


Burial Depth Estimation for Partially Embedded Rocks Using Scanning Laser Doppler Vibrometry and Neural Networks

Yiyan Ruan and Erik Komendera

Abstract— Estimating the burial depth of partially embedded rocks is a key capability for energy-limited robotic excavation, where unnecessary digging or repeated probing can be prohibitively costly. This paper presents a vibration-based estimation pipeline that uses impulse excitation and scanning laser Doppler vibrometry to characterize the soil-coupled dynamic response of a partially embedded rock, then maps that response to burial depth using a neural network. Physics-based spring–mass models from the geotechnical literature establish that a single frequency measurement recovers only the burial-depth-to-mass ratio, h/m . These models do not fully account for the nonlinear, strain-dependent soil behavior and coupled vibration modes present in practice, motivating a data-driven approach that extracts richer features from multi-point scanned data. Eleven measurable features, including resonance parameters, half-power damping ratio, and a spatial vibration decay metric, are mapped to burial depth by a Multilayer Perceptron trained with strict Leave-One-Rock-Out cross-validation, achieving, to the authors' knowledge, the first vibration-based burial-depth estimation without prior knowledge of rock mass, with 25.3% mean absolute percentage error (MAPE) and 1.03 cm root mean squared error (RMSE) on 16 concrete blocks spanning 180 measurements.



Index Terms— Frequency response, geophysical measurements, laser Doppler vibrometry, neural networks, robot sensing systems, vibration measurement.

I. INTRODUCTION

FUTURE lunar surface missions under the Artemis program will require autonomous robots to prepare terrain, remove obstacles, and support long-duration infrastructure while operating in unstructured regolith with limited human supervision [1]–[3]. Rocks in these environments are both navigation obstacles and excavation hazards. Before attempting removal, a robot must estimate whether a rock is within the vehicle's handling capacity, which depends not only on its visible size but also on its total mass and the degree to which it is anchored by the surrounding regolith.

This work complements our prior research published in IEEE SENSORS 2025, which addressed the visible portion of this problem by reconstructing partially buried rocks from RGB-D images [4] and estimating exposed geometry, volume, and center of gravity. However, camera imagery and geometry reconstruction cannot observe the buried portion of a rock. As a result, they cannot determine whether a visually similar rock is lightly covered and easy to remove, or deeply embedded and mechanically constrained by the soil.

Burial depth is therefore a critical missing variable for

robotic excavation planning. Directly measuring it would require digging around the rock, defeating the purpose of pre-excitation assessment. Vibration response offers a non-excavating alternative: a partially embedded rock behaves as a rigid body coupled to the surrounding soil, and its resonance frequency, mobility, damping, and spatial mode shape are governed by the soil–rock contact condition. Because this contact condition changes with burial depth, the measured vibration response can provide information about the hidden embedded portion of the rock.

Building on this idea, this paper proposes a data-driven approach that foregoes analytical inversion. Rather than inverting an analytical model, which requires soil stiffness and rock mass values that are unavailable in the field, we treat the multi-point vibration response as a rich feature space and learn the mapping to burial depth from labeled examples. Physical effects such as strain-dependent stiffness, coupled vibration modes, and depth-dependent damping produce correlated and repeatable signatures in the frequency response that a data-driven model can exploit even when no single effect is analytically invertible. The central question addressed in this work is whether paired multi-point vibration measurements, combined with exposed-geometry inputs, can recover absolute burial depth for held-out rocks without direct knowledge of rock mass or soil stiffness.

The main contributions of this work are:

Y. Ruan and E. Komendera are with the Department of Mechanical Engineering, Virginia Polytechnic Institute and State University, Blacksburg, VA, USA (e-mail: yiyanr97@vt.edu; komendera@vt.edu).

Code and data: <https://github.com/JXNICHOLAS/rock-vib-burial-depth>.

- A vibration-based framework for estimating the burial depth of partially embedded rocks, motivated by lunar robotic excavation and site-preparation needs.
- A feature extraction methodology from multi-point frequency response data, including half-power damping ratio and a spatial vibration decay metric from scanning vibrometry.
- An eleven-input neural network regressor evaluated with strict Leave-One-Rock-Out (LORO) cross-validation to test generalization to unseen rocks.
- A controlled vibration dataset collected from 16 cut concrete blocks with known geometry, density, and ground-truth burial depth, spanning 180 individual measurements.

II. RELATED WORK

Vibration-based rock assessment has been studied primarily in the context of geohazard monitoring on unstable slopes. Ma et al. [5], [6] demonstrated that a remotely positioned laser Doppler vibrometer can relate vibration frequency and response amplitude to bonding condition and block stability. Du et al. [7] and Xie et al. [8] subsequently derived quantitative stability coefficients from natural frequency by modeling the rock-slope connection as a three-dimensional spring system. In the geophysics community, ambient-vibration monitoring of rock columns and cliff compartments has established that resonance frequency tracks constraint-condition changes over time [9], [10]. These studies establish vibration response as an effective indicator of rock constraint and stability, but they focus on long-term monitoring or stability classification rather than single-shot burial-depth estimation.

The most directly related analytical work is the spring-mass model of Jia et al. [11], which relates the damped natural frequency to rock mass m , burial depth h , geometry, and soil reaction coefficient k . Subsequent extensions applied this model to perilous rocks on soil slopes [12] and, most recently, to rapid boulder identification with experimentally validated safety-factor formulas [13]. As discussed in Section III, these formulations relate frequency to the ratio h/m , so recovering absolute depth requires additional information about mass, density, or soil-contact properties. Fukata et al. [14] addressed a closely related problem, estimating the in-soil penetration depth of railway-side boulders, using vibration measurement combined with 3D finite-element eigenvalue analysis and independently measured soil properties.

On the sensing side, laser Doppler vibrometry (LDV) enables non-contact, spatially resolved vibration measurement and is well established for modal testing and scanning vibration analysis [15], [16]. In the structural health monitoring community, neural networks trained on vibration features have been widely used for learning nonlinear mappings from dynamic response to structural parameters such as damage state, stiffness, and boundary conditions [17], [18], supporting the data-driven regression approach adopted here.

To the authors' knowledge, however, no published method estimates absolute burial depth from vibration data when buried geometry, effective soil stiffness, contact condition, and rock mass are simultaneously unknown.




Fig. 1. Simplified lumped spring-mass-damper model of a partially embedded rock. The rock is represented by a rigid mass m , and the surrounding soil or regolith is approximated by effective horizontal and vertical stiffness and damping terms, (k_x, ζ_x) and (k_z, ζ_z) . Burial depth changes these effective constraints, altering the measured vibration response.

III. PROBLEM ANALYSIS

A. Identifiability from Frequency Alone

Jia et al. [11] model a partially embedded block as a rigid body on a linear spring-damper system, as shown in Fig. 1. Their formulation gives the horizontal damped fundamental frequency as

$$f_n = \frac{\sqrt{1 - \zeta^2}}{2\pi} \sqrt{\frac{\left(\frac{a}{h_e}\right)^2}{\left(\frac{a}{h_e}\right)^2 + 1}} \sqrt{\frac{khb}{m}} \quad (1)$$

where a and b are the rock cross-section dimensions along and perpendicular to the vibration direction, m is the rock mass, h is the burial depth, h_e is the exposed height, ζ is the damping ratio, and k is the effective soil stiffness per unit contact area. Rearranging for the unknowns gives:

$$\frac{h}{m} = \frac{4\pi^2 f_n^2 (a^2 + h_e^2)}{(1 - \zeta^2) k b a^2} \quad (2)$$

Equation (2) exposes the fundamental identifiability constraint: h and m appear only as a ratio. A single frequency measurement, no matter how precise, cannot separately recover burial depth when mass is unknown. Moreover, even if all parameters in Equation (2) were independently measured, the linear spring-mass assumption underlying these equations cannot capture the nonlinear effects described in Section III-B, making closed-form inversion unreliable in practice.

Since a single-direction measurement constrains only one cross-section of the rock, a natural extension is to add a second measurement in an orthogonal direction under the same burial condition. In the second direction the cross-sectional dimensions swap roles in Equation (1): a becomes b and b becomes a . Dividing the two frequency equations cancels k , m , h , and ζ :

$$\left(\frac{f_{n,1}}{f_{n,2}}\right)^2 = \frac{a(b^2 + h_e^2)}{b(a^2 + h_e^2)} \quad (3)$$

Since a , b , and h_e are all measurable from the exposed face, the frequency ratio constrains only already-observable quantities; the burial depth h has cancelled out entirely. Moreover, this derivation assumes that both the effective stiffness k and damping ratio ζ are identical in both directions; as discussed

in Section III-B, soil stiffness is strain-dependent, so the two directions, excited at different force levels or over different contact areas, may mobilize different effective stiffnesses, preventing the clean cancellation that produces Equation (3). A second orthogonal measurement therefore cannot recover h through frequency ratios alone. The paired measurement nonetheless provides additional value by doubling the feature set, reducing noise, and introducing difference features that encode cross-sectional asymmetry, which is shape information inaccessible from any single direction.

B. Complications Beyond the Linear Model

Three physical effects cause the actual vibration response to deviate substantially from Equation (1):

Strain-dependent stiffness. Equation (1) treats k as a constant, but granular soils exhibit pronounced modulus reduction with increasing strain amplitude, consistent with the well-documented G/G_{\max} degradation curve of sand [19]. This means k is not merely unknown; it varies *within* a single measurement session depending on excitation level. As a result, although Equation (1) predicts that f_n increases monotonically with burial depth, only 36

However, for an ideal single-degree-of-freedom (SDOF) spring–mass system with fixed mass and damping ratio, the product of peak mobility M_p (the maximum value of the frequency response function, in (m/s)/N) and natural frequency is independent of stiffness [20]:

$$M_p \times f_n = \frac{1}{4\pi \zeta m} \quad (4)$$

In this idealized case, if the soil softens under a hard strike, f_n decreases while M_p increases, making $M_p \times f_n$ less sensitive to stiffness variation than either quantity alone. This relationship motivates providing both M_p and f_n as separate inputs so that the neural network can learn their combined dependence on burial condition.

Coupled translation–rocking modes. The spring–mass model assumes pure horizontal translation, but a partially embedded rock also exhibits a rocking mode in which it pivots about its buried base [21], [22]. By Chasles’ theorem [23], the rigid-body velocity field decomposes into a uniform translation and a rotation that varies linearly with height:

$$u(y) = U_{\text{trans}} + \theta \cdot y \quad (5)$$

where U_{trans} is the translational amplitude, θ is the rocking angular amplitude, and y is the height above the sand surface. The ratio θ/U_{trans} changes systematically with burial depth, contact geometry, and soil state. It is not captured by any scalar frequency measurement but is directly observable in the spatial variation of the FRF magnitude across scan points, motivating the spatial log-slope feature β introduced in Section IV-A.

Geometric and damping uncertainty. Equation (1) assumes a rectangular contact area bh and a damping ratio ζ that is independent of burial condition. In practice, irregular rock geometry, chamfered edges, and uneven settling cause the effective contact area to deviate from bh in a depth-dependent way, and ζ can also change with burial condition

TABLE I
ELEVEN-INPUT FEATURE VECTOR

#	Symbol	Physical meaning	Unit
1	\bar{w}	Mean cross-section width	cm
2	Δw	Width directional asymmetry	cm
3	h_e	Exposed height above sand	cm
4	\bar{f}_n	Mean resonance frequency	Hz
5	Δf_n	Frequency directional asymmetry	Hz
6	\bar{M}_p	Mean peak mobility	(m/s)/N
7	ΔM_p	Mobility directional asymmetry	(m/s)/N
8	$\bar{\zeta}$	Mean damping ratio	–
9	$\Delta \zeta$	Damping directional asymmetry	–
10	$\bar{\beta}$	Mean spatial vibration decay slope	1/cm
11	$\Delta \beta$	Decay slope directional asymmetry	1/cm

as more soil is engaged in dissipation. These factors introduce systematic, condition-dependent bias that an analytical model cannot correct without ground-truth contact measurements.

Taken together, these effects make direct inversion of Equation (1) unreliable, but they do not remove the burial-depth information from the measurement. Instead of relying on a single resonance frequency, the proposed approach uses features extracted from the FRF and spatial scan, allowing the regression model to learn the consistent response changes associated with burial depth. This motivates the feature extraction in Section IV and the neural network regression in Section VI.

IV. FEATURE EXTRACTION

For each measurement, the scanning laser Doppler vibrometer records velocity at multiple points on the exposed rock face, averaging multiple hammer impacts at each point. The H_1 frequency response function (FRF), $H_1(\omega) = S_{vf}(\omega)/S_{ff}(\omega)$, is computed from the velocity–force cross-spectrum S_{vf} and force auto-spectrum S_{ff} . The FRFs are spatially averaged across scan points to obtain the spectral features, while the point-wise FRF magnitudes are retained to compute the spatial decay feature. From these FRFs and direct geometric measurements, eleven features (Table I) are extracted as inputs to the neural network.

A. Spectral and Spatial Features from the FRF

Four features are extracted from each measurement’s FRF data. Fig. 2 illustrates the spectral quantities used to define the first three.

The **resonance frequency** f_n is the average of the peak frequency f_{peak} (frequency at maximum FRF magnitude) and the half-power center frequency $f_{\text{hp}} = (f_1 + f_2)/2$:

$$f_n = \frac{f_{\text{peak}} + f_{\text{hp}}}{2} \quad (6)$$

Averaging the two estimates provides robustness against noise spikes that can shift f_{peak} alone.

The **peak mobility** M_p is the maximum value of the averaged FRF, in (m/s)/N, encoding how freely the rock vibrates at resonance: heavier or more deeply constrained rocks generally produce lower mobility. Across the 16 blocks, M_p varies by a factor of 18 (0.06–1.1 (m/s)/N).




Fig. 2. Spatially averaged mobility FRF for Rock 5 at 60% burial. The peak mobility M_p and its half-power level $M_p/\sqrt{2}$ define the crossing frequencies f_1 and f_2 , from which the half-power center frequency $f_{hp} = (f_1 + f_2)/2$ and the peak frequency f_{peak} are used to compute the resonance-frequency estimate and damping ratio.

The **damping ratio** ζ is derived from the same half-power crossings f_1 and f_2 shown in Fig. 2, where $|H_1|$ falls to $M_p/\sqrt{2}$:

$$\zeta = \frac{f_2 - f_1}{2 f_{hp}} \quad (7)$$

This standard half-power bandwidth approximation [20] is valid for lightly-to-moderately damped systems; at the higher damping ratios observed here, the approximation introduces bias, but the neural network learns from the extracted values regardless of the analytical accuracy of Equation (7). Physically, deeper burial can increase the grain-to-grain contact area between rock and soil, dissipating more energy and often raising ζ . In our dataset, damping ratios range from 0.32 to 0.57.

The **spatial log-slope** β exploits the fact that the scanning LDV measures the FRF at 3–5 points on the exposed face (one at the top-right corner, one at the top-left corner, and one to three points in a vertical column along the center, as shown in Fig. 3) rather than at a single location. For each scan point, the peak FRF magnitude near f_n is extracted within a $\pm 2\%$ frequency window centered on f_n . The spatial log-slope is:

$$\beta = \frac{\ln(M_{p,top} / M_{p,bot})}{\Delta y} \quad (8)$$

where $M_{p,top}$ is the mean peak mobility of the two highest scan points, $M_{p,bot}$ is the peak mobility of the lowest scan point, and $\Delta y = y_{top} - y_{bot}$ is the vertical separation. The log-ratio makes β insensitive to uniform amplitude scaling, so it emphasizes the relative spatial decay of the response rather than the absolute mobility level. Physically, β captures the mode shape gradient described in Section III-B: from the coupled translation–rocking decomposition (Equation (5)), pure translation yields uniform mobility at all heights ($\beta \approx 0$), while a rocking component adds a contribution proportional to height above the pivot, increasing mobility toward the top (positive β). Because the relative contribution of each mode

changes systematically with burial depth, β encodes burial-depth information that no scalar frequency measurement can provide.

B. Geometric Features from Direct Measurement

Two features come from direct physical measurement rather than the vibrometer. The **exposed height** h_e is the visible vertical extent of the rock above the sand surface; it appears explicitly in the spring–mass model (Equation (1)) and, for a robotic system, would be obtained from RGB-D reconstruction [4]. The **cross-section width** w is the horizontal extent of the block face perpendicular to the vibrometer’s line of sight, corresponding to the dimensions a and b in the analytical model.

C. Two-Direction Encoding

At each burial depth, two measurements are taken in orthogonal directions, each exciting a different face of the block, yielding 90 paired samples in total from the 180 individual measurements. Each measurement yields the same four scalar features (f_n , M_p , ζ , β). The pair is encoded as:

$$\bar{x} = \frac{x_1 + x_2}{2}, \quad \Delta x = x_1 - x_2 \quad (9)$$

Mean features are intended to capture the average rock response, which is closely related to burial depth, while **difference features** capture directional asymmetry, which provides shape-related information. The cross-section widths of the two excited faces are encoded identically as \bar{w} and Δw .

The mean-and-difference decomposition of paired signals is a classical technique in paired-measurement statistics [24]. We apply this decomposition to paired multi-directional vibration features, where it provides a physics-motivated separation: the symmetric subspace (\bar{x}) isolates the burial-depth signal shared across directions, while the antisymmetric subspace (Δx) captures cross-sectional geometry.

An alternative is to pass the per-direction values directly as raw pairs ($f_{n,1}$, $f_{n,2}$, $M_{p,1}$, $M_{p,2}$, ...). Raw pairs, however, are order-dependent: which direction is labeled “1” is arbitrary, so the network must spend capacity learning an invariance that the mean/difference decomposition enforces by construction. Set-based architectures such as DeepSets [25] can learn such a decomposition from data, but require substantially more training samples than the 90 pairs available here; Section VII-C confirms that the hand-crafted encoding outperforms all learned alternatives at this dataset scale.

V. EXPERIMENTAL SETUP

A. Test Specimens

Sixteen cut concrete blocks are used as test specimens, spanning dimensions 3.2–14.0 cm per side, masses 412–2749 g, and densities 1.92–2.03 g/cm³. The masses are reported to characterize specimen density only and are not provided to the regression model. Concrete serves as a repeatable, rigid laboratory surrogate with rock-like density and stiffness, providing controlled geometry and ground-truth burial depth. This follows established practice in vibration-based rock assessment,




Fig. 3. A Polytec PSV-500 scanning laser Doppler vibrometer measures the velocity response of a concrete block partially buried in play sand.

where controlled-geometry specimens provide the foundation for subsequent field deployment [5], [11], [12].

Each block is partially buried in a box filled with commercial silica play sand (poorly graded quartz, median grain size $D_{50} \approx 0.5$ mm) as shown in Fig. 3. Two orientations are tested per block: with the tallest dimension (z) vertical and with the shortest dimension (x) vertical. At a given burial percentage, these yield different absolute depths (for example, a block with $x = 5.0$ cm and $z = 10.0$ cm at 40% burial gives $h = 2.0$ cm when x is vertical and $h = 4.0$ cm when z is vertical), increasing dataset diversity without fabricating additional specimens. Each block is tested at multiple burial levels per orientation, ranging from 15% to 75% of the vertical dimension. Below roughly 15–20%, the block slides rather than vibrates under hammer excitation; above roughly 75%, too little face remains exposed for both actuation and scanning. Across all blocks and conditions, burial depths range from 0.99 to 8.40 cm, exposed heights from 1.23 to 11.90 cm, and resonance peak frequencies from 19.4 to 45.0 Hz.

B. Measurement Equipment and Procedure

Vibration measurements are performed with a Polytec PSV-500 scanning laser Doppler vibrometer (SLDV), which provides non-contact velocity measurement via the Doppler shift of a reflected laser beam (Fig. 3). The PSV-500 uses a visible helium-neon (HeNe) laser optimized for highly reflective surfaces. Because concrete has low surface reflectivity, retroreflective tape was applied at each scan point to maintain signal quality.

A Kistler 9724A5000 instrumented impulse hammer provides the excitation force, with its force transducer connected to the vibrometer’s reference channel so the software can compute the H1 frequency response function directly. The hammer uses a quartz Piezotron sensing element with a 0–5000 N force range and 1 mV/N sensitivity; a plastic impact tip was used for all measurements. At each burial depth, the vibrometer scans 3–5 points on the exposed face following the layout described in Section IV-A, with three hammer impacts averaged per scan point to produce each FRF. Two measurements are taken per burial depth, one per perpendicular face, forming the paired inputs described in Section IV. All FRF data are computed within 0–400 Hz; all observed resonances and half-power crossings fell below 49.4 Hz.

C. Applicability to Lunar Environments

The experimental setup differs from the lunar surface in gravity, soil composition, and specimen geometry. While the specific learned weights would need to be retrained with data collected in the target environment, the underlying physical principles (soil-coupled resonance, mode-shape dependence on embedment, and amplitude-dependent stiffness) are not Earth-specific, and the sensing modality, feature-extraction pipeline, and network architecture are directly applicable to other soil or regolith conditions.

VI. NEURAL NETWORK METHODOLOGY

A. Leave-One-Rock-Out Cross-Validation

Given the small dataset, standard random-split cross-validation would allow measurements from the same block at different burial depths to appear in both training and test sets, inflating performance estimates. We adopt *strict Leave-One-Rock-Out (LORO)* cross-validation: in each of the 16 folds, *all* measurements from a single physical block (across all burial depths and experimental conditions) are withheld as the test set. The model is trained on the remaining 15 blocks.

B. Architecture

A Multilayer Perceptron (MLP) regressor maps the eleven-dimensional feature vector to burial depth in centimeters. The architecture consists of two hidden layers with 16 and 8 neurons, tanh activations, and an ℓ_2 regularization coefficient $\alpha = 5.0$. The Adam optimizer is used with an initial learning rate of 5×10^{-4} , early stopping after 300 iterations with no improvement (tolerance 10^{-6}), and a maximum of 50,000 iterations. This configuration was selected from a grid search over hidden-layer sizes, activation functions, and regularization strengths as a compact model with strong mean LORO performance. Due to the small dataset size, results exhibit moderate sensitivity to random initialization. All input features are standardized to zero mean and unit variance using training-set statistics; the target variable h (cm) is not scaled.

VII. RESULTS AND DISCUSSION

The results are organized around three questions: whether the proposed FRF-derived features improve burial-depth estimation, which physical quantities carry the strongest predictive signal, and whether the two-direction mean/difference encoding is preferable to learned alternatives. Together, these analyses show that the best performance comes from combining visible geometry with damping, spatial vibration decay, and paired-direction information in a compact feature representation.

A. Feature Ablation Study

The first experiment evaluates whether the additional vibration features proposed in this work improve burial-depth prediction under strict LORO validation. Table II compares the proposed two-direction feature ablation against a single-direction benchmark under the same LORO protocol, reporting mean \pm std over 20 fixed random seeds (0–19). The first four

TABLE II

FEATURE ABLATION: LORO CROSS-VALIDATION PERFORMANCE

$h_e, \bar{w},$ \bar{f}_n, \bar{M}_p ($+\Delta$)	$\bar{\zeta},$ $\Delta\zeta$	$\bar{\beta},$ $\Delta\beta$	Dim.	MAPE (%) mean \pm std	RMSE (cm) mean \pm std
<i>Two-direction</i>					
✓			7	33.4 \pm 0.6	1.220 \pm 0.014
✓	✓		9	27.0 \pm 0.3	1.047 \pm 0.009
✓		✓	9	26.5 \pm 0.2	1.081 \pm 0.005
✓	✓	✓	11	25.3 \pm 0.4	1.029 \pm 0.009
<i>Single-direction</i>					
✓ [†]	✓ [†]	✓ [†]	7	31.6 \pm 0.6	1.216 \pm 0.009

[†]Geometry ($h_e, \bar{w}, \Delta w$) retained; vibration quantities (f_n, M_p, ζ, β) are single-direction values with no Δ counterpart.

rows use the mean/difference encoding and progressively add damping and spatial-decay features to the baseline geometry, resonance-frequency, and peak-mobility inputs. The final row reports a separate single-direction benchmark and is included to show the value of paired directional encoding.

The baseline model uses resonance frequency, peak mobility, exposed height, and cross-section width, encoded using the same two-direction mean/difference representation as the final model. Adding the mean and asymmetry damping features, $\bar{\zeta}$ and $\Delta\zeta$, reduces MAPE from 33.4% to 27.0%. Adding the corresponding spatial log-slope features, $\bar{\beta}$ and $\Delta\beta$, gives a comparable improvement, reducing MAPE to 26.5%, confirming that the spatial variation of vibration amplitude carries burial-depth information beyond what the spectral peak captures.

The strongest result is obtained when damping and spatial-decay features are used together. The full 11-feature model reduces MAPE from 33.4% to 25.3%, a 24% relative improvement, and reduces RMSE from 1.22 to 1.03 cm. The improvement beyond the β -only model indicates that the two features are complementary: ζ captures energy dissipation in the frequency domain, while β captures how the vibration field varies over the exposed rock face. This supports the central premise of the method: the resonance peak location alone is an incomplete descriptor of burial depth, which is also reflected in the width of the resonance and the spatial distribution of motion across the rock face.

The single-direction benchmark provides a useful reference point. It treats each single-direction measurement as an independent sample with seven features: the four FRF features (f_n, M_p, ζ, β), exposed height h_e , and the mean and difference of the two cross-section widths ($\bar{w}, \Delta w$). As discussed in Section III, these richer FRF features make single-direction estimation possible, giving 31.6% MAPE. The full two-direction model, which uses the same four FRF quantities under the paired mean/difference encoding, achieves 25.3% MAPE, showing that directional asymmetry provides additional shape-related information beyond what any single measurement direction can capture.

Fig. 4(a) plots predicted versus true burial depth for the 11-input model under LORO cross-validation. The estimates remain close to the identity line across the full 0.99–8.40 cm burial-depth range, spanning the diverse block sizes and burial

TABLE III

PERMUTATION FEATURE IMPORTANCE: 11-INPUT MODEL

Rank	Feature	Δ RMSE (cm) mean \pm std	Δ RMSE (%) mean \pm std
1	h_e	1.539 \pm 0.028	321 \pm 6
2	\bar{w}	1.202 \pm 0.032	250 \pm 7
3	$\bar{\zeta}$	0.848 \pm 0.024	177 \pm 5
4	$\bar{\beta}$	0.820 \pm 0.031	171 \pm 6
5	Δw	0.631 \pm 0.031	131 \pm 6
6	\bar{f}_n	0.328 \pm 0.017	68 \pm 3
7	ΔM_p	0.322 \pm 0.020	67 \pm 4
8	\bar{M}_p	0.264 \pm 0.013	55 \pm 3
9	Δf_n	0.240 \pm 0.012	50 \pm 3
10	$\Delta\zeta$	0.105 \pm 0.006	22 \pm 1
11	$\Delta\beta$	0.103 \pm 0.009	21 \pm 2

conditions present in the dataset. Fig. 4(b) shows the corresponding total-volume estimates computed from the predicted burial depth and the known cross-sectional geometry. The volume scatter is tighter because the predicted burial depth affects only one dimension of the derived volume, while the other block dimensions are directly measured.

B. Feature Sensitivity Analysis

The ablation study shows that the full feature set improves prediction accuracy, but it does not show which inputs the trained model relies on most strongly. To interpret the learned regressor, feature importance is assessed by permutation importance [26]: each feature is randomly shuffled one at a time across 20 random permutations per seed, and the mean increase in RMSE relative to the unshuffled baseline is recorded. The model is trained on the full paired dataset for this analysis; results are reported as Δ RMSE in centimeters and as a percentage of the unshuffled model's RMSE, which is normalized by the full-dataset RMSE rather than the LORO RMSE in Table II. This analysis is used only to rank feature sensitivity and is not reported as an independent generalization estimate. Table III reports the full ranking.

The largest sensitivity is observed for h_e and \bar{w} , both of which are direct geometric measurements. This is expected from the analytical model in Equation (1): exposed height constrains the visible portion of the rock, while cross-section width affects the effective mass and contact geometry involved in the vibration response. Their dominant ranking confirms that the network is not learning burial depth from vibration data alone, but from vibration features interpreted in the context of visible geometry.

The next two features are the quantities introduced in this work: $\bar{\zeta}$ and $\bar{\beta}$. Both rank above the mean resonance frequency \bar{f}_n , with $\bar{\zeta}$ increasing RMSE by 0.85 ± 0.02 cm when shuffled and $\bar{\beta}$ by 0.82 ± 0.03 cm. Their rankings are close and within each other's standard deviation, so no strong ordering between them can be claimed. This ranking supports the ablation result in Table II: the spatial vibration decay and half-power bandwidth contain burial-depth information that is not captured by the resonance peak location alone.

The remaining features clarify the role of the two-direction encoding. The mean vibration quantities carry most of the




Fig. 4. LORO cross-validation predictions vs. ground truth for the 11-input model (seed = 0). (a) Burial depth h (MAPE = 25.4%, RMSE = 1.03 cm). (b) Total rock volume V (MAPE = 9.2%, RMSE = 70 cm³). Dotted lines represent a perfect prediction $\pm 25\%$.

burial-depth signal, while the difference features contribute at lower magnitude by encoding cross-sectional asymmetry. Among the difference features, Δw and ΔM_p rank highest, while Δf_n , $\Delta \beta$, and $\Delta \zeta$ contribute more modestly. This is consistent with the strain-dependent stiffness discussed in Section III-B: frequency differences between the two directions can be affected by excitation-level variation, whereas the mean damping, mobility, and spatial-decay responses provide more stable indicators of the rock–soil interaction. Thus, the sensitivity analysis refines the ablation result: the paired measurement is useful primarily because it stabilizes the mean response and adds secondary shape asymmetry. This also explains why single-direction measurement remains viable but less accurate: it preserves the main FRF and geometry cues, but removes the paired asymmetry information. The single-direction benchmark in Table II carries a 6 percentage point penalty in MAPE (31.6% vs. 25.3%), confirming that burial-depth estimation remains feasible from a single strike direction when paired measurement is impractical.

C. Input Encoding and Architecture Comparison

The sensitivity analysis suggests that the symmetric and antisymmetric parts of the paired measurement play different roles: mean features carry the dominant burial-depth signal, while difference features provide secondary shape information. The next experiment tests whether this structure should be imposed explicitly through the proposed mean/difference encoding, or learned directly from raw paired inputs using more flexible architectures. Table IV compares the final model against alternative encodings and neural network architectures, all evaluated under identical LORO conditions with the same 11 feature values per paired measurement.

The linear and polynomial baselines first show that the mapping from the proposed features to burial depth is nonlinear. Linear regression gives 37.6% MAPE, while the degree-2

polynomial model performs even worse at 42.6% MAPE. The polynomial expansion produces 77 nonconstant basis functions from only 90 total paired samples, making it prone to overfitting under LORO validation. The 11-feature MLP therefore improves performance not simply by adding parameters, but by providing a regularized nonlinear mapping that remains small enough for the dataset.

All neural model results in Table IV are reported as mean \pm std over seeds 0–19 to account for sensitivity to random initialization; linear models are deterministic and report a single value. All neural alternatives were tuned by the same grid search procedure used for the proposed model, sweeping architecture size and regularization strength; the comparisons are made at each model’s best configuration. Despite this, every learned encoding performs worse than the proposed mean/difference representation. The raw-pairs model reaches 31.9% MAPE, indicating that the network can partially learn the paired structure but must spend capacity discovering the mean/difference relationship from limited data. DeepSets [25] and the Siamese model reach 36.2% and 34.9% MAPE at their grid-search optimum, still trailing the proposed model by 9–11 percentage points. Their pooling and concatenation operations do not impose the directional asymmetry as directly as the mean/difference encoding, and at the current dataset scale they do not have enough data to learn it reliably from scratch. Transformer models also underperform, with MAPE between 29.2% and 30.6%, because their additional capacity is not supported by the current dataset size.

Overall, the architecture comparison is consistent with the ablation and sensitivity studies. The best result is not obtained by making the architecture more complex, but by encoding the physics of the measurement in a compact form before regression. At the present dataset scale, the hand-crafted mean/difference decomposition in Equation (9) provides a better bias than learned set-based or attention-based alternatives.

TABLE IV
INPUT ENCODING AND ARCHITECTURE COMPARISON (LORO CV; NEURAL MODELS: SEEDS 0–19)

Model	Variant	Configuration	MAPE (%) mean \pm std	RMSE (cm) mean \pm std
Linear	Mean/diff	Linear regression	37.6	1.43
	Mean/diff	Polynomial (deg. 2)	42.6	1.62
MLP	Mean/diff (proposed)	(16,8) tanh, $\alpha=5$	25.3 ± 0.4	1.029 ± 0.009
	Raw pairs ($f_{n,1}, f_{n,2}, \dots$)	(16,8) tanh, $\alpha=5$	31.9 ± 0.3	1.226 ± 0.009
Set-based	DeepSets (mean-pool)	$\phi=(4), \rho=(4), \text{wd}=10^{-2}$	36.2 ± 1.9	1.573 ± 0.062
	Siamese (concat-pool)	$\phi=(4), \rho=(4), \text{wd}=10^{-2}$	34.9 ± 2.7	1.521 ± 0.118
Transformer	Transformer-S	$d=16$, 1 layer, 2 heads	29.2 ± 2.7	1.423 ± 0.126
	Transformer-M	$d=32$, 2 layers, 4 heads	30.6 ± 2.2	1.538 ± 0.063

Larger and more diverse datasets may eventually make those architectures useful, but the current results favor a small MLP with physically structured inputs.

D. Limitations

The preceding results show that burial-depth information is recoverable from soil-coupled vibration measurements, but they also indicate where the present model is most constrained. The strongest features are geometric quantities, and the learned vibration trends are based on a controlled set of concrete blocks in one sand condition. This leaves three main limitations that must be addressed before the method can be considered general across natural rocks or lunar regolith.

First, all sixteen concrete blocks have near-identical densities ($1.92\text{--}2.03\text{ g/cm}^3$), so the network has not learned to discriminate across rock types. Lunar rocks span a wider range, from anorthosite at approximately 2.4 g/cm^3 to basalt at approximately 3.0 g/cm^3 [27], and density directly affects the frequency response through the mass term in Equation (1). Second, all training data were collected with a single sand type, so generalization across different compaction states, grain sizes, and regolith compositions remains to be validated. Third, the dataset contains only 16 blocks with controlled rectangular geometry, which is small relative to the feature dimensionality and does not capture the irregular shapes of natural rocks.

Larger and more diverse datasets are expected to improve both prediction accuracy and feature reliability. They would also provide a more appropriate setting for the set-based and transformer architectures in Table IV, which currently overfit or underperform because the available paired dataset is too small. In this sense, these results are a controlled proof of concept, defining both the value of the proposed feature set and the data diversity needed for future field or lunar-analog validation.

VIII. CONCLUSION

This work shows that the hidden burial depth of a partially embedded rock can be inferred from its soil-coupled vibration response using a small set of physically interpretable features. Starting from the limitation of frequency-only spring–mass models, the proposed pipeline combines impulse excitation, scanning laser vibrometry, half-power damping, spatial vibration decay, and two-direction mean/difference encoding in an

eleven-feature neural-network regressor. Under strict Leave-One-Rock-Out validation, the model predicts burial depth for unseen concrete blocks with 25.3% MAPE and 1.03 cm RMSE, while the damping and spatial-decay features reduce MAPE by 24% relative to a geometry, frequency, and mobility baseline. These results indicate that burial depth is encoded not only in the resonance peak, but also in the bandwidth and spatial distribution of the vibration response.

More broadly, this work establishes that non-contact vibration sensing can reveal subsurface constraints that are invisible to cameras and geometry reconstruction. Although the learned weights must be retrained for the target environment, the sensing modality and feature-extraction pipeline are applicable across granular-media conditions. On the Moon, integrating this pipeline with robotic excitation and RGB-D geometry estimation [4] would allow an autonomous excavator to assess burial condition before committing energy to removal. On Earth, the same approach may support slope-stability assessment, boulder hazard classification, and construction-site rock characterization.

Future work will focus on broadening specimen diversity, validating the approach across regolith analogs and compaction states, and integrating automated excitation with RGB-D perception in a robotic excavation pipeline. We will extend the method to irregular natural rocks of varying size, shape, and density, where semantic rock-type classification from imagery may provide a useful density prior. We will also replace the manual impulse hammer with a robot-mounted automated striker or solenoid actuator and evaluate whether single-direction measurement can approach two-direction accuracy with a larger training set. These extensions would move the method toward a perception-to-action workflow in which predicted burial depth and uncertainty support decisions to remove, avoid, or re-measure a rock.

The dataset and code used in this study are available at [28].

ACKNOWLEDGMENT

The authors would like to thank Dr. Oumar Barry for providing the test equipment, Dr. Suyi Li for guidance in this work, and the members of the Virginia Tech Field and Space Experimental Robotics (FASER) Laboratory, Kareem Basit, Duke Cassels-Smith, and Timothy Perez for their assistance in data collection.

REFERENCES

- [1] S. Giel, J. Hurrell, S. Santra, A. Mishra, K. Uno, and K. Yoshida, "Design and development of a modular bucket drum excavator for lunar ISRU," in *Proc. Int. Conf. Space Robotics (iSpaRo)*, 2025.
- [2] R. Lee, B. Younes, A. Pletta, J. Harrington, R. Q. Wong, and W. Whitaker, "CraterGrader: Autonomous robotic terrain manipulation for lunar site preparation and earthmoving," in *Proc. Robotics: Science and Systems (RSS)*, 2024.
- [3] J. Walther, R. L. Johns, H. Kolvenbach, V. T. Bickel, and M. Hutter, "Autonomous construction of lunar infrastructure with in-situ boulders," *Frontiers in Space Technologies*, vol. 5, p. 1345337, 2024.
- [4] Y. Ruan and E. Komendera, "Lunar rock reconstruction with RGB-D images," in *2025 IEEE SENSORS*, 2025, pp. 1–4.
- [5] G. Ma, K. Sawada, H. Saito, F. Uehan, and A. Yashima, "Study on evaluating rock block stability by using a remotely positioned laser doppler vibrometer," *International Journal of GEOMATE*, vol. 2, no. 4, pp. 247–252, 2012.
- [6] G. Ma, K. Sawada, and A. Yashima, "Experimental study of the applicability of the remotely positioned laser doppler vibrometer to rock-block stability assessment," *Rock Mechanics and Rock Engineering*, vol. 48, no. 2, pp. 787–802, 2015.
- [7] Y. Du, M. Xie, Y. Jiang, B. Li, Y. Gao, and Q. Liu, "Safety monitoring experiment of unstable rock based on natural vibration frequency," *Rock and Soil Mechanics*, vol. 37, no. 10, pp. 3035–3040, 2016.
- [8] M. Xie, W. Liu, Y. Du, and H. Wang, "The evaluation method of rock mass stability based on natural frequency," *Advances in Civil Engineering*, vol. 2021, p. 6652960, 2021.
- [9] P. Bottelin, D. Jongmans, L. Baillet, T. Lebourg, D. Hantz, C. Lévy, O. Le Roux, H. Cadet, L. Lorier, J.-D. Rouiller *et al.*, "Spectral analysis of prone-to-fall rock compartments using ambient vibrations," *Journal of Environmental and Engineering Geophysics*, vol. 18, no. 4, pp. 205–217, 2013.
- [10] C. Colombero, L. Baillet, C. Comina, D. Jongmans, E. Larose, J. Valentin, and S. Vinciguerra, "Integration of ambient seismic noise monitoring, displacement and meteorological measurements to infer the temperature-controlled long-term evolution of a complex prone-to-fall cliff," *Geophysical Journal International*, vol. 213, no. 3, pp. 1876–1897, 2018.
- [11] Y. Jia, G. Song, T. Jiang, G. Yu, H. Liu, and X. Pan, "Buried depth calculation of the slope of the unstable rock based on natural vibration frequency," *Frontiers in Earth Science*, vol. 10, p. 929825, 2022.
- [12] Y. Jia, G. Song, L. Wang, T. Jiang, J. Zhao, and Z. Li, "Research on stability evaluation of perilous rock on soil slope based on natural vibration frequency," *Applied Sciences*, vol. 13, no. 4, p. 2406, 2023.
- [13] Y. Jia, Y. Liang, G. Song, H. Wang, M. Xie, S. Hua, Z. Chen, W. Ma, and T. Niu, "Study on rapid identification of soil slope boulders based on natural vibration frequency," *Scientific Reports*, vol. 16, p. 3924, 2026.
- [14] T. Fukata, T. Taniguchi, and S. Shibuya, "A proposal of estimation method for rockfall risk degree based on vibration measurement of the boulder," *Journal of Japan Society of Civil Engineers, Ser. C (Geosphere Engineering)*, vol. 69, no. 1, pp. 140–151, 2013.
- [15] S. J. Rothberg, M. S. Allen, P. Castellini, D. Di Maio, J. J. J. Dirckx, D. J. Ewins, B. J. Halkon, P. Muysshondt, N. Paone, T. Ryan, H. Steger, E. P. Tomasini, S. Vanlanduit, and J. F. Vignola, "An international review of laser doppler vibrometry: Making light work of vibration measurement," *Optics and Lasers in Engineering*, vol. 99, pp. 11–22, 2017.
- [16] P. Castellini, M. Martarelli, and E. P. Tomasini, "Laser Doppler Vibrometry: Development of advanced solutions answering to technology's needs," *Mechanical Systems and Signal Processing*, vol. 20, no. 6, pp. 1265–1285, 2006.
- [17] D. A. Sofge, "Structural health monitoring using neural network based vibrational system identification," in *Proc. ANZIS '94 Australian New Zealand Conf. Intelligent Information Systems*, 1994, pp. 91–94, accessible via arXiv:0705.3669.
- [18] B. Puruncajas, Y. Vidal, and C. Tutivén, "Vibration-response-only structural health monitoring for offshore wind turbine jacket foundations via convolutional neural networks," *Sensors*, vol. 20, no. 12, p. 3429, 2020.
- [19] B. O. Hardin and V. P. Drnevich, "Shear modulus and damping in soils: Measurement and parameter effects," *Journal of the Soil Mechanics and Foundations Division*, vol. 98, no. SM6, pp. 603–624, 1972.
- [20] D. J. Ewins, *Modal Testing: Theory, Practice and Application*, 2nd ed. Baldock, UK: Research Studies Press, 2000.
- [21] G. Gazetas, "Analysis of machine foundation vibrations: State of the art," *International Journal of Soil Dynamics and Earthquake Engineering*, vol. 2, no. 1, pp. 2–42, 1983.
- [22] A. S. Veletsos and Y. T. Wei, "Lateral and rocking vibrations of footings," *Journal of the Soil Mechanics and Foundations Division, ASCE*, vol. 97, no. SM-9, pp. 1227–1248, 1971.
- [23] M. Chasles, "Note sur les propriétés générales du système de deux corps semblables entr'eux et placés d'une manière quelconque dans l'espace; et sur le déplacement fini ou infiniment petit d'un corps solide libre," *Bulletin des Sciences Mathématiques, Astronomiques, Physiques et Chimiques*, vol. 14, pp. 321–326, 1830.
- [24] J. M. Bland and D. G. Altman, "Statistical methods for assessing agreement between two methods of clinical measurement," *The Lancet*, vol. 327, no. 8476, pp. 307–310, 1986.
- [25] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R. Salakhutdinov, and A. J. Smola, "Deep sets," in *Advances in Neural Information Processing Systems*, vol. 30. Curran Associates, Inc., 2017.
- [26] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [27] W. S. Kiefer, R. J. Macke, D. T. Britt, A. J. Irving, and G. J. Consolmagno, "The density and porosity of lunar rocks," *Geophysical Research Letters*, vol. 39, no. 7, p. L07201, 2012.
- [28] Y. Ruan and E. Komendera, "rock-vib-burial-depth," GitHub, 2026. [Online]. Available: <https://github.com/JXNICHOLAS/rock-vib-burial-depth>



Yiyan Ruan received the B.S. degree in mechanical engineering from Virginia Polytechnic Institute and State University, Blacksburg, VA, USA, where he is currently pursuing the Ph.D. degree in mechanical engineering. His research interests include multi-axis additive manufacturing, 3D scene understanding, robotic automation, and robotic perception for planetary exploration.



Erik Komendera is an Assistant Professor of Mechanical Engineering at Virginia Polytechnic Institute and State University in Blacksburg, VA. He earned his Ph.D. from the University of Colorado in Boulder, CO. He was previously a researcher at NASA Langley Research Center in Hampton, VA. His research interests include robotic perception, Bayesian estimation, operations planning, autonomous construction, in-space assembly and servicing, and robot-enhanced additive and subtractive manufacturing. His current and previous collaborators include NASA, DOD, NSF, and the space industry.