 4.0 License.

CHI '24, May 11–16, 2024, Honolulu, HI, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0330-0/24/05
<https://doi.org/10.1145/3613904.3642286>

1 INTRODUCTION

Sound wave harnessing to levitate and manipulate objects has captured the interest of the Human-Computer Interaction (HCI) community due to its revolutionary possibilities. Acoustic levitation uses acoustic radiation pressure to create mid-air traps for levitating objects. This is achieved through phase retrieval algorithms, such as IBP [24], NAIVE [34], and GS-PAT [34], which shape the activation signals for a phased array of transducers (PAT).

This technology has spurred novel applications within the HCI and Graphics community in the past 10 years, like levitation-based volumetric display [18, 28, 31, 34], data physicalisation [13, 30], physical interaction with mid-air contents [2, 20], contactless assembly and printing [7, 8], etc. Platforms like OpenMPD [27] have democratized access to levitation systems, offering a comprehensive software tool for Unity-based application development. Yet, as we continue to innovate and push the boundaries of displays, assemblies, and interactions at the dynamic application level, there's a growing demand for more robust levitation to achieve widely accessible and user-friendly applications, especially given the increasing number of levitation points and evolving trajectories in the real world, which puts forward higher standards in phase retrieval algorithms to meet the demands of more complex levitation applications.

Existing multi-point phase retrieval algorithms predominantly rely on simulated outcomes and overlook potential discrepancies when applied to real-world dynamic levitation scenarios. Dynamic levitation introduces myriad complexities that lead to unpredictable failures. Recent empirical studies, such as those by ArticLev [8] and DataLev [13], underscore this challenge, reporting variability in success rates during assembly and animation phases. These findings illuminate the pressing need for enhanced strategies, especially as the complexity of 3D movements escalates with an increasing number of traps.

A significant hurdle in advancing the study of dynamic multi-point levitation has been the absence of a dedicated dataset to foster rigorous analysis and algorithm development. Recognizing this gap, our initial endeavor is the formulation of a comprehensive dataset. With over 180,000 data points sourced from two distinct phase retrieval solvers (NAIVE and GSPAT), this dataset serves as the backbone of our data-driven approach. By combining the

insights of our study with the potential of this dataset, we aspire to spark a surge in innovative acoustic levitation research, driving both dataset expansion and the exploration of more novel research avenues.

Building on this dataset, we introduce "StableLev", a multi-particle stability prediction and enhancement pipeline tailored to harness the power of dynamic levitation data. With the goal of pinpointing and subsequently rectifying unstable segments within levitation trajectories, StableLev employs a sophisticated combination of recurrent neural networks (RNN) like long short-term memory networks (LSTM) or gated recurrent unit (GRU) with AutoEncoders (AE) and variational AutoEncoders (VAE). This intricate model architecture yields a prediction performance of approximately 90%, improving the reliability and precision of dynamic levitation stability predictions.

In this paper, we make three key contributions that advance the pursuit of stable and reliable acoustic levitation:

- (1) The curation of a first-of-its-kind comprehensive dataset, amassing over 180,000 data points from two distinct solvers, acting as a catalyst for data-driven approaches in future research.
- (2) The development and introduction of "StableLev," a state-of-the-art stability prediction and enhancement pipeline. By integrating RNNs with VAE and AE, our model achieves a performance of 90% (F -score = 0.9) in predicting and rectifying unstable levitation trajectories.
- (3) A series of detailed evaluations demonstrating StableLev's efficacy and applicability in real-world scenarios, ensuring dynamic levitation with unmatched stability.

These advancements not only enhance the robustness of current acoustic levitation systems but also provide a robust foundation for future innovations in the field.

2 RELATED WORKS

2.1 Applications with Acoustic Levitation

Recent advancements in acoustic levitation have positioned it as a versatile interface across diverse fields. For the purpose of physical displays, floating charts [30] employed the positions of expanded polystyrene (EPS) particles to depict mid-air scatter plots, effectively transforming static data into interactive, tangible forms. This novel representation demonstrated the ability of acoustic levitation to enable immersive data physicalizations. To represent more intricate shapes, both LeviProps [28] and ArticLev [8] introduced optimized levitation structures and assembly processes. They integrated complex primitives such as threads and fabrics, offering the possibility of multi-material, dynamic physical displays. Moreover, DataLev [13] harnessed levitation to produce reconfigurable, multimodal data physicalizations with enhanced materiality. By combining different materials, DataLev not only presents data but also provides multi-sensory feedback, enriching the user experience.

LeviCursor [2] introduced a mechanism for manipulating and stabilizing a levitating particle. Focusing on indirect interactions, it offers a unique method where distance-based interactions between a finger and the particle create a way to select and interact with levitated particles. TipTrap [20] improved the interaction by enabling closer proximity of the finger to the levitated particle,

utilizing sound scattering from the finger to create a levitation trap for direct, co-located interaction.

Leveraging the benefits of contactless manipulation, acoustic levitation has found novel applications in food delivery and 3D printing. In food delivery systems [38], acoustic levitation has been used to levitate and deliver food (like miniature burgers) as well as drinks (like gin and tonic), promising to redefine the gastronomic experience by allowing chefs to play with food textures and presentations [39]. Additionally, in the domain of 3D printing, it aids in maneuvering UV resin and sticks, opening doors to innovative design possibilities and structures that were previously challenging or impossible to achieve [7].

2.2 Acoustic Levitation Principles and Advances

Initial acoustic levitation studies employed a single ultrasonic transducer on a Langevin horn directing towards a reflective surface [42]. These "single-axis levitators" produce standing wave acoustic fields, allowing small particles to levitate at points with minimal acoustic pressure. Here, the Gor'kov potential (U) is minimum, and the acoustic radiation force $\mathbf{F} = -\nabla U$, given by the potential's negative gradient [4], becomes zero, thus holds particles in the trapping positions. Particles that are close to standing wave antinodes get pushed towards the nearest potential well [1].

Recently, phased arrays of transducers (PAT) have replaced the Langevin horn in acoustic levitation devices [29]. They offer more acoustic energy, arbitrary movement of levitated particles beyond the antinodes of a fixed standing wave, and the generation of multiple acoustic traps [24]. A PAT-based levitator includes a total of N transducers arranged into an array or set of arrays [25]. The phase delays $\phi \in \mathbb{R}^N$ of these transducers give rise to different complex acoustic fields $\mathbf{p} \in \mathbb{C}^M$ at M points in space. For the m^{th} point in space, $m = [1, \dots, M]$, this complex pressure is given by

$$p_m = \sum_{n=1}^N f_{m,n} \cdot e^{i\phi_n} \quad (1)$$

where $f_{m,n}$ describes the acoustic transmission of a transducer n to a point m , given by the piston model [34]. Using phase retrieval algorithms, one can optimize the transducers' phases ϕ to generate acoustic fields that include local standing wave patterns that correspond to acoustic traps (Figure 1). To guarantee the generation of proper sinusoidal standing wave patterns, i.e. the interchange between acoustic pressure minima and maxima, it has been quite common to compute phases that maximize the Laplacian of the Gor'kov potential (known as trapping stiffness) [7, 25, 28]. When trapping stiffness is maximum, the acoustic radiation forces converge to the trapping positions.

However, many applications such as volumetric displays [12, 18, 34] require high update rates for the transducer phases (i.e., $>10\text{kHz}$), and computing trapping stiffness via finite difference derivatives involves computing acoustic pressure at many points per trap (i.e., 55 points in [17]) using Eq 1. Instead, direct minimization of the Gor'kov potential (or its simplification) that only considers spatial derivatives in the principal acoustic wave propagation direction [17] has provided more viable alternatives for fast multi-point levitation even in the presence of scattering objects.

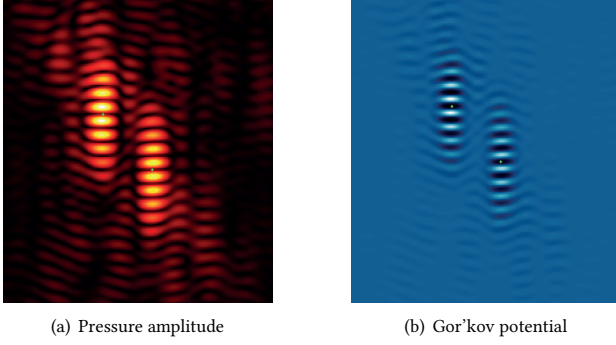


Figure 1: Illustration of the simulated sound field generated by Top-Bottom PATs with Two Traps: (a) Depicts the pressure amplitude distribution around the trap. (b) Depicts the distribution of Gor'kov Potential around the trap. Particles (green dots in (a), (b)) are suspended at the sound wave's antinode, coinciding with the minimum Gor'kov potential.

The most computationally efficient acoustic field optimization for levitation is to compute acoustic pressure only at the trapping location, aiming to create high amplitude focal points that can be converted into acoustic tweezers via standard levitation signatures [25]. In symmetric levitation setups (e.g. when no rigid obstacles are involved), the acoustic pressure of focal points linearly correlates with their trapping stiffness after the application of levitation signatures[34].

2.3 Acoustic Levitation Stability

While advances have been made in acoustic levitation, particles frequently fall or stray from desired traps during experiments. Single-axis levitation studies have explored particle stability concerning size and medium viscosity [10], and others have analyzed levitated droplet oscillation [16]. However, while PAT levitation largely draws from single-axis principles that assume a plane standing wave field, the exploration of instabilities in dynamic levitation, especially involving multiple particles moving in free space, remains limited.

Recently, it has been shown that for large displacements and thus fast movements of a single particle in a PAT levitator, the assumption that forces are linear in the vicinity of the trapping position is not valid anymore. Instead, the trapping stiffness becomes non-linear, which explains the period-doubling bifurcation of levitated particles [11]. Furthermore, in contrast to single-point levitation, which usually involves shifts of a focusing phase map to move particles in 3D space, multi-point levitation requires optimization of transducers' phases as described in the previous section. So far, this optimization has been time-invariant, i.e. the transducer phases ϕ^t for different movement time steps $t = [1, \dots, T]$, are optimized independently, which leads to high phase changes $\Delta\phi^t$ among transducers between movement frames. High abrupt phase changes lead to amplitude fluctuations in the transducers' emission [37]. That is, the delivered acoustic energy diminished and is not sufficient to keep particles levitated in mid-air.

On the other hand, current research on acoustic levitation has only resolved particle misplacements but not particle drops. In an HCI context, LeviCursor used a motion capture system to avoid particle placements at the (weaker) secondary traps of the acoustic tweezer standing wave pattern [2], while LeviProps performed a simulated annealing to find the trapping positions of highest trapping stiffness to hold an acoustically transparent fabric in mid-air [28]. Finally, other studies have been occupied with numerical optimization of single-particle trajectories [31], so that the showcased experimental trajectories better match the desired ones, effectively reducing particle misplacements.

3 LEVITATION DATASET

In this section, we detail the creation of our hybrid levitation dataset, which combines simulated and experimental data. This approach addresses the limitations of existing phase optimizations that excel in static simulations but falter in dynamic scenarios. We conducted levitation experiments on various multi-particle trajectories, recording particle positions and categorizing outcomes. Our research offers the first extensive levitation dataset based on rigorous feature extraction. This paves the way for deeper insights into acoustophoretic platforms and inspires future ML-driven applications.

Figure 2 graphically depicts our dataset generation process. This section outlines our levitation setup and the creation of multi-particle trajectories used for both experiments and model training/evaluation. We subsequently explain the intricacies of acoustic field propagation, optimization, and the hardware setup for capturing particle positions in transducer phase control experiments. We conclude this section by addressing the pre-processing steps for the analytical and observed features, setting the stage for the data-driven models discussed later.

3.1 Multi-particle Trajectories

We utilize a path planning algorithm [5] to generate feasible motion trajectories for multiple particles. We use a top-bottom 16×16 PAT levitation setup which has been adopted as the standard configuration for levitating multiple particles [8, 12, 13, 18, 20, 27, 28, 34]. While levitating 2 or 4 particles is straightforward, higher particle numbers often lead to more unsuccessful attempts. This arises from diminished acoustic energy per trap and increased particle occupation in the working volume. Additionally, due to the transducers' directivity, traps generated away from the central axis are also marginally weaker. A recent study using the same setup [13]) shows that the success rates of 3D animation with 4, 6, and 8 particles are 90%, 60%, and 40%, respectively. Considering the balance of difficulty, we focus on generating data for 6-particles in the system.

Initially, particles are randomly assigned 3D start and end positions within the working volume. We then create collision-free trajectories, maintaining minimum horizontal and vertical distances of 1.4cm and 3cm between particles. Each particle moves at a unique velocity, with a maximum of $v_{\max} = 0.1\text{ms}^{-1}$ as in [13]. The path planning algorithm provides checkpoints and constant speeds between them for each particle. With the PAT's update rate of 10 kHz, we interpolate trapping positions between waypoints, determining

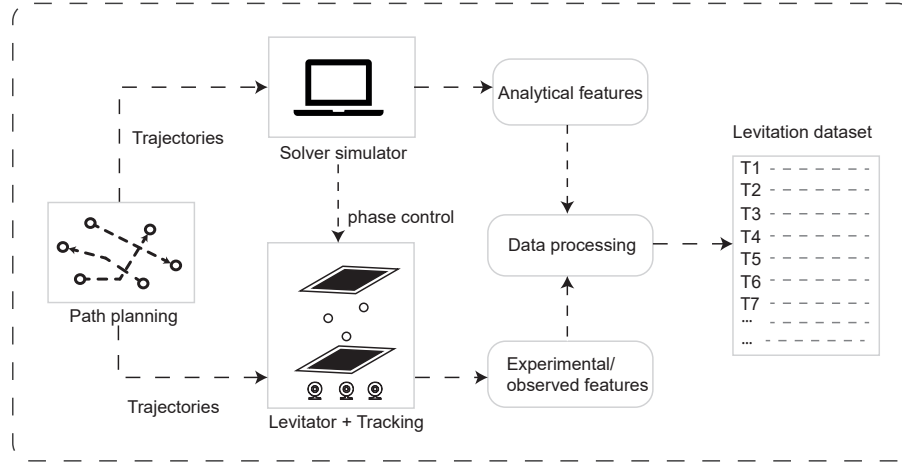


Figure 2: The building procedure of hybrid levitation dataset.

acoustic traps. This leads to a set of position data for all particles at each time step $t = [1, \dots, T]$.

3.2 Analytical Features

In this section, we acquire simulated data for multi-particle trajectories in the form of time series. First, we compute the acoustic transmission between transducers and trapping positions (see Eq. 1) and then optimize the transducers' phases ϕ^t for each time step t . To generate a varied levitation dataset, we employ different multi-point phase retrieval solvers integrated within the OpenMPD developing platform [27], such as the NAIVE and the GS-PAT [34] algorithms. Both algorithms use one pressure point per trap to estimate the phase at the focal points (or traps) and compute the phases of the transducers ϕ^t using backward propagation. Considering the attributes of different algorithms, the NAIVE algorithm does not optimize the focal points' phases but rather assumes that all points will share the same phase. For this reason, it is very common that among multiple particles, there will be some outliers of lower focal point amplitude [34]. On the other hand, GS-PAT iteratively optimizes the estimated point phases, and the generated focal amplitudes are generally higher in simulation. However, optimizing phases leads to large phase changes $\Delta\phi^t$ between successive time steps, which causes transducer amplitude fluctuations [37], and thus weaker traps and particle drops. Thus, we merge the analytical features generated by different phase retrieval solvers to reflect different issues that arise in levitation experiments.

Our analytical features mostly include concise data representations based on the complex acoustic field generated by the computed transducers' phases ϕ^t and Eq. 1. In this way, we can acquire smaller size data representations by computing the pressure p^t only at the trapping positions, or more physics-oriented data for levitation, like trapping stiffness S^t at these points. Similarly, we can compute the phase change $\Delta\theta^t$ between the computed focal points among time steps, as any large phase changes $\Delta\phi^t$ in the transducer domain will transfer to the (much fewer) focal points. Notably, to incorporate the periodicity of phase values (i.e., $\phi^t, \theta^t \in [-\pi, \pi]$) into phase

changes, we calculate the absolute phase change as shown in the equation 2.

$$|\Delta\theta^t| = \min\left(\left|\theta^t - \theta^{(t-1)}\right|, 2\pi - \left|\theta^t - \theta^{(t-1)}\right|\right) \quad (2)$$

3.3 Experimental / Observed Features

Using the motion trajectories from Section 3.1, we conduct levitation experiments to observe actual trajectories and assess motion stability. The multi-point levitation solvers continuously adjust the trajectory trap positions at a 10kHz rate using optimized transducer phases. We pair this with the OptiTrack Flex motion capture system, consisting of infrared cameras equipped with LEDs, to track real-time motion trajectories (See Figure 3). These cameras capture reflections from levitated particles to determine their 3D positions. Six cameras, placed at varied angles, cover the levitation volume. Despite the system's 120Hz tracking rate being slower than the levitator's 10kHz update, it's apt for our purpose since the particles move slowly, keeping the dataset manageable.

To minimize the impact of external factors leading to unexpected drops, before initiating the experiment, we inspect every transducer to confirm they are functioning and ensure no wind disturbance is present. During experiments, we use 2 mm-diameter EPS particles for levitation and mitigate the particle initialization displacement (to a secondary trap) by using the tracking system's feedback. We also ensure the PAT operates at a consistent performance without overheating (i.e., taking a break and allowing it to cool every 30 minutes).

We track each group of motion trajectories three consecutive times and record actual trajectories. When a displacement between the target and the actual (i.e., captured) trajectories of a particle becomes larger than 10mm, we consider that particle to have dropped. If all three attempts are completed without any particle dropping, we label the group as 'stable' (normal); otherwise, we label it as 'unstable' (abnormal). Those outcome features of motion performance will be used in later model training and evaluation. The tracking process takes approximately 40 hours.




Figure 3: Motion capture system tracks multi-particle trajectories in the mid-air.

3.4 Data-processing and Dataset Composition

To align the camera’s coordinates with the levitator’s, we applied an affine transformation to the position data. Before trajectory experiments, we established this transformation matrix by levitating a particle at 27 predefined locations (a 3D grid of $3 \times 3 \times 3$) and recording its positions using the motion capture system. This matrix then corrects the captured data to the levitator’s coordinate system, reducing position-tracking biases.

When tracking multiple particles, the motion capture system can mismatch particle positions, complicating the task of matching particle indices to position data. We address this by employing the Hungarian algorithm [21] to optimize assignments based on proximity between target and captured positions at each time step, ensuring accurate tracking trajectories. Additionally, due to optical variations and particle occlusion, some positions might be missed. These gaps are filled using interpolation to provide a comprehensive trajectory.

We merged the analytical and refined experimental data to craft our final dataset, which consists of time-series data from 200 groups spanning 902 time steps each. These groups are broken down into 90 from the NAIVE solver and 110 from the GS-PAT solver. This dataset encompasses analytical features like the Gor’kov potential (U), trapping stiffness (S), focal point amplitude (a), phase (θ), amplitude change (Δa), point phase change ($\Delta \theta$), and average transducer phase change ($\Delta \phi$). The outcome labels are detailed in Table 1. Notably, while we used the same motion trajectories for the two solvers, the resulting features and instabilities varied. The phase instabilities of GS-PAT are more frequent and can affect all acoustic traps, leading to more abnormal groups. In contrast, the NAIVE algorithm shows higher success rates as its amplitude discrepancies usually concern individual traps. Our hybrid dataset is available online ¹.

4 STABLELEV

While research has noted instabilities and drops in multi-point levitation[8, 13], there’s no documented study predicting such behavior during dynamic levitation of multiple particles. Furthermore,

¹<https://github.com/Lei-Oriana/StableLev>

	Normal	Abnormal
NAIVE	75	15
GS-PAT	32	78

Table 1: Observed outcome labels of running 200 group levitation trajectories by NAIVE and GS-PAT solvers.

strategies for improving stability in dynamic settings are uncharted in the current literature. Addressing this gap, we introduce StableLev, a data-driven solution for optimizing multi-particle stability. This method unfolds in three phases:

From our dataset and expertise, we pinpoint essential levitation characteristics (*stage: feature curation*). Using diverse deep neural network models, we spot anomalies in unstable trajectories (*stage: anomaly detection*). We rectify detected anomalies, bolstering motion steadiness (*stage: anomaly amendment*).

4.1 Feature Curation

In our analysis, we utilized a feature correlation heatmap (Figure 4) to determine the relationships between each analytical feature in our dataset (section 3.4). The heatmap color scale indicates the strength of feature correlations. Notably, Gor’kov potential (U), stiffness (S), and focal point amplitude (a) emerge as tightly interrelated, all indicating trap intensity. Among these, we prioritize focal point amplitude (a) due to its computational efficiency, needing just a singular pressure value per trap, unlike the Gor’kov-associated features that demand more complex computations[25].




Figure 4: Correlation heatmap of potential analytical features from the transducers and traps, including Gor’kov potential (U), trapping stiffness (S), focal point amplitude (a), phase (θ), amplitude change (Δa), point phase change ($\Delta \theta$) and average transducer phase change ($\Delta \phi$).

Additionally, the heatmap reveals phase change ($\Delta \theta$) as another critical factor impacting trap intensity. Rapid phase transitions during motion adjustments might induce transducer emission fluctuations [37], an aspect not reflected in intensity-related analytical features. This realization underscores the significance of the point phase change ($\Delta \theta$) as an additional feature that can provide valuable information for anomaly detection.

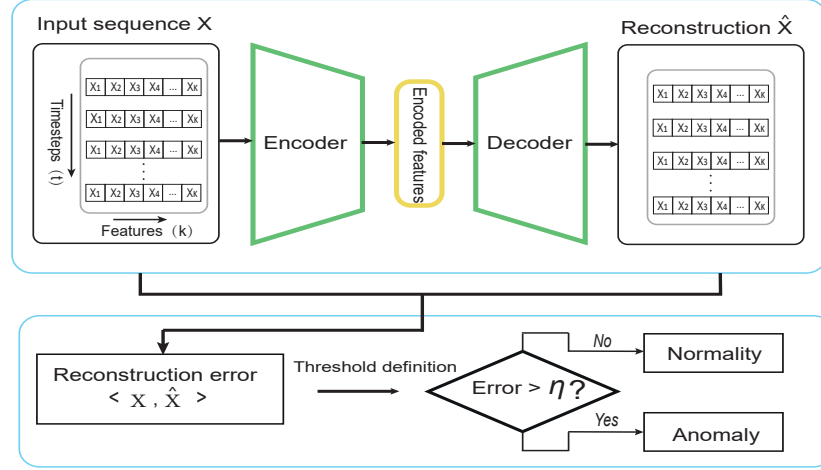


Figure 5: Anomaly detection model overview using AutoEncoder-based deep neural networks.

4.2 Anomaly Detection

Having identified the focal point amplitude (a) and point phase change ($\Delta\theta$) as pivotal features, our next aim is to detect instances of particle drops during the dynamic levitation process. We characterize these drops as anomalies in a time-series dataset, posing unique challenges due to their unbalanced occurrences and unpredictable behaviors [9, 32].

Deep learning has shown remarkable capabilities in learning underlying features to detect anomalies [32, 41]. AutoEncoders (AE) emerge as a promising solution to this anomaly detection challenge. AEs excel at uncovering non-linear correlations in datasets, crucial for recognizing subtle, unpredicted deviations. Their encoder-decoder mechanism efficiently reduces data dimensionality, emphasizing crucial features while filtering out noise. This ensures accurate anomaly identification even within intricate datasets. [6, 35].

To further bolster anomaly detection capabilities, we integrate deep AEs with the temporal dynamics of long-short-term memory (LSTM) and gated recurrent unit (GRU) architectures and the robustness of deep generative models such as variational AutoEncoders (VAE):

LSTM: A specialized form of RNN, LSTMs adeptly handle long-term dependencies in sequential data. Equipped with memory cells and gate units, they can filter noise and retain significant patterns, rendering them particularly effective for our anomaly detection challenge [15].

GRU: A variant of RNNs, GRUs address gradient issues inherent in traditional RNNs. Their memory cells and gated units, including the update and reset gates, make them adept at processing complex time-series data and detecting anomalies [14, 36].

VAE: Variational AutoEncoders blend probabilistic generative models with deep neural network capacities. Their encoders output conditional probability distributions, thus allowing superior data reconstruction, making them valuable for modeling standard behaviors in anomaly detection [23, 43].

Using these building blocks, we propose three hybrid anomaly detector models, namely LSTM AE, GRU AE, and LSTM VAE, all

built on the AE framework. Figure 5 shows our hybrid anomaly detector designs, where we can represent the encoder and decoder units with different components, such as LSTM or GRU layers. For LSTM VAE, we represent the encoder unit by LSTM layers followed by the distribution function (mean and variance) to learn the encoder features and decoder unit by LSTM layers, respectively. A detailed exploration of each unit’s functionality can be found in Homayouni et al [19]. Table 2 presents the specifications and hyper-parameters of hybrid AE models.

Utilizing our time-series dataset, our models train on the selected features of each levitated particle. Through the AE’s encoder-decoder framework, our approach discerns patterns in stable trajectories and pinpoints anomalies in unstable ones. The results showcasing the efficacy of our anomaly detection approach are detailed in Figure 6, 7 and Section 5.1.

4.3 Anomaly Amendment

Following our anomaly detection process, we take corrective measures in the anomaly regions to rectify potential instabilities. Of the two prominent analytical features critical for detecting anomalies in real levitation trajectories, namely the focal point amplitude and phase change, we prioritize rectifying anomalies linked to amplitude (see results in Section 5.2).

Though our AE models predict absolute phase change values ($\Delta\theta$), and given that our dataset also includes the focal point phase (θ), we could technically estimate phase information for amending anomalies. However, this estimation process introduces complexities. Namely, it necessitates additional constraints to ensure that the AE-predicted phase changes remain minimal across consecutive time frames. Presently, significant phase changes are managed by interpolating over the entire phase change range across various time frames. This interpolation, however, hampers the speed of focal point generation and transitions [37].

In practice, multi-point levitation solvers determine transducer phases based on a consistent target point amplitude for each trap, aiming for uniform intensity across traps. However, our anomaly

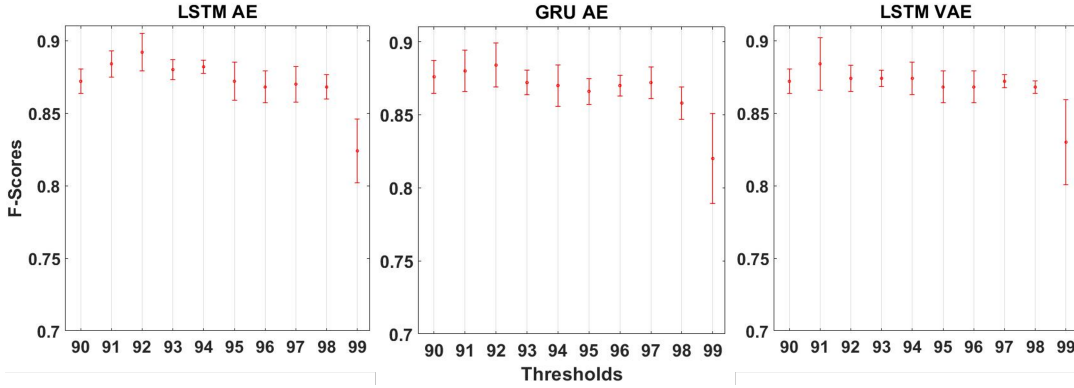


Figure 6: Performance F-score error plots for each k-fold hybrid AE model at different reconstruction error threshold values.

detection and observation results indicate that this ideal criterion often remains unachieved in real experiments, leading to unintended particle drop. A direct solution to this is adjusting the amplitude of specific, unstable particles within identified anomaly regions. Given the anomaly time windows, we can trace back to the corresponding positions along the trajectory, where we create target trap points. At these points, we deliberately increase the target amplitude for unstable particles while maintaining the amplitude for stable ones.

5 RESULTS

5.1 Hybrid AE Models Training and Anomaly Detection

To train our hybrid models, we first represent the time series data set of the selected features into a sequence of time windows (T_w), where the selected features of all traps at time step t of each time window can be depicted as $X_{1,t}, X_{2,t}, \dots, X_{k,t}$, where, $k = 1, \dots$, total features of all traps.

We explored various "look-back periods" or time window sizes (10, 50, 100, 200, and 500) to train our AE models. Given the abrupt variations in our feature dataset's short time steps and the 200 groups of sequences, each with 902 time steps, larger windows proved inefficient and unwieldy. Hence, we settled on a time window size ($(T_w) = 20$) to train our model.

We split the dataset from Section 3.4 into training (77 normal samples: 55 NAIVE, 22 GSPAT) and test sets, which contain both normal (20 NAIVE, 10 GSPAT) and abnormal samples (15 NAIVE, 78 GSPAT). Given the differing scales of the selected features, we apply the linear Min-Max scaling method, transforming the feature values as,

$$X = \frac{X_{\text{original}} - \min_{\text{normal}}}{\max_{\text{normal}} - \min_{\text{normal}}} \quad (3)$$

Here, X ranges between $[0-1]$, with \min_{normal} and \max_{normal} denoting the training dataset's minimum and maximum values. Utilizing 5-fold cross-validation [40], we train hybrid models on sequential data X , encoding with LSTM, GRU, or VAE, then reconstructing to \hat{X} . The aim is minimizing the mean squared reconstruction error ($R_L = \text{MSE}(X - \hat{X})$).

Time windows or look-back period T_w	20
K-fold	5
Number of layers for encoder/decoder	2
Number of memory units or neurons in each layer	64(layer1) 32(layer2)
Activation function	tanh
Reconstruction error threshold η	92%

Table 2: Hyperparameters of hybrid AutoEncoder models.

Utilizing trained models, we predicted on the test dataset, leveraging the reconstruction error threshold η to distinguish sequence types. To select the proper threshold, we determine it with the F-score as the metric of choice[41]. For our detection purpose, we prioritize Precision (finding as many anomalies as possible) and Recall (minimizing the missed anomalies) as our evaluation metrics, since F-Score usually makes a balance between Precision and Recall, also useful when the distribution of class is imbalanced. So we adopt the F-score as our metric to select the threshold to distinguish between normal and abnormal. In K-fold cross-validation, we assessed variability and uncertainty in predictions of each fold, displaying performance as error bars on F-score values. Our analyses, shown in Figure 6, reveal that a 90% - 99% threshold range saw consistent performance, with mean F-scores approximately between 0.80 and 0.90.

We opt for the LSTM AE hybrid model to enhance stability and improve dynamic levitation. With a 92% threshold, the model achieves a mean F-score of 0.9 (with corresponding Precision: 86%, Recall:95%), demonstrating roughly 90% accuracy in detecting anomalies on the test dataset. In our test set, most of the actual abnormal groups (88 groups) are correctly predicted as abnormal (true positive), while a few (5 groups) fall into a false negative category. Some actual normal groups (14 groups) are predicted as abnormal (false positive), and 16 actual normal groups are predicted as normal (true negative). Among the true positive groups, we present a few examples where feature anomalies precede actual particle drop events (e.g., large position displacement captured by the camera) in Figure 7. Notably, a single anomaly step does not necessarily lead to one

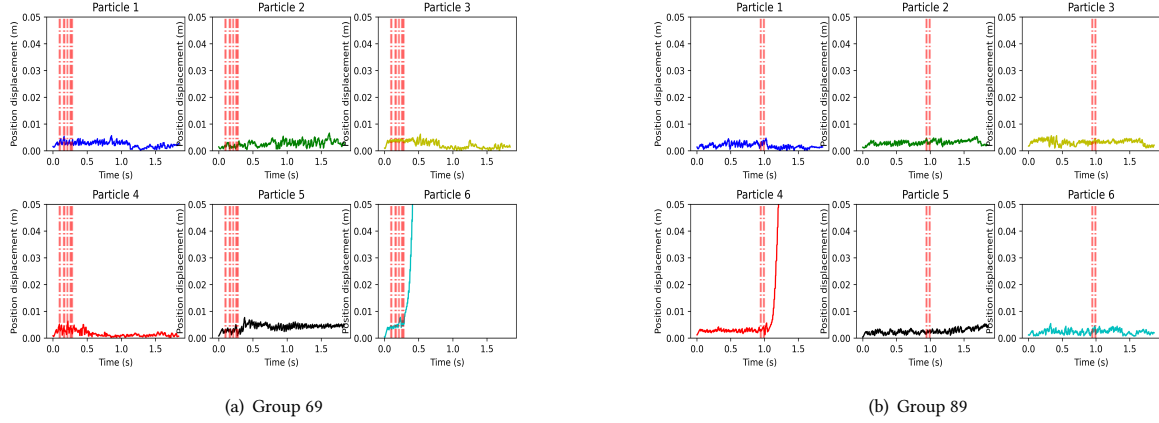


Figure 7: The position displacements between target trajectories and real trajectories, with the red dashed lines indicating the predicted anomaly time windows. Particle 6 dropped in group 69 and particle 4 in group 89.

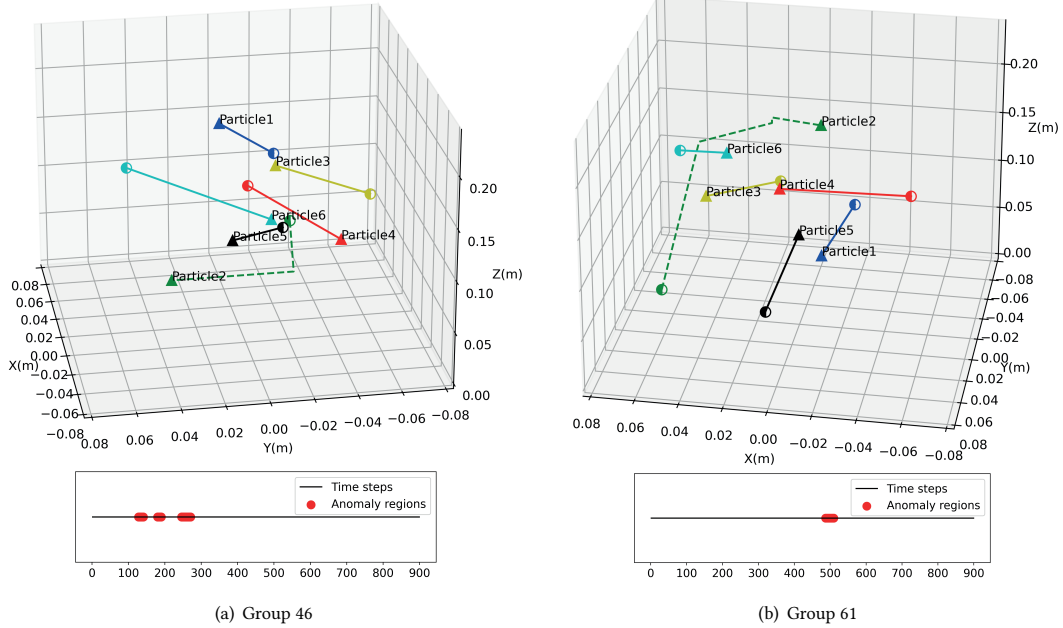


Figure 8: Two abnormal groups succeeded after adjustment to the target point amplitude of dropping particles. In group 46 (a) and group 61 (b), only particle 2 dropped when moving along the target trajectories (dashed line). Anomaly time windows (in red) are predicted by the anomaly detection model.

(or more) particle drops and we often observe an accumulation of anomalies before drop events, as indicated by the red dashed lines in Figure 7.

5.2 Stability Enhancement

Here, we present a few anomalous groups (46, 61, 62, and 67 in Figure 8, 9) as examples and report the stability enhancement through the amplitude amendment approach proposed in Section 4.3.

First, we used the processed tracking trajectories in our dataset and compared them to the target trajectories, identifying which particles dropped (i.e., particle 2 in Figure 8). Also with the anomaly time regions, when setting the phase retrieval solver parameters, we change the unstable trap's target amplitude to higher than other previous stable particles (traps). Note that we gradually increase the target amplitude and find a proper increase. Here, with 30% of amplitude enhancement at predicted anomalous regions, we

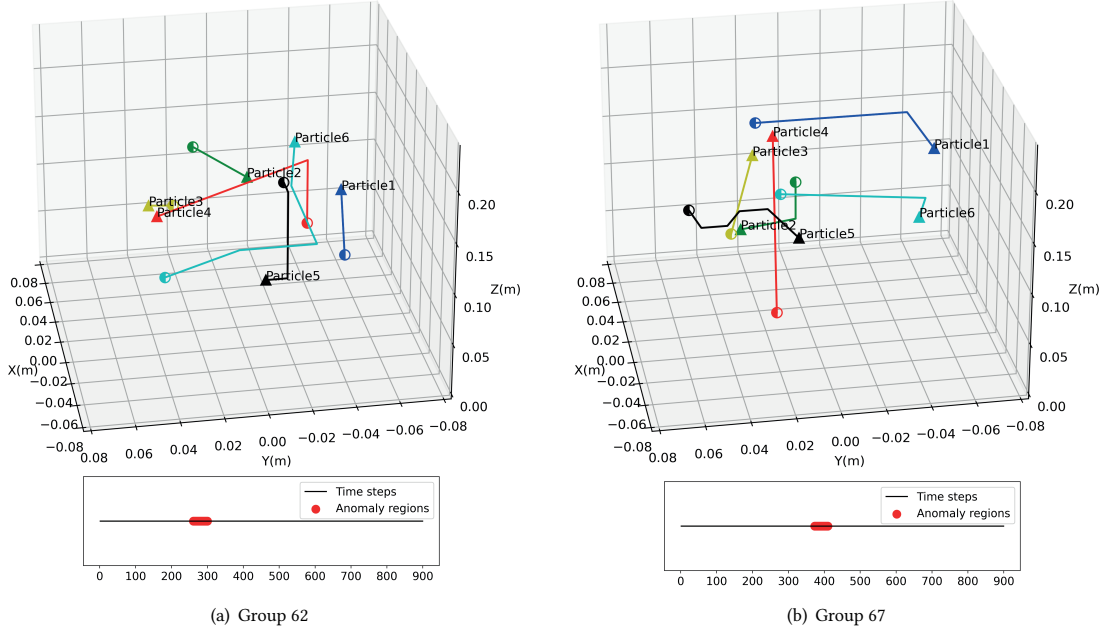


Figure 9: Two abnormal groups succeeded after adjusting the target point amplitude, including particles beyond those exclusively categorized as 'dropping. (a) In group 62, all trajectories get stable after lowering the target amplitude of particle 4. (b) In group 67, all trajectories get stable after lowering the target amplitude of particle 1 and increasing the target amplitude of particle 2.

repeatedly ran the same trajectory for consecutive 3 times as we did in Section 3, and particle 2 in both groups 46 and 61 arrived at the endpoint of the trajectory without dropping.

In our examination of amplitude modifications for groups 62 and 67, we noted that exclusively increasing the target amplitude of the unstable trap does not reliably enhance stability. Across three repeated tests of a few groups, we found it is uncertain whether a particular particle will consistently drop, complicating the identification of the problematic trap. However, during this period of anomaly, we have the option to either reduce the amplitude of the stronger trap, increase the amplitude of the relatively weaker one, or employ a combination of both strategies to comprehensively address this instability. In group 62 (see Figure 9), after lowering the trap amplitude of particle 4 by 20%, no drop occurred. Likewise, by lowering the trap amplitude of particle 1 and increasing the trap amplitude of particle 2 by 20% within the suggested anomaly time region, we prevented the drop from happening in group 67.

6 DISCUSSION

6.1 Improvement of Stability Detection and Enhancement

Harnessing our domain knowledge and utilizing the feature correlation heat map, we have identified two crucial features that are instrumental in achieving trap stability during the dynamic levitation of multi-particles. Our current model, based on these features, attained an F-score of 0.9, but incorporating additional correlated intensity features like Gor'kov potential (U) and stiffness (S) could further enhance learning and improve model performance. For

instance, the selection of Gor'kov potential (U) with focal point amplitude (a) in combination with point phase change ($\Delta\theta$) can unveil previously hidden patterns that can greatly enhance the model's efficacy. However, as features like stiffness are more computationally demanding, the trade-off between the increase in feature dimensions and the associated training costs should be taken into account.

One limitation of our study is its emphasis on anomaly amendment using only one primary feature: the focal point amplitude, which is straightforward and feasible for existing phase retrieval algorithms to achieve in real time. Beyond intuitively tuning amplitude, we envision that the reconstructed time-series sequences by our AE models are inherent indicators of stable amplitude and phase change. By further utilizing that information, it is possible to tune both amplitude and phase changes to get robust levitation performance. Our dataset already encompasses time-series data on varied features. Therefore, the holistic approach to repairing anomalies could encompass amplitude, phase, and others, marking a potential avenue for future extensions of our work.

6.2 Generalizability of Dataset and Finding

We observed that various phase retrieval solvers influence stability differently due to their inherent properties. To reflect this, we selected two representative solvers that showcase a range of levitation situations, thereby improving the generalizability of our dataset. Meanwhile, our data is exclusively derived from the setup with 16×16 top-bottom PATs, which are commonly used in levitation applications and have shown some performance issues. In the future, to extend the dataset in different setups (e.g., 8×8 top-bottom

PATs, V-shaped PATs), researchers can start by establishing the levitator through the OpenMP framework [27] which is becoming a standard hardware and solver solution in this field (e.g., adopted in UIST student innovation contest). Once set up, researchers can replicate the dataset composition process and train their models with proper analytical and experimental features.

6.3 Levitation Stability on Various Object Structures

In this work, we mainly discussed levitation stability on individual EPS particles, and those observing and experimental features are based on particle performance. Apart from point objects, other object structures like threads [8] or fabrics [28] are often employed for levitation. To consider the stability of such levitation structures, compared to the individual particles, collective effects (e.g., a set of traps maintains a fixed relative distance and shares the same velocity) should be considered when analyzing the features and making amendments. Single unstable traps can inevitably affect the neighbor traps, especially when they are physically connected by threads or fabrics. Therefore, more types of analytical features and experimental constraints will be introduced for multi-"object" levitation stability.

6.4 Levitation Stability on Various Object Materials

Note that levitation stability also varies depending on different materials. As previous comparisons [13] suggest, liquid particles achieve less stable motion compared to EPS solid particles under the same number of traps and moving velocity. Properties such as density and surface tension of liquid particles [38] play a crucial role in how particles respond to acoustic pressure and acoustic radiation force. These forces can cause deformations, changes in shape, or even particle destruction, which are all dynamic aspects of particle behavior during levitation. Therefore, when characterizing the levitation stability of different object materials, it is necessary to consider more underlying physics properties and different motion behaviors, going beyond the simple binary distinction of whether they drop or remain suspended.

6.5 Further Data-driven Explorations

This paper presents StableLev, the first data-driven approach tailored for dynamic multi-point acoustic levitation. It's crucial to recognize that prior deep learning endeavors (e.g., [22]) targeting acoustic phase-modulating devices have primarily hinged on simulated data and focused on creating complex acoustic fields, often visualized as images. However, such acoustic wavefront shaping strategies are predominantly suited for static particle manipulations, a prime example being a one-step acoustic fabrication, as documented in [26]. In stark contrast, StableLev stands out with its capability to process time series data to successfully bolster the stability of levitating particles in motion during experiments.

StableLev represents an initial foray into leveraging data for pinpointing and controlling dynamic acoustic fields. Acoustic levitation inherently grapples with non-linear acoustic phenomena, which the Gor'kov theory has sought to streamline by assuming

enduring acoustic waves, essentially static in nature. Our open-sourced hybrid levitation dataset and methodology can inspire novel research that can, for example, seek explicit equations for the non-linear dynamics of multi-particle levitation, similar to ongoing research on data-driven discovery of governing physics [3] or hardware-in-the-loop modeling of interactive devices [33].

7 CONCLUSION

Acoustic levitation, a burgeoning domain, promises groundbreaking advancements in human-computer interaction. Despite its transformative potential, challenges in stable multi-point dynamic levitation persist. This paper addresses these challenges by introducing "StableLev," a data-driven methodology that utilizes a hybrid levitation dataset that we created as a blend of simulations and real-world data. With an emphasis on anomaly detection through AutoEncoders, we effectively pinpoint and rectify unstable levitation trajectories. Achieving an F -score of 0.9, StableLev paves the way for more reliable acoustic levitation, fostering robustness in contemporary systems and serving as a foundational pillar for future innovations.

ACKNOWLEDGMENTS

This work was supported by the EU-H2020 through their ERC Advanced Grant (No. 787413), by the Royal Academy of Engineering through their Chairs in Emerging Technology Program (CIET18/19), by the EPSRC through their Prosperity Partnership Program (EP/V037846/1), and by the UKRI Frontier Research Guarantee Grant (EP/X019519/1).

REFERENCES

- [1] Marco AB Andrade, Nicolás Pérez, and Julio C Adamowski. 2018. Review of progress in acoustic levitation. *Brazilian Journal of Physics* 48 (2018), 190–213. <https://doi.org/10.1007/s13538-017-0552-6>
- [2] Myroslav Bachynskyi, Viktorija Paneva, and Jörg Müller. 2018. LeviCursor: Dexterous interaction with a levitating object. In *Proceedings of the 2018 ACM International Conference on Interactive Surfaces and Spaces*. Association for Computing Machinery, New York, NY, USA, 253–262. <https://doi.org/10.1145/3279778.3279802>
- [3] Steven Brunton, Joshua Proctor, and Nathan Kutz. 2016. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences of the United States of America* 113, 15 (2016). <https://doi.org/10.1073/pnas.1517384113>
- [4] Henrik Bruus. 2012. Acoustofluidics 7: The acoustic radiation force on small particles. *Lab on a Chip* 12, 6 (2012), 1014–1021. <https://doi.org/10.1039/C2LC21068A>
- [5] Jingkai Chen, Jiaoyang Li, Chuchu Fan, and Brian C Williams. 2021. Scalable and safe multi-agent motion planning with nonlinear dynamics and bounded disturbances. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Association for the Advancement of Artificial Intelligence, 1101 Pennsylvania Ave, NW, Suite 300, Washington, DC 20004 USA, 11237–11245. <https://doi.org/10.1609/aaai.v35i13.17340>
- [6] Zhaomin Chen, Chai Kiat Yeo, Bu Sung Lee, and Chiew Tong Lau. 2018. Autoencoder-based network anomaly detection. In *2018 Wireless telecommunications symposium (WTS)*. IEEE, Phoenix, AZ, USA, 1–5. <https://doi.org/10.1109/WTS.2018.8363930>
- [7] Iñigo Ezcurdia, Rafael Morales, Marco AB Andrade, and Asier Marzo. 2022. LeviPrint: Contactless Fabrication using Full Acoustic Trapping of Elongated Parts. In *ACM SIGGRAPH 2022 Conference Proceedings*. Association for Computing Machinery, New York, NY, USA, 1–9. <https://doi.org/10.1145/3528233.3530752>
- [8] Andreas Rene Fender, Diego Martinez Plasencia, and Sriram Subramanian. 2021. ArticLev: An integrated self-assembly pipeline for articulated multi-bead levitation primitives. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3411764.3445342>
- [9] Tharindu Fernando, Harshala Gammulle, Simon Denman, Sridha Sridharan, and Clinton Fookes. 2021. Deep learning for medical anomaly detection—a survey. *ACM Computing Surveys (CSUR)* 54, 7 (2021), 1–37. <https://doi.org/10.1145/3464423>

- [10] D Foresti, M Nabavi, and D Poulikakos. 2012. On the acoustic levitation stability behaviour of spherical and ellipsoidal particles. *Journal of Fluid Mechanics* 709 (2012), 581–592. <https://doi.org/10.1017/jfm.2012.350>
- [11] Tatsuki Fushimi, TL Hill, Asier Marzo, and BW Drinkwater. 2018. Nonlinear trapping stiffness of mid-air single-axis acoustic levitators. *Applied Physics Letters* 113, 3 (2018), 034102. <https://doi.org/10.1063/1.5034116>
- [12] Tatsuki Fushimi, Asier Marzo, Bruce W Drinkwater, and Thomas L Hill. 2019. Acoustophoretic volumetric displays using a fast-moving levitated particle. *Applied Physics Letters* 115, 6 (2019), 064101. <https://doi.org/10.1063/1.5113467>
- [13] Lei Gao, Pourang Irani, Sriram Subramanian, Gowdham Prabhakar, Diego Martinez Plasencia, and Ryuji Hirayama. 2023. DataLev: Mid-air Data Physicalisation Using Acoustic Levitation. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3544548.3581016>
- [14] Xundong Gong, Shibo Liao, Fei Hu, Xiaoqing Hu, and Chunshan Liu. 2022. Autoencoder-Based Anomaly Detection for Time Series Data in Complex Systems. In *2022 IEEE Asia Pacific Conference on Circuits and Systems (APCCAS)*. IEEE, Shenzhen, China, 428–433. <https://doi.org/10.1109/APCCAS55924.2022.10090260>
- [15] Edan Habler and Asaf Shabtai. 2018. Using LSTM encoder-decoder algorithm for detecting anomalous ADS-B messages. *Computers & Security* 78 (2018), 155–173. <https://doi.org/10.1016/j.cose.2018.07.004>
- [16] Koji Hasegawa and Manami Murata. 2022. Oscillation Dynamics of Multiple Water Droplets Levitated in an Acoustic Field. *Micromachines* 13, 9 (2022), 1373. <https://doi.org/10.3390/mi13091373>
- [17] Ryuji Hirayama, Giorgos Christopoulos, Diego Martinez Plasencia, and Sriram Subramanian. 2022. High-speed acoustic holography with arbitrary scattering objects. *Science advances* 8, 24 (2022), eabn7614. <https://doi.org/10.1126/sciadv.abn7614>
- [18] Ryuji Hirayama, Diego Martinez Plasencia, Nobuyuki Masuda, and Sriram Subramanian. 2019. A volumetric display for visual, tactile and audio presentation using acoustic trapping. *Nature* 575, 7782 (2019), 320–323. <https://doi.org/10.1038/s41586-019-1739-5>
- [19] Hajar Homayouni, Sudipto Ghosh, Indrakshi Ray, Shlok Gondalia, Jerry Duggan, and Michael G Kahn. 2020. An autocorrelation-based LSTM-Autoencoder for anomaly detection on time-series data. In *2020 IEEE international conference on big data (big data)*. IEEE, Atlanta, GA, USA, 5068–5077. <https://doi.org/10.1109/BigData50022.2020.9378192>
- [20] Eimontas Jankauskis, Sonia Elizondo, Roberto Montano Murillo, Asier Marzo, and Diego Martinez Plasencia. 2022. TipTrap: A Co-located Direct Manipulation Technique for Acoustically Levitated Content. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3526113.3545675>
- [21] Harold W Kuhn. 1955. The Hungarian method for the assignment problem. *Naval research logistics quarterly* 2, 1-2 (1955), 83–97. <https://doi.org/10.1002/nav.3800020109>
- [22] Qin Lin, Jiaqian Wang, Feiyan Cai, Rujun Zhang, Degang Zhao, Xiangxiang Xia, Jinping Wang, and Hairong Zheng. 2021. A deep learning approach for the fast generation of acoustic holograms. *The Journal of the Acoustical Society of America* 149, 4 (2021), 2312–2322. <https://doi.org/10.1121/10.0003959>
- [23] Shuyu Lin, Ronald Clark, Robert Birke, Sandro Schönborn, Niki Trigoni, and Stephen Roberts. 2020. Anomaly detection for time series using vae-lstm hybrid model. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, Barcelona, Spain, 4322–4326. <https://doi.org/10.1109/ICASSP40776.2020.9053558>
- [24] Asier Marzo and Bruce W Drinkwater. 2019. Holographic acoustic tweezers. *Proceedings of the National Academy of Sciences* 116, 1 (2019), 84–89. <https://doi.org/10.1073/pnas.1813047115>
- [25] Asier Marzo, Sue Ann Seah, Bruce W Drinkwater, Deepak Ranjan Sahoo, Benjamin Long, and Sriram Subramanian. 2015. Holographic acoustic elements for manipulation of levitated objects. *Nature communications* 6, 1 (2015), 8661. [https://doi.org/10.1038/ncomms9661\(2015\)](https://doi.org/10.1038/ncomms9661(2015))
- [26] Kai Melde, Heiner Kremer, Minghui Shi, Senne Seneca, Christopher Frey, Ilia Platzman, Christian Degel, Daniel Schmitt, Bernhard Schölkopf, and Peer Fischer. 2023. Compact holographic sound fields enable rapid one-step assembly of matter in 3D. *Science Advances* 9, 6 (2023). <https://doi.org/10.1126/sciadv.adf6182>
- [27] Roberto Montano-Murillo, Ryuji Hirayama, and Diego Martinez Plasencia. 2023. OpenMPD: A low-level presentation engine for Multimodal Particle-based Displays. *ACM Transactions on Graphics* 42, 2 (2023), 1–13. <https://doi.org/10.1145/3572896>
- [28] Rafael Morales, Asier Marzo, Sriram Subramanian, and Diego Martínez. 2019. LeviProps: Animating levitated optimized fabric structures using holographic acoustic tweezers. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 651–661. <https://doi.org/10.1145/3332165.3347882>
- [29] Yoichi Ochiai, Takayuki Hoshi, and Jun Kekimoto. 2014. Three-dimensional mid-air acoustic manipulation by ultrasonic phased arrays. *PLoS one* 9, 5 (2014), e97590. <https://doi.org/10.1371/journal.pone.0102525>
- [30] Themis Omirou, Asier Marzo Perez, Sriram Subramanian, and Anne Roudaut. 2016. Floating charts: Data plotting using free-floating acoustically levitated representations. In *2016 IEEE Symposium on 3D User Interfaces (3DUI)*. IEEE, Greenville, SC, USA, 187–190. <https://doi.org/10.1109/3DUI.2016.7460051>
- [31] Viktorija Paneva, Arthur Fleig, Diego Martinez Plasencia, Timm Faulwasser, and Jörg Müller. 2022. OptiTrap: Optimal trap trajectories for acoustic levitation displays. *ACM Transactions on Graphics* 41, 5 (2022), 1–14. <https://doi.org/10.1145/3517746>
- [32] Guansong Pang, Chunhua Shen, Longbing Cao, and Anton Van Den Hengel. 2021. Deep learning for anomaly detection: A review. *ACM computing surveys (CSUR)* 54, 2 (2021), 1–38. <https://doi.org/10.1145/3439950>
- [33] Yifan Peng, Sueyon Choi, Nitish Padmanaban, and Gordon Wetzstein. 2020. Neural holography with camera-in-the-loop training. *ACM Transactions on Graphics* 39, 6 (2020), 1–14. <https://doi.org/10.1145/3414685.3417802>
- [34] Diego Martinez Plasencia, Ryuji Hirayama, Roberto Montano-Murillo, and Sriram Subramanian. 2020. GS-PAT: high-speed multi-point sound-fields for phased arrays of transducers. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 138–1. <https://doi.org/10.1145/3386569.3392492>
- [35] Jamal Saeedi and Alessandro Giusti. 2023. Semi-supervised visual anomaly detection based on convolutional autoencoder and transfer learning. *Machine Learning with Applications* 11 (2023), 100451. <https://doi.org/10.1016/j.mlwa.2023.100451>
- [36] Changcheng Sun, Zhiwei He, Huipin Lin, Linhui Cai, Hui Cai, and Mingyu Gao. 2023. Anomaly detection of power battery pack using gated recurrent units based variational autoencoder. *Applied Soft Computing* 132 (2023), 109903. <https://doi.org/10.1016/j.asoc.2022.109903>
- [37] Shun Suzuki, Masahiro Fujiwara, Yasutoshi Makino, and Hiroyuki Shinoda. 2020. Reducing amplitude fluctuation by gradual phase shift in midair ultrasound haptics. *IEEE transactions on haptics* 13, 1 (2020), 87–93. <https://doi.org/10.1109/TOH.2020.2965946>
- [38] Chi Thanh Vi, Asier Marzo, Damien Ablart, Gianluca Memoli, Sriram Subramanian, Bruce Drinkwater, and Marianna Obrist. 2017. Tastyfloats: A contactless food delivery system. In *Proceedings of the 2017 ACM International Conference on Interactive Surfaces and Spaces*. Association for Computing Machinery, New York, NY, USA, 161–170. <https://doi.org/10.1145/3132272.3134123>
- [39] Chi Thanh Vi, Asier Marzo, Gianluca Memoli, Emanuela Maggioni, Damien Ablart, Martin Yeomans, and Marianna Obrist. 2020. LeviSense: A platform for the multisensory integration in levitating food and insights into its effect on flavour perception. *International Journal of Human-Computer Studies* 139 (2020), 102428. <https://doi.org/10.1016/j.ijhcs.2020.102428>
- [40] Hoang Lan Vu, Kelvin Tsun Wai Ng, Amy Richter, and Chunjiang An. 2022. Analysis of input set characteristics and variances on k-fold cross validation for a Recurrent Neural Network model on waste disposal rate estimation. *Journal of environmental management* 311 (2022), 114869. <https://doi.org/10.1016/j.jenvman.2022.114869>
- [41] Yuanyuan Wei, Julian Jang-Jaccard, Fariza Sabrina, Wen Xu, Seyit Camtepe, and Aeryn Dunmore. 2023. Reconstruction-based LSTM-Autoencoder for Anomaly-based DDoS Attack Detection over Multivariate Time-Series Data. arXiv:2305.09475 [cs.CR]
- [42] RR Whymark. 1975. Acoustic field positioning for containerless processing. *Ultrasonics* 13, 6 (1975), 251–261. [https://doi.org/10.1016/0041-624X\(75\)90072-4](https://doi.org/10.1016/0041-624X(75)90072-4)
- [43] Chunkai Zhang, Shaocong Li, Hongye Zhang, and Yingyang Chen. 2020. VELC: A New Variational AutoEncoder Based Model for Time Series Anomaly Detection. arXiv:1907.01702 [cs.LG]