Modeling Lunar Surface Charging Using Physics-Informed Neural Networks

Niloofar Zendehdel
Mechanical & Aerospace Eng. Dept.
Missouri University of Science & Technology
Rolla, USA
nzx9d@mst.edu

Adib Mosharrof
Dept. of Computer Science
University of Kentucky
Lexington, USA
adib.mosharrof@uky.edu

Katherine Delgado
Dept. of Computer Science
University of Kentucky
Lexington, USA
kadelgado@uky.edu

Daoru Han
Mechanical & Aerospace Engr. Dept.
Missouri University of Science & Technology
Rolla, USA
handao@mst.edu

Xin Liang
Dept. of Computer Science
University of Kentucky
Lexington, USA
xliang@uky.edu

Tong Shu
Dept. of Computer Science & Engr.
University of North Texas
Denton, USA
tong.shu@unt.edu

***Abstract*—Modeling the electric potential profile above the lunar surface is critical for understanding surface charging and interactions with the space environment. Traditional methods like Particle-in-Cell (PIC) simulations are highly accurate but computationally expensive. To address this, we propose a hybrid approach using a Multi-Layer Perceptron (MLP) architecture in both data-driven neural networks and Physics-Informed Neural Networks (PINNs). The PINN component incorporates physical laws directly into the training process, ensuring physical consistency, while the data-driven component captures complex patterns. This combination offers a significant reduction in computational cost compared to PIC methods while maintaining high modeling accuracy. Our results show that the proposed method effectively represents the electric potential profile above the lunar surface, even with limited data.**

***Index Terms*—Physics-Informed Neural Network, Lunar Surface Charging, electric potential profile.**

I. Introduction

The Moon, as an airless celestial body, presents a unique environment where its surface is continually exposed to the surrounding space. This exposure leads to intricate interactions between the lunar surface and the near-lunar plasma environment, making its study pivotal for understanding the broader dynamics of space-surface interactions. Plasma, often termed the fourth state of matter, comprises charged particles, namely ions and electrons. Its presence around the Moon is a result of various cosmic phenomena, including solar winds and cosmic rays. When this plasma, coupled with solar ultraviolet (UV) radiation, interacts with the lunar surface, it induces electrical charges. This charging mechanism varies depending on the Moon's position relative to the Sun. On the dayside, where the lunar surface is directly exposed to solar UV radiation, photocharging becomes the primary mechanism. This process results in the emission of photoelectrons from the surface, leading to a net positive charge. The region immediately above the surface, where this positive charge dominates, is termed the "photoelectron sheath." This sheath is relatively thin, extending roughly up to a meter above the surface.

In contrast, the nightside of the Moon, shielded from direct solar UV radiation, experiences a different charging mechanism. Here, the influx of plasma electrons from the surrounding space causes the lunar surface to acquire a net negative charge. This negatively charged region gives rise to the "Debye sheath," which can be much thicker than its dayside counterpart, extending up to a kilometer. The extent of this sheath is governed by the Debye length, a fundamental parameter in plasma physics. It quantifies the distance over which electric fields in a plasma get screened or attenuated, ensuring quasi-neutrality. Understanding the Debye length is crucial as it provides insights into the electric field's behavior near the surface, influencing various phenomena like dust levitation and surface arcing (fig 1) [1].

The interaction between the lunar surface and plasma is further complicated by various factors such as solar UV intensity, ambient plasma variations, surface composition, magnetic anomalies, and the lunar wake. Additionally, the Moon's passage through the Earth's magnetosphere introduces complexities, especially in the plasma sheet, which is more tenuous and hotter than the solar wind. Surface charging is not merely a theoretical concern; it has practical implications for lunar exploration. The differential charging of objects on the surface could lead to unanticipated electrical discharges, and the transport of charged dust presents potential hazards. Understanding the near-lunar plasma environment is essential for manned and robotic surface exploration activities, as well as for scientific observations such as lunar-based astronomy.

Analytical solutions, rooted in mathematical principles, often adopt simplifying assumptions to make the problem tractable. While they provide a clear mathematical framework, these solutions might overlook certain complexities inherent to the real-world scenario. Particle-in-cell (PIC) methods, as detailed in [3], stand out by simulating the collective

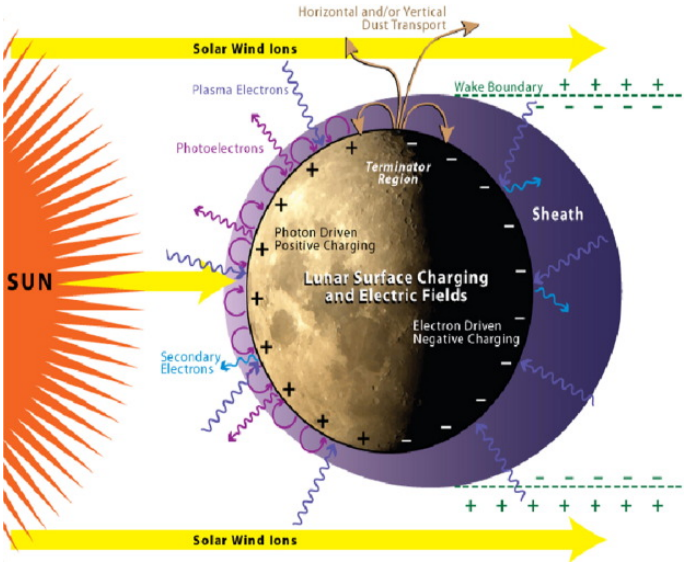

Fig. 1: The Lunar Dust-Plasma Environment [2]

behavior of individual charged particles as they interact with electromagnetic fields. PIC operates by dividing the simulation space into cells and tracking the motion of charged particles through these cells, calculating particle dynamics based on local electromagnetic fields. This fine-grained approach, though offering a detailed representation, comes at the cost of high computational overhead due to the necessity of simulating each particle individually and updating fields at every time step. The computational burden becomes significant because accurate simulations require a large number of particles and fine temporal and spatial resolutions to resolve complex charge distributions and interactions. Consequently, PIC methods, while precise, are resource-intensive, making them less practical for large-scale or real-time applications.

In contrast, neural networks have shown remarkable success in learning and generalizing complex patterns from data. Neural networks provide a more computationally efficient solution by approximating these complex relationships through learned weights and activation functions, making predictions using only matrix multiplications rather than solving differential equations at each time step. Once trained, they can generate accurate predictions almost instantaneously, resulting in a significantly reduced computational load compared to PIC.

In this paper, we employ a Multi-Layer Perceptron (MLP) based Physics-Informed Neural Networks (PINNs), which embed the governing physical laws directly into the training process. This incorporation of physics ensures that the model adheres to known physical constraints, even in regions with sparse data, while avoiding the need for particle-based simulations. In this study, we leverage a combination of PINNs and traditional data-driven neural networks to model the electric potential profile above the lunar surface. The PINN component enforces the physical consistency by minimizing the residuals of the governing equations, while the data-driven component captures complex patterns that might be overlooked by purely physics-based models. This hybrid approach allows us to achieve a balance between computational efficiency and physical accuracy, significantly reducing the computational cost associated with traditional PIC methods while maintaining high fidelity in modeling electric potential profiles. The contributions of this paper are:

- Proposing a hybrid modeling approach that combines PINNs and data-driven models to accurately capture the electric potential profile with reduced computational cost.
- Providing an analysis of different loss functions and optimization strategies to enhance the efficiency and performance of neural network models in space environment simulations.

The rest of the paper is organized as follows: In Section 2, we formulate the problem of surface charging. Section 3 introduces physics-informed neural networks (PINNs) as the method to solve this problem, detailing the architecture and training process that enable the network to capture the underlying physics of surface charging. In Section 4, results are demonstrated, showcasing the effectiveness of the proposed neural network approach in accurately modeling the electric potential profile above the lunar surface. Finally, Section 5 concludes the paper, summarizing the key findings and highlighting the contributions of this work to the field of space environment modeling.

II. Problem Formulation

The governing equation pivotal to understanding surface potential within a plasma environment is the Poisson equation. Rooted in the core principles of electromagnetism, this equation emerges from Gauss's law for electricity. The Poisson equation serves as a mathematical representation that captures the relationship between the distribution of electric charge in space and the resulting electric potential. It is expressed as:

$$\nabla^2\phi = -\rho/\epsilon_0 \tag{1}$$

Where ϕ represents the electric potential, a scalar quantity that describes the amount of potential energy a unit charge would have at a specific location in space, ρ is the charge density, which quantifies the amount of electric charge per unit volume, and ϵ_0 is the permittivity of free space. With these definitions in place, equation (1) can be further transformed and represented as an ordinary differential equation (ODE) [5]:

$$\hat{\phi}" = \left(1 + \frac{2\hat{\phi}}{M^2}\right)^{-1/2} - e^{-\hat{\phi}} \tag{2}$$

Where $\hat{\phi} = \dfrac{-e\phi}{KT_e}$ and $M = \dfrac{u_0}{C_s} = \dfrac{u_0}{\sqrt{\dfrac{KT_e}{m_i}}}$, u_0 is the drifting velocity, K is Boltzmann constant, m_i is mass of ion, T_e is electron temperature. Given the equation (2), we can integrate both sides once by multiplying them by $\hat{\phi}'$

$$\frac{1}{2}\left(\hat{\phi}'^2 - \hat{\phi_0}'^2\right) = M^2\left[\left(1 + \frac{2\hat{\phi}}{M^2}\right)^{\frac{1}{2}} - 1\right] + e^{-\hat{\phi}} - 1 \tag{3}$$

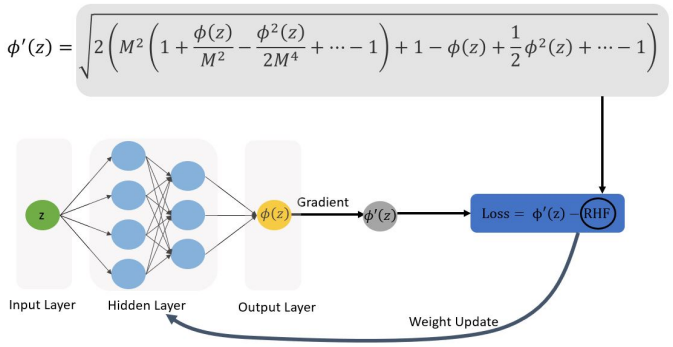

Fig. 2: PINN structure using residual loss

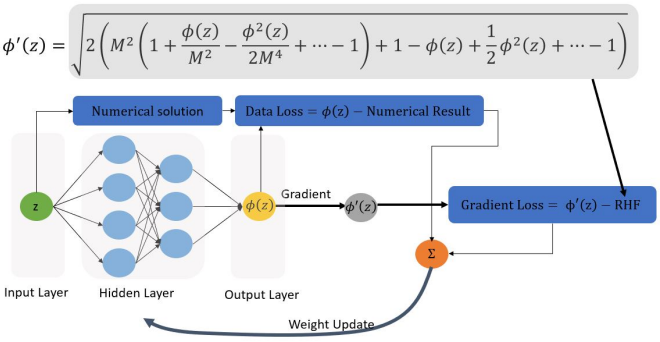

Fig. 3: PINN structure using residual loss and data loss

Using the Taylor series, we can expand the right-hand terms for $\hat{\phi} << 1$:

$$\hat{\phi}' = \sqrt{2\left(M^2\left[1+\frac{\hat{\phi}}{M^2}-\frac{\hat{\phi}^2}{2M^4}+\cdots-1\right]+1-\hat{\phi}+\frac{\hat{\phi}^2}{2}+\cdots-1\right)} \tag{4}$$

Our objective is to find $\hat{\phi}$ so that it satisfies the ordinary differential equation (ODE) represented by equation (4). This function represents the electric potential profile above the lunar surface, capturing the essential physics of the problem.The next sections will detail the specific architecture and training process of the neural network, demonstrating how it is designed to find $\hat{\phi}$ that satisfies the given ODE.

III. Method

In the context of modeling the electric potential profile above the lunar surface, we employ a MLP-based Physics-Informed Neural Networks (PINNs), which integrate known physical laws directly into the training process [6]. Unlike traditional neural networks, which are primarily data-driven and learn patterns directly from the training data, PINNs are designed to satisfy not only the data constraints but also the governing differential equations that describe the physical system. This dual nature of PINNs allows them to produce predictions that are consistent with both the observed data and the known physics of the problem, thereby reducing the risk of generating physically implausible solutions [7].

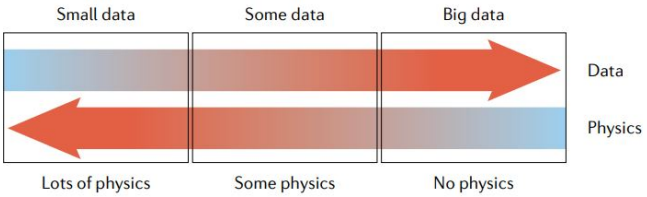

Fig. 4: Data and Physics Scenarios [7]

This is achieved by introducing a loss function that can consist of two main components: the residual loss [8] and the data loss [9].

- The residual loss ensures that the predictions made by the neural network are in alignment with the governing physical equations of the system. Specifically, it quantifies the discrepancy between the derivatives of the network's output (as computed through automatic differentiation) and the terms present in the differential equations that describe the system. This component of the loss function acts as a regularizer, guiding the network to learn solutions that are not only consistent with the data but also with the underlying physics.
- The data loss, on the other hand, is designed to measure the difference between the network's predictions and the actual observed data. This is typically computed as a mean squared error between the network's output and the known target values from the dataset. The data loss ensures that the network is effectively learning from the available data and is able to generalize well to new, unseen data points.

By combining these two components into a single loss function, PINNs are able to leverage the strengths of both data-driven learning and physics-based modeling. This hybrid approach allows PINNs to effectively interpolate and extrapolate in regions where data might be sparse, while still maintaining physical plausibility. It also provides a natural way to incorporate domain knowledge into the learning process, which can be especially valuable when dealing with complex systems where data is expensive or difficult to obtain [10]–[12].

Furthermore, the use of PINNs in this context is particularly advantageous as it allows for the incorporation of complex boundary and initial conditions directly into the training process. This is in contrast to traditional numerical methods, which often require separate handling of such conditions. With PINNs, these conditions are naturally embedded within the loss function, ensuring that the trained model inherently respects these critical constraints.

For the purpose of this research, we approached the problem in three stages as depicted in (Fig 4): first, we used conventional neural networks and data loss to model the problem. Secondly, As the fundamental PINN training formulation is unsupervised and does not require labeled data [13], such as results from other simulations, we considered the PINN

structure with the loss function being the residual loss only (Fig 2):

$$L_{res} = \hat{\phi}'_{NN} - \sqrt{2\left(M^2\left[1+\frac{\hat{\phi}_{NN}}{M^2}-\frac{\hat{\phi}^2_{NN}}{2M^4}+...-1\right]+1-\hat{\phi}_{NN}+\frac{\hat{\phi}^2_{NN}}{2}+...-1\right)}, \tag{5}$$

where $\hat{\phi}_{NN}$ is the prediction from the neural network.

This focus on the residual loss ensured that the model aligned with the known physics of lunar surface charging, providing a foundation that adheres to the underlying physical principles. We then refined the approach by considering the loss function as a combination of both residual loss and data discrepancy (Fig 3), and compared the results with the results of only data-driven network and fundamental PINN:

$$L_{data} = \frac{1}{N}\Sigma_{i=1}^{N}(\hat{\phi}_{NN,i} - \hat{\phi}_i)^2, \tag{6}$$

where N is the number of data points, $\hat{\phi}_i$ is the actual observed value. The overall loss function for the PINN can then be represented as:

$$L_{total} = \alpha L_{res} + \beta L_{data}, \tag{7}$$

where α and β are weighting factors that determine the relative importance of the residual loss and data loss, respectively.

This combined loss function aimed to achieve a balanced and accurate representation of the electric potential profile above the lunar surface, leveraging both the physical constraints and the available data.

IV. Results

In our endeavor to model the electric potential profile above the lunar surface, we initially explored the use of conventional neural networks. For this approach, the primary objective was to minimize the data loss, which measures the discrepancy between the network's predictions and the actual observed data. This data-centric approach served as our baseline, allowing us to compare the performance of conventional neural networks with that of more sophisticated models. The results of this initial exploration, as depicted in Figure 5, provided valuable insights into the capabilities and limitations of using a purely data-driven approach for this complex physical system.

We then employed neural networks comprising linear layers and non-linear activation functions. Given that the input data to the network consists of floating-point numbers representing continuous variables, the use of complex layers such as convolutional layers was deemed unnecessary. Instead, we focused on employing simple linear layers coupled with non-linear activation functions. This approach allowed us to approximate the underlying functions governing the electric potential profile effectively.

Our exploration included both shallow and deeper neural network architectures. Figure 6:(a) and Figure 6: (b) illustrate the predictions produced by these two distinct models. Interestingly, the results reveal that both the shallow and deeper networks yield similar predictions. This observation suggests that, for this specific problem, increasing the complexity of the network does not necessarily translate into significant improvements in predictive accuracy.

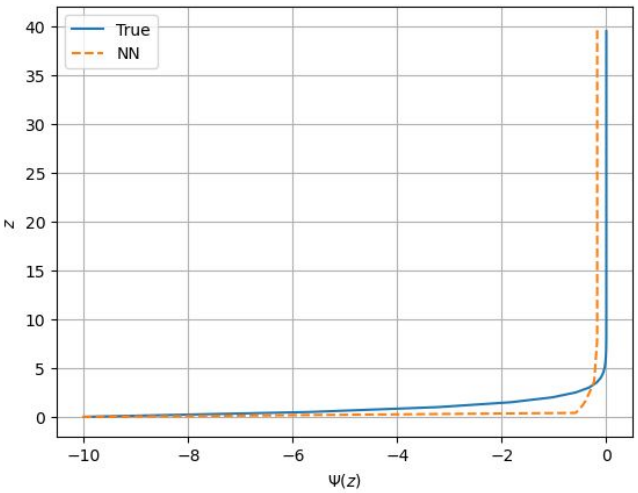

Fig. 5: Conventional (Data-Driven) Neural Network

This finding is particularly noteworthy as it challenges the common notion that deeper networks are always superior for complex tasks. In our case, it appears that a simpler, more computationally efficient architecture is capable of capturing the essential characteristics of the electric potential profile above the lunar surface.

A. Activation Functions

Activation functions are fundamental components in neural networks, introducing non-linearity into the model, which allows it to capture complex relationships in the data. Without activation functions, the neural network would simply be a linear regression model, which has limited expressiveness. Among the myriad of activation functions available, the sigmoid, rectified linear unit (ReLU), and hyperbolic tangent (tanh) stand out as the most popular and widely used in various deep learning architectures.

- The sigmoid function, characterized by its classic S-shaped curve, maps any real-valued number to a value between 0 and 1. This makes it particularly useful in scenarios like binary classification, where we need to predict probabilities. However, one of its drawbacks is the vanishing gradient problem, which can impede the network's learning in certain situations.
- The tanh function is quite similar to the sigmoid but maps its input to a range between -1 and 1. This zero-centered nature often makes it preferable over the sigmoid in hidden layers of neural networks.
- ReLU, with its simple design, has become one of the default choices for many deep learning models. It outputs the input directly if it's positive; otherwise, it outputs zero. Its efficiency and computational simplicity make it a favorite, though it's not without its challenges, such as

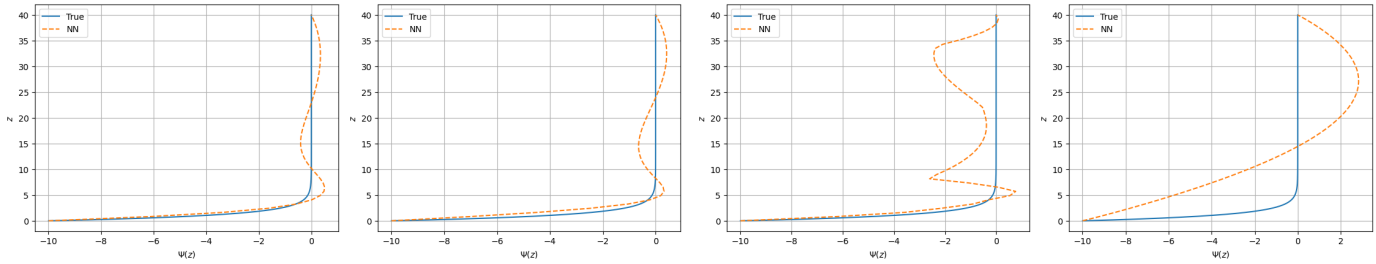

(a) Network depth: shallow, activation function: Sigmoid, Optimizer: Adam

(b) Network depth: deep, activation function: Sigmoid, Optimizer: Adam

(c) Network depth: shallow, activation function: ReLU, Optimizer: Adam

(d) Network depth: shallow, activation function: Sigmoid, Optimizer: SGD

Fig. 6: Comparison of networks

the "dying ReLU" problem where neurons can sometimes get stuck during training.

As depicted in Figure 6:(a) and Figure 6:(c), there is a noticeable difference in the performance of models using these two functions. The model employing the ReLU activation (Figure 6:(c)), produces predictions that are less smooth compared to the one using the sigmoid activation, evident from Figure 6:(a). The jagged nature of the predictions from the ReLU model might be due to its characteristic of outputting zero for any negative input, which can sometimes lead to less smooth approximations, especially when the model is trying to capture intricate patterns.

On the contrary, the sigmoid function, which maps its input to a value between 0 and 1, produces smoother predictions. This smoothness can be attributed to the sigmoid's S-shaped curve, which allows for a more gradual transition between values. In the context of our study, this smooth transition appears to be more aligned with the underlying patterns of the data.

Given these observations, it becomes evident that the sigmoid activation function is better suited for the specific task explored in this paper. While both activation functions have their merits and are widely used in various domains, the choice of function should be informed by the nature of the data and the specific requirements of the task at hand.

B. Optimizer

Optimizers are the driving force behind the learning process of neural networks. They navigate the vast parameter space to find the optimal values that reduce the error between the predicted and actual outputs. By adjusting the model's weights in response to the computed gradients, optimizers ensure that the network converges to a solution that best fits the data. Among the plethora of optimizers available, Adam [15] and SGD (Stochastic Gradient Descent) [16] are two of the most widely used. Each optimizer has its unique approach to updating the model's parameters, and their effectiveness can vary depending on the specific problem and dataset.

As illustrated in Figure 6:(a) and Figure 6:(d), there's a discernible difference in the performance of models trained using these two optimizers. The model trained with the

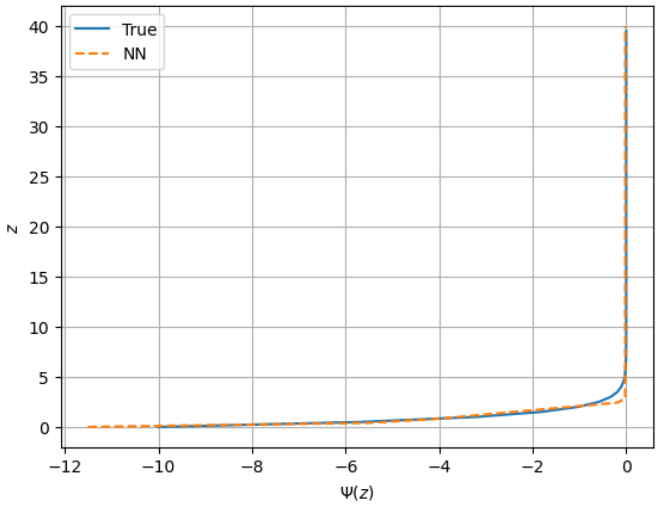

Fig. 7: Model performance using data loss and residual loss

Adam optimizer, as shown in Figure 6:(a), produces a more refined and smoother prediction curve. This can be attributed to Adam's adaptive learning rate mechanism, which adjusts the learning rate for each parameter based on the historical gradient information. This feature allows Adam to converge faster and often results in better generalization.

On the other hand, SGD (Figure 6:(d)), being a more traditional and straightforward optimization method, updates the model's parameters using a fixed learning rate. While it has been the go-to optimizer for many years, in certain scenarios, especially with complex models and non-convex loss surfaces, SGD might struggle to find the optimal solution as efficiently as some of its more modern counterparts like Adam.

Given the observed results, it's evident that the Adam optimizer offers a more effective optimization strategy for the task at hand. However, it's essential to note that the choice of optimizer should always be based on empirical evidence, as different tasks and datasets might favor different optimization techniques.

C. Loss Function

Initially, we adopted a physics-centric approach, where the loss function was purely based on the physics governing the lunar surface charging. This approach, while rooted in established physical principles, aimed to ensure that the model's predictions were in harmony with the known behaviors of the lunar surface in a plasma environment. The results from this model were promising, showcasing its capability to grasp and replicate the fundamental characteristics of the electric potential profile.

However, recognizing the potential of data-driven insights, we ventured into a more holistic approach. In this third model, we amalgamated the residual loss with a data loss component. This hybrid loss function was designed to strike a balance between the theoretical physics of the problem and the empirical evidence from the data. The outcome was enlightening. The model, guided by this combined loss function, demonstrated enhanced performance, capturing the intricacies of the electric potential profile with greater precision. It became evident that by synergizing the foundational knowledge of physics with the empirical patterns in the data, the model could achieve a more comprehensive and accurate depiction of lunar surface charging.

Figure 7 visually encapsulates the comparative performance of the two models. It underscores the value of integrating both physics-based reasoning and data-driven insights, highlighting the efficacy of the combined loss function in addressing the complex problem of lunar surface charging.

V. Conclusion

In conclusion, this research explored modeling the electric potential profile above the lunar surface using Physics-Informed Neural Networks. We evaluated three loss models: residual loss, data loss, and a hybrid approach combining both. The hybrid model, integrating data and residual loss, performed best by leveraging empirical data and adhering to physics principles for robust and precise modeling.

The study also compared shallow and deep neural networks, finding no significant advantage to deeper networks in this context, emphasizing the need to tailor architecture to specific problems. Moreover, we assessed various optimizers, highlighting their impact on performance.

These insights advance understanding of lunar surface charging and neural network modeling. By evaluating loss functions, network depth, and optimization techniques, we identified areas for further research and demonstrated the potential of neural networks in space environment modeling, bridging theory, and practical application.

As we advance, our focus will be on incorporating uncertainty into our models by varying M values within a specified range and evolving them to process both M and Z inputs. By doing so, we anticipate a richer, more comprehensive understanding of the system dynamics, leading to predictions that are not only accurate but also encompass the multifaceted nature of the lunar electric potential profile.

Acknowledgment

This research is sponsored by the National Science Foundation under Grant No. OAC-2306184 with the University of North Texas and Grant No. OAC-2330364 with the University of Kentucky. The experiments presented here were carried out with computing resources provided by the Center for High Performance Computing Research at Missouri University of Science and Technology through an NSF grant OAC-1919789 and Kummer Innovation and Entrepreneurship Doctoral Fellowship at Missouri University of Science and Technology and by the NSF grant CMMI-1954548.

References

[1] Freeman, J. W., and M. Ibrahim. "Lunar electric fields, surface potential and associated plasma sheaths." In Lunar Science Institute, Conference on Interactions of the Interplanetary Plasma with the Modern and Ancient Moon, Lake Geneva, Wis., Sept. 30-Oct. 4, 1974. The Moon, vol. 14, Sept. 1975, p. 103-114., vol. 14, pp. 103-114. 1975.

[2] Chandran, SB Rakesh, C. L. Veenas, L. R. Asitha, B. Parvathy, K. R. Rakhimol, A. Abraham, S. R. Rajesh, A. P. Sunitha, and G. Renuka. "Potential–Current characteristics of lunar surface at average solar wind conditions." Advances in Space Research 70, no. 2 (2022): 546-555.

[3] Han, Daoru, Xiaoming He, David Lund, and Xu Zhang. "PIFE-PIC: Parallel immersed finite element particle-in-cell for 3-D kinetic simulations of plasma-material interactions." SIAM Journal on Scientific Computing 43, no. 3 (2021): C235-C257.

[4] Wei, Xinpeng, Jianxun Zhao, Xiaoming He, Zhen Hu, Xiaoping Du, and Daoru Han. "Adaptive Kriging method for uncertainty quantification of the photoelectron sheath and dust levitation on the lunar surface." Journal of Verification, Validation and Uncertainty Quantification 6, no. 1 (2021): 011006.

[5] Chen, Francis F. Introduction to plasma physics and controlled fusion. Vol. 1. New York: Plenum press, 1984.

[6] Lagaris, Isaac E., Aristidis Likas, and Dimitrios I. Fotiadis. "Artificial neural networks for solving ordinary and partial differential equations." IEEE transactions on neural networks 9, no. 5 (1998): 987-1000.

[7] Karniadakis, G.E., Kevrekidis, I.G., Lu, L. et al. Physics-informed machine learning. Nat Rev Phys 3, 422–440 (2021). https://doi.org/10.1038/s42254-021-00314-5

[8] Cuomo, S., Di Cola, V.S., Giampaolo, F. et al. Scientific Machine Learning Through Physics–Informed Neural Networks: Where we are and What's Next. J Sci Comput 92, 88 (2022). https://doi.org/10.1007/s10915-022-01939-z

[9] Thakur, Sukirt, Maziar Raissi, Harsa Mitra, and Arezoo Ardekani. "Temporal Consistency Loss for Physics-Informed Neural Networks." arXiv preprint arXiv:2301.13262 (2023).

[10] Giampaolo, F., De Rosa, M., Qi, P. et al. Physics-informed neural networks approach for 1D and 2D Gray-Scott systems. Adv. Model. and Simul. in Eng. Sci. 9, 5 (2022). https://doi.org/10.1186/s40323-022-00219-7

[11] Raissi, Maziar, Paris Perdikaris, and George E. Karniadakis. "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations." Journal of Computational physics 378 (2019): 686-707.

[12] Lawal, Zaharaddeen Karami, Hayati Yassin, Daphne Teck Ching Lai, and Azam Che Idris. 2022. "Physics-Informed Neural Network (PINN) Evolution and Beyond: A Systematic Literature Review and Bibliometric Analysis" Big Data and Cognitive Computing 6, no. 4: 140. https://doi.org/10.3390/bdcc6040140

[13] Markidis, Stefano. "The old and the new: Can physics-informed deep-learning replace traditional linear solvers?." Frontiers in big Data 4 (2021): 669097.

[14] Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-yao Huang, Zhihui Li, Xiaojiang Chen, and Xin Wang. 2021. A Comprehensive Survey of Neural Architecture Search: Challenges and Solutions. ACM Comput. Surv. 54, 4, Article 76 (May 2022), 34 pages. https://doi.org/10.1145/3447582

[15] Kingma, Diederik P., and Jimmy Ba. "Adam: A method for stochastic optimization." arXiv preprint arXiv:1412.6980 (2014).

[16] Robbins, Herbert, and Sutton Monro. "A stochastic approximation method." The annals of mathematical statistics (1951): 400-407.