Optimal Energy Management of a Fast Charging Service Station with Physics-Informed Neural Networks

Francesco Liberati, Emanuele De Santis, Mohab M. H. Atanasious, Alessandro Di Giorgio
Department of Computer, Control, and Management Engineering Antonio Ruberti,
Via Ariosto 25, 00185 Rome, Italy
Email: {liberati, edesantis, atanasious, digiorgio}@diag.uniroma1.it

***Abstract*—In this paper, we investigate the use of Physics-Informed Neural Networks (PINNs) for the solution of an optimal control problem related with the optimal control of a stationary electric storage system (ESS) installed in a service station for plug-in electric vehicles (PEVs). The ESS is used to balance the PEVs charging power, in order to mitigate the impact on the grid, and keep low the power flow at the point of connection with the grid. The proposed PINN is trained in order to learn the optimality conditions of the optimal control problem so that, after training, it can provide the solution with no significant computation effort. This one represents a promising alternative to the analytical computation of the optimal control (which is possible only in very simple settings), and to the solution of the optimal control problem with numerical methods, which requires significant time in the more complex and realistic settings. Numerical simulations are presented to evaluate the effectiveness of the trained PINN in solving the optimal control problem.**

***Index Terms*—Physics-Informed Neural Networks, Optimal Control, PEV Service Station, Energy Storage, Two-Point Boundary Value Problem**

I. Introduction

Increasing the number of dedicated service stations to provide fast charging service to plug-in electric vehicles (PEVs) is essential to sustain the widespread of electromobility. Together with the improvement of battery capacity, and the development of battery swapping technology, ultra rapid fast charging is obviously one of the few means to solve the range anxiety issue and favor the massive adoption of PEVs. The first electric service stations have been recently already deployed across the globe (see, e.g., [1] and [2]). Given the need of providing PEV users with a very rapid charging service, the power charging levels are expected to be very high in service stations, when considering several PEVs simultaneously charging, easily reaching the megawatt levels in the near future. For this reason, it is convent to install an electric storage system (ESS) in the service area, to be controlled for the purpose of balancing the PEV charging power, alleviating the stress on the main grid to which the service station is connected. The ESS is expected to discharge power in the times of the day when there is an intense demand of charging power from the PEVs, and then to recharge at periods of low demand (e.g., at night). In this paper, we investigate the use of physics-informed neural networks (PINNs) to solve the problem of optimal control of the ESS in the service station. Traditionally, this optimal control problem has been tackled either analytically (by computing the exact optimal solution to the problem, when possible), or numerically, by applying methods such as shooting [3], or by leveraging on passivity [4]. Both traditional approaches have drawbacks. Indeed, the exact analytical solution can be found only in very simple settings (in the present case, when the constraint on the state and the powers are ignored), while it is not known in general for the complete formulation of the problem (the one which includes all the applicable constraints). On the other hand, the numerical approaches, which can be used to solve the fully constrained problem, can be very slow to compute the solution or difficult to set in order to avoid numerical issues.

In our previous paper [5], we computed the analytical optimal solution to the problem, considering a simplified setting in which the ESS state constraints are ignored. The analytical solution was computed simply by solving the necessary optimality conditions related to the optimal control problem, which were found by applying the Pontryagin maximum principle. These optimality conditions form together a so-called two-point boundary value problem (TPBVP), i.e., a set of two differential equations (one for the state of the storage and the second for the costate of the optimal control problem), and a set of two boundary condition, one written for the initial problem time (the initial condition of the storage), and the second written in the final problem time (the so-called transversality condition). In [5], no costraints were considered, which allowed one to solve the TPBVP in a closed form. In this paper instead, we train a PINN to learn the optimal solution.

In general, PINNs are neural networks which are trained using a loss function that embeds the differential equations

Corresponding author Mohab M. H. Atanasious (email: atanasious@diag.uniroma1.it).
F. Liberati and E. De Santis are co-first authors of the paper.
This work was supported by the SAPIENZA—ATENEO 2023 "Advanced Predictive and AI-Based Control Algorithms With Application to Logistic Management and eHealth" under Project RM123188F7C85F26.

that describe the physical process they should emulate [6]. In the present paper, the differential equations to learn instead represent (together with the two boundary conditions of the TPBVP) the optimality conditions associated with the optimal control problem to solve. In this way, after sufficient training, the optimal solution can be found simply by querying in real time the trained neural network, with negligible computation time. In the simulations, we compare the solution achieved by the trained PINN with the solution found with a traditional numerical shooting method (indeed, they are in general the ones used in the complex cased in which there is no closed-form solution). The idea of using PINNs to solve optimal control problems has recently been proposed in other fields of application. Several applications are found in the aerospace domain. For example, in [7], Schiassi et al. use PINNs to compute optimal manuvers for space vehicles, such as planar orbit transfers. They apply the theory of functional connections to ensure the analytical satisfaction of the boundary conditions of the TPBVP. The approach is refined in [8] to find the fuel-optimal trajectories for space manouvers. The PINN is named Pontryagin neural network since it learns the conditions arising from the Pontryagin maximum principle. In [9], Furfaro et al. apply the methodology to the solve a missile pitch-plane autopilot optimal control problem. In this case, the authors use the PINN to learn the optimal control policy of an infinite-horizon control problem, letting the PINN learn the Hamilton-Jacobi-Bellman partial differential equations that describe the optimal solutions. Again, the theory of functional connections is used. Since it learns the Hamilton-Jacobi-Bellman equation, the PINN is called Bellman neural network.

Beyond aerospace, PINNs have been recently proposed for solving optimal control problems in biology. For instance, in [10] PINNs are applied for the optimal control in bacterial cultivation, for the optimal control of mold and fungicide, and for design of optimal cancer chemotherapy strategies. In conclusion, PINNs are emerging as a powerful tool to derive fast solvers for optimal control problems, overcoming the traditional numerical based solving approaches, which often cannot be implemented in real-time applications due to excessive computation times.

In this paper, we propose to apply the machinery of PINNs to address a relevant optimal control problem in the smart grid domain, given by the optimal management of an ESS in a service station for PEV fast charging. To the best of our knowledge, this is one of the first applications in this domain. This paper is a first step, laying the foundations for future works in which the proposed solution will be expanded and tested, as explained in the following.

The rest of the paper is organized as follows: Section II describes the reference scenario and problem formulation, Section III briefly recalls PINNs and details the PINN-based solution, Section IV presents the simulation results. Finally, Section V concludes the paper and outlines the future works.

II. Reference Scenario and Problem Formulation

The infrastructure setup is shown in Fig. 1. The ESS and the charging stations of the service station are directly connected to the grid through their own power electronics. Let $p(t)$ denote the power flowing at time t at the point of connection (POC) with the grid, $u(t)$ the ESS charging/discharging power (positive if recharging), $w(t)$ the cumulative PEV charging power (non-negative), $x(t)$ the ESS level of energy. The power flowing at the POC is given by:

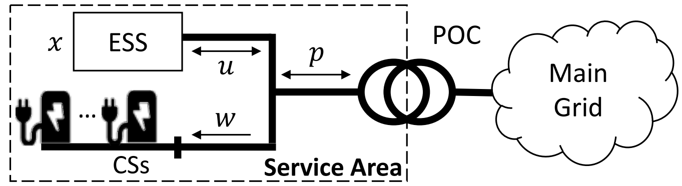

Fig. 1. Architecture of the service station.

$$p(t) = u(t) + w(t). \tag{1}$$

The ESS can be modelled as a lossy integrator:

$$\dot{x}(t) = \eta\left(u(t)\right)u(t), \tag{2}$$

where $\eta(u(t))$ is the efficiency function.

Finally, variables $p(t)$, $u(t)$ and $x(t)$ have lower and upper admissible limits:

$$u_{\min} \le u(t) \le u_{\max}, \tag{3}$$

$$p_{\min} \le p(t) \le p_{\max}, \tag{4}$$

$$x_{\min} \le x(t) \le x_{\max}. \tag{5}$$

The goal of the control problem is to control the ESS power $u(t)$ over a given time span (e.g., a day) in order to balance the PEV recharging power, so that the power profile at the POC is flattened, and the storage state of charge remains close to a reference value (e.g., 50% full charge), so to always keep good control margins.

The optimal control can be therefore found by minimizing the following target function:

$$J = \frac{s}{2}(x(t_f) - x_{\text{ref}})^2 + \int_{t_0}^{t_f} \left(\frac{q}{2}(x(t) - x_{\text{ref}})^2 + \frac{r}{2}p(t)^2\right) dt, \tag{6}$$

subject to (1)-(5), where t_0 is the initial time, t_f the final time, x_{ref} is the reference energy level, and q, r, and s are weighting parameters for the running and terminal costs.

After defining the Hamiltonian as:

$$\mathcal{H} := \frac{q}{2}(x(t) - x_{\text{ref}})^2 + \frac{r}{2}p(t)^2 + \lambda(t)\eta(u(t))u(t), \tag{7}$$

where $\lambda(t)$ is the co-state variable, the necessary optimality conditions can be found by applying the Pontryagin's minimum principle, as shown in [5], and, after including (1) in (7), they lead to the following TPBVP (when ESS losses are neglected - $\eta(u(t)) = 1$):

$$\dot{x}(t) = u(t), \tag{8}$$

$$\dot{\lambda}(t) = -q[x(t) - x_{\text{ref}}], \tag{9}$$

$$u(t) = -\frac{\lambda(t)}{r} - w(t), \tag{10}$$

with boundary conditions $x(t_0) = x_0$ and $\lambda(t_f) = s\left(x\left(t_f\right) - x_{\text{ref}}\right)$. This TPBVP can be solved in closed form when the box constraints on x, p and u are ignored, and is typically solved using numerical methods instead when they are retained. In the next section, we present the PINN-based approach to learn the optimal control solution.

III. PINN-based Optimal Control

PINNs learn the solution of a set of differential equations (typcically modelling a physical system) by incorporating them in the loss function used for training. In this work, we use a PINN to approximate the solution of the TPBVP presented in the previous section. As it is a first work in this application domain, we decided to ignore the box constraints and consider a lossless storage (i.e., $\eta(u(t)) = 1 \quad \forall u$), so to have benchmark results to be tested against in future works, where additional constraints on the control, on the state, and more accurate storage models are considered. The PINN is trained to predict the optimal state $x(t)$ and the co-state $\lambda(t)$, solution of the above TPBVP, over the time interval $[t_0, t_f]$. Specifically, the input of the PINN is given by the initial state x_0 (i.e., the state of the ESS at t_0), and by the time t at which we want to evaluate the optimal solution. The output is given by the optimal solution $x(t)$ and $\lambda(t)$ (solution of the TPBVP) predicted by the PINN. Once the optimal value of $\lambda(t)$ is obtained from the PINN, then the optimal control $u(t)$ to command to the ESS is given by (10). A scheme clarifying the use of the PINN to compute the storage optimal control is reported in Fig. 2. At every time t_0 when the optimal control to be applied to the storage over the future time span $[t_0, t_f]$ (e.g., from current time to end of the day) needs to be computed, the current initial state of the storage is measured, and a prediction of the future PEV demand $w(t)$ over $[t_0, t_f]$ is acquired. Then, the control for every time $t \in [t_0, t_f]$ is computed by feeding the trained PINN with the initial state, the time when to compute the solution, and the entire curve of w in $[t_0, t_f]$. The PINN returns the optimal costate $\lambda(t)$ at t, from which the optimal storage control to apply is computed immediately using (10), as illustrated in the figure. Notice that there is the question of how to pass the entire curve w to the PINN. In this paper, for simplicity we train the PINN for a given curve w. In general, to keep the number of the PINN inputs low, thus increasing the learning performances, a good strategy is to compute the Fourier expansion of w, and pass the coefficients of the expansion up to a given truncation order. Then training over different curves w can be performed by simply exploring the space of the Fourier coefficients. To let the PINN learn the differential equations and the boundary conditions of the TPBVP associated with the optimal control problem, two main types of terms are considered in the loss function used for training of the PINN.

- Residuals associated to the differential equations of the TPBVP: to enforce the dynamics $\dot{x} - u = 0$ and $\dot{\lambda} + q(x - x_{\text{ref}}) = 0$.
- Residuals associated to the boundary conditions of the TPBVP: To ensure $x(0) = x_0$ and $\lambda(t_f) = s(x(t_f) - x_{\text{ref}})$.

The total loss is a weighted sum of the above residual terms:

$$\mathcal{L} = \alpha_{r_1}\mathcal{L}_{r_1} + \alpha_{r_2}\mathcal{L}_{r_2} + \alpha_{\text{bc}_2}\mathcal{L}_{\text{bc}_2}, \tag{11}$$

where:

$$\mathcal{L}_{r_1} = \frac{1}{N}\sum_{i=1}^{N}\left[\dot{x}(t_i) + \frac{\lambda(t_i)}{r} - w(t_i)\right]^2, \tag{12}$$

$$\mathcal{L}_{r_2} = \frac{1}{N}\sum_{i=1}^{N}\left[\dot{\lambda}(t_i) + q(x(t_i) - x_{\text{ref}})\right]^2, \tag{13}$$

$$\mathcal{L}_{\text{bc}_1} = \frac{1}{N_{\text{bc}_1}}\sum_{i=1}^{N_{\text{bc}_1}}[x(t_0) - x_0]^2, \tag{14}$$

$$\mathcal{L}_{\text{bc}_2} = \frac{1}{N_{\text{bc}_2}}\sum_{i=1}^{N_{\text{bc}_2}}\left[\lambda(t_f) - s(x(t_f) - x_{\text{ref}})\right]^2. \tag{15}$$

Notice that the expression of the optimal control is directly written in the first loss term. Notice also that the derivatives present in the first two loss terms are the derivatives of the output of the PINN with respect to one of the inputs (t). During training, they are computed by differentiating through the PINN, a procedure called automatic differentiation [11]. Hence, all the terms in the loss function are inputs of the PINN, outputs of the PINN, derivatives of the outputs, or values assumed fixed (the weights of the target function and the reference state). During training, the loss is evaluated, and the parameters of the PINN are updated in order to reduce the loss. In detail, the residual loss of the differential equations is evaluated at N different collocation points, N different times $\{t_i\}_{i=1,\dots,N}$ between t_0 and t_f. The first two loss terms are the average of the residual of the differential terms at the collocation points. As for the last two loss terms, instead, the residual loss is evaluated at N_{bc_1} and N_{bc_2} different points, with temporal coordinate fixed respectively to t_0 and t_f, since they are the boundary condition residuals.

Once trained, the PINN provides a continuous approximation of the optimal values of $x(t)$ and $\lambda(t)$, from which the optimal control $u(t)$ is derived from the formula of the optimal control (10). In the next section, the simulation results are presented. The PINN solution is compared to the one obtained with the classic shooting methods.

IV. Simulations

The proposed PINN-based solution was implemented in Python on Google Colab, using the PyTorch library for neural network training, and the SciPy library for numerical integration. The PINN was trained using the Adam optimizer with a learning rate scheduler.

We simulate one day of operation. The simulated aggregated PEV demand $w(t)$ has two peaks, one around midday and another in the evening, reflecting a plausible PEV charging pattern (the curve $w(t)$ is displayed below in Fig. 5). The main simulation parameters are reported in Tab. I.

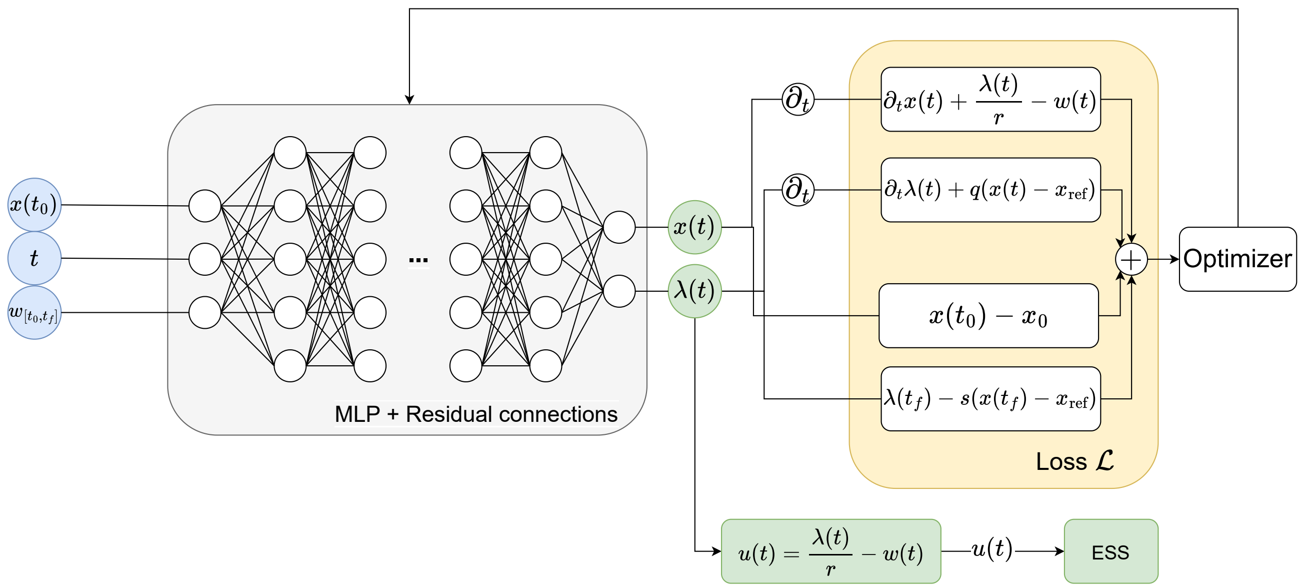

Fig. 2. Scheme to compute the optimal storage control with PINN.

TABLE I
MAIN SIMULATION PARAMETERS

Parameter	Value
Initial time (t_0)	0 hours
Final time (t_f)	24 hours
Reference energy (x_{ref})	500 kWh
Initial state range (x_0)	[400, 600] kWh
State cost weight (q)	1.0
Control cost weight (r)	10.0
Terminal cost weight (s)	10.0
Number of collocation points (N)	2000
Number of training epochs	100,000
Optimizer	AdamW
Learning rate	0.001
Learning rate scheduler	StepLR (step 10,000, gamma 0.5)
Initial loss weight (α_{r_1})	10^2
Initial loss weight (α_{r_2})	10^1
Initial loss weight (α_{bc_1})	10^2
Initial loss weight (α_{bc_2})	10^{-3}
Network architecture	5 layers, 200 neurons/layer
Activation function	Tanh
Residual connections	Yes

Neural network with residual connections is used to increase the computational power of the neural network itself. The mathematical formulation is described as follow:

$$\begin{aligned} z_1 &= f\left(W_1^T z_0 + b_1\right) \\ z_k &= f\left(W_k^T z_{k-1} + b_k\right) + z_{k-1}, \quad \forall k \in [2, l-1], \\ z_l &= f\left(W_l^T z_{l-1} + b_l\right) \end{aligned} \tag{16}$$

where l is the number of layers of the networks, $f(\cdot)$ is the activation function, W_k are the weights of the $k-th$ layer, z_0 is the input of the network, b_k is the bias of the $k-th$ layer and z_k is the output of the $k-th$ layer.

The PINN was trained for 100,000 epochs, with the initial energy x_0 randomly sampled between 400 and 600 kWh during training, for the PINN to be able to generalize across different values of the initial state (a wider range can be chosen as well). The training process took 20.87 minutes. Figure 3 shows the evolution of the training loss over epochs, including also the detail of each loss term (a moving average of the loss is plotted, for better visualization). In addition, the weights of the loss function are adapted dynamically across the training by giving a higher weight (importance) to the worst loss terms. This process is performed each 1000 epochs At the end of the training process, the PINN already shows very good convergence (it could be further improved with longer training). To evaluate the performance of the PINN and its ability to generalize over different values of the initial ESS state, we run simulations considering five different ESS initial states $x_0 \in \{400, 450, 500, 550, 600\}$ kWh. For these initial states, Fig. 4 plots the evolution of the ESS state and

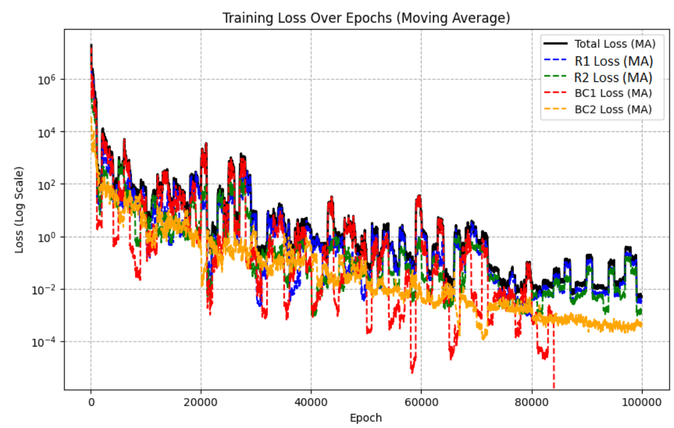

Fig. 3. Training loss over 100000 training epochs.

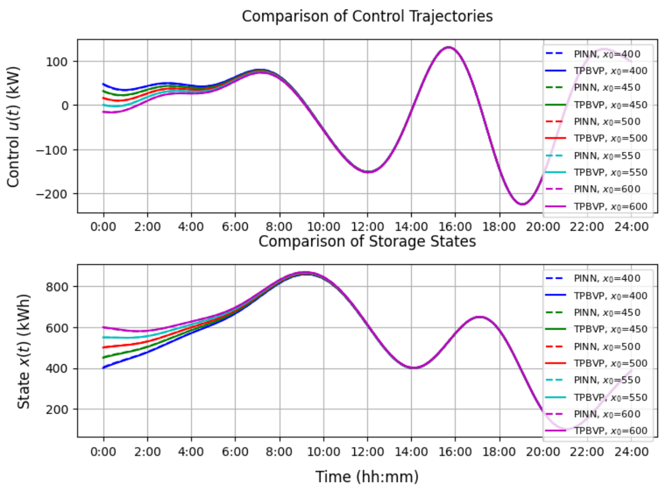

Fig. 4. Comparison of control trajectories $u(t)$ (top) and storage states $x(t)$ (bottom) between PINN and TPBVP shooting solution for different initial states x_0.

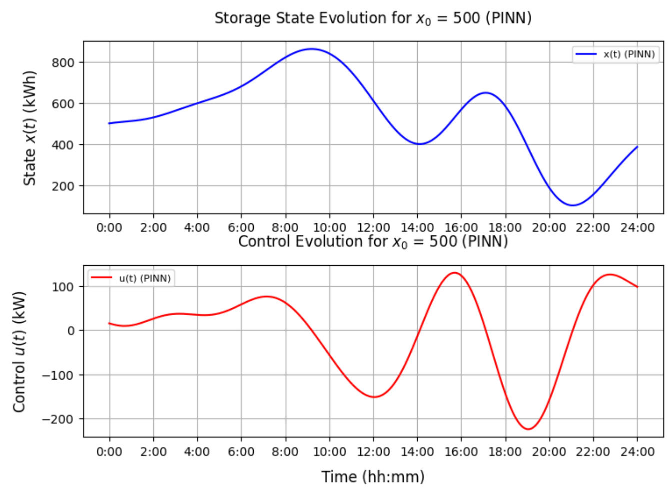

Fig. 6. State $x(t)$ (top) and control $u(t)$ (bottom) evolution from PINN for $x_0 = 500$ kWh.

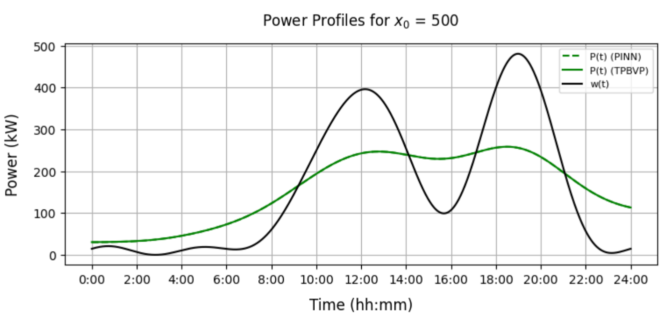

Fig. 5. Power profiles for $x_0 = 500$ kWh, comparing $P(t)$ from PINN and TPBVP shooting solution with the demand $w(t)$.

TABLE II
VALUE OF THE OBJECTIVE FUNCTION FOR DIFFERENT INITIAL STATES x_0

x_0	PINN	TPBVP	% Gap (PINN vs TPBVP)
400	4491238.11	4490515.07	0.0161
450	4461190.93	4463163.92	-0.0442
500	4443109.14	4443718.50	-0.0137
550	4432703.41	4432178.82	0.0119
600	4426110.81	4428544.87	-0.0550

ESS control, as computed by the trained PINN, and through a standard shooting method. The simulations confirm that the PINN is very accurate in computing the optimal solution of the ESS control problem, as the PINN and the TPBVP shooting solutions overlap. Next, we present the power flow results for one of the above five initial states. We take the initial state equal to the reference state, $x_0 = 500$ kWh. Figure 5 plots the power flow at the point of connection between the service station and the main grid, $p(t)$, and, for comparison, the aggregated demand from the PEVs, $w(t)$. The figure displays both the power profile computed using the optimal control from the PINN, and the one computed using the optimal control from the shooting method. The figure again confirms an excellent alignment between the result of the two methods (the PINN and the shooting method), and shows the effectiveness of the optimal control problem in flattening the power profile (the peaks of w are significantly reduced). Figure 6 presents the state $x(t)$ and control $u(t)$ evolution computed by the PINN. Finally, to provide an absolute measure of the performance achieved by the PINN, Tab. II presents these the comparison of the total cost achieved by the PINN solution (i.e., the value of the target function (6) attained by the solution), against the total cost achieved by the optimal control computed with the numerical solution of the TPBVP with the shooting method. The gap between the two is very low, confirming the good training of the PINN (in some cases, PINN performs better than shooting, due to the numerical nature of the latter).

V. CONCLUSIONS

This paper discussed the use of a PINN to solve the problem of optimally controlling an electrical storage in a service station for PEV fast charging. The PINN was employed to learn the state-costate differential equations and the boundary conditions characterizing the optimal solution, and derived from the Pontryagin maximum principle. Once trained, the PINN takes in input the value of the storage initial state, and the time for which we want to compute the optimal solution, and returns the optimal value of the costate at that time, which then can be used to compute the optimal storage control to apply. The advantage is that this process has extremely low computational time (almost instantaneous), as opposed to the solution of the problem with numerical methods, which can imply times not compatible with real time control. The simulation results confirm that the trained PINN returns a solution which is practically equivalent to the optimal one computed numerically.

Future works will regard the solution of the problem considering all the constraints on the storage state of charge and on all the power flows (which were neglected in this paper)), and considering as well losses in the storage (which is considered lossless in this paper). Also, while the PINN for this paper, assumes a fixed profile for the demand from PEVs (w), in future works the demand will be set as an input (for instance, by passing to the PINN the coefficients of a Fourier expansion), so that the optimal solution can be quickly re-computed every time the forecast of the future PEV demand changes. Moreover, specific learning-based demand forecasting mechanisms (as in [12]) can be investigated.

ACKNOWLEDGEMENT

The authors would like to acknowledge the use of Grok 3.beta [13] for writing and debugging the simulation code.

REFERENCES

[1] Shell. Shell opens its largest EV charging station in China. *Last accessed on 15/03/2024*. [Online]. Available: https://www.shell.com/energy-and-innovation/mobility/mobility-news/shell-opens-largest-ev-charging-station.html

[2] Hampton, S., Azzouz, L., Fawcett, T., Grunewald, P., Howey, D., Kumtepeli, V., Mould, T., Rose, T. Energy Superhub Oxford: Final report - Europe's most powerful electric vehicle charging hub heading to Oxford. *Last accessed on 15/03/2024*. [Online]. Available: https://energysuperhuboxford.org/

[3] M. A. Arefin, M. A. Nishu, M. N. Dhali, and M. H. Uddin, "Analysis of reliable solutions to the boundary value problems by using shooting method," *Mathematical Problems in Engineering*, vol. 2022, no. 1, p. 2895023, 2022.

[4] M. Mattioni and P. Borja, "Digital passivity-based control of underactuated mechanical systems," *Automatica*, vol. 173, p. 112096, Mar. 2025. [Online]. Available: http://dx.doi.org/10.1016/j.automatica.2024.112096

[5] F. Liberati, A. Di Giorgio, and G. Koch, "Optimal stochastic control of energy storage system based on pontryagin minimum principle for flattening pev fast charging in a service area," *IEEE Control Systems Letters*, vol. 6, pp. 247–252, 2022.

[6] M. Raissi, P. Perdikaris, and G. 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*, vol. 378, pp. 686–707, 2019. [Online]. Available: https://doi.org/10.1016/j.jcp.2018.10.045

[7] E. Schiassi, A. D'Ambrosio, K. Drozd, F. Curti, and R. Furfaro, "Physics-informed neural networks for optimal planar orbit transfers," *Journal of Spacecraft and Rockets*, vol. 59, no. 3, pp. 834–849, 2022.

[8] A. D'Ambrosio and R. Furfaro, "Learning fuel-optimal trajectories for space applications via pontryagin neural networks," *Aerospace*, vol. 11, no. 3, p. 228, 2024.

[9] R. Furfaro and A. D'Ambrosio, "Increasing autonomy of aerospace systems via pinn-based solutions of hjb equation," in *AIAA SCITECH 2024 Forum*. AIAA, 2024.

[10] W. Jiang, "Application of physics-informed neural networks in optimal control problems of biological models," in *2024 6th International Conference on Frontier Technologies of Information and Computer (ICFTIC)*, 2024, pp. 1617–1623.

[11] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, "Automatic differentiation in machine learning: a survey," *Journal of machine learning research*, vol. 18, no. 153, pp. 1–43, 2018.

[12] A. Pimpinella, F. D. Giusto, A. E. C. Redondi, L. Venturini, and A. Pavon, "Forecasting busy-hour downlink traffic in cellular networks," in *ICC 2022 - IEEE International Conference on Communications*. IEEE, May 2022, p. 4336–4341. [Online]. Available: http://dx.doi.org/10.1109/ICC45855.2022.9838982

[13] xAI, "Grok 3 beta," https://x.ai/news/grok-3, February 2025.