POLITECNICO DI TORINO
Repository ISTITUZIONALE

Data-driven control of wave energy systems using random forests and deep neural networks

Original
Data-driven control of wave energy systems using random forests and deep neural networks / Pasta, E., Carapellese, F., Faedo, N., Brandimarte, P.. - In: APPLIED OCEAN RESEARCH. - ISSN 0141-1187. - 140:(2023), pp. 1-15. [10.1016/j.apor.2023.103749]

Availability:
This version is available at: 11583/2982585 since: 2023-09-29T07:57:15Z

Publisher:
Elsevier

Published
DOI:10.1016/j.apor.2023.103749

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

26 November 2026

Applied Ocean Research 140 (2023) 103749



Contents lists available at ScienceDirect

Applied Ocean Research

journal homepage: www.elsevier.com/locate/apor



Research paper

Data-driven control of wave energy systems using random forests and deep neural networks

Edoardo Pasta [a,*], Fabio Carapellese [a], Nicolás Faedo [a], Paolo Brandimarte [b]

[a] Marine Offshore Renewable Energy Lab., Department of Mechanical and Aerospace Engineering (DIMEAS), Politecnico di Torino, Corso Duca degli Abruzzi 24, Turin (TO), 10129, Italy

[b] Department of Mathematical Sciences (DISMA), Politecnico di Torino, Corso Duca degli Abruzzi 24, Turin (TO), 10129, Italy

ARTICLE INFO

Keywords:
Wave energy
Optimal control
Data-driven control
Random forests
Regression trees
Deep neural networks

ABSTRACT

In the field of sustainable energy and alternatives to fossil fuels, wave energy is generating an increasing interest due to its untapped potential. However, the levelised cost of energy of wave energy systems is still not able to compete with other renewable technologies, mainly due to high costs associated with their conversion process. In this context, the development of *energy-maximising* control strategies plays an important role towards the economic viability of wave energy technology, by optimising the overall energy conversion, hence contributing towards minimising the associated cost of energy. State-of-the-art control systems for wave energy converters are mostly *model-based*, exploiting *control-oriented* models of the device to compute the applied control actions with a limited computational burden. Nevertheless, these representations of the system are simplified, and often based upon unrealistic assumptions, such as small motion around the zero equilibrium position, which are inherently invalidated during device operations and that can lead to large uncertainties, resulting in suboptimal power absorption. For these reasons, in this paper, a purely *data-driven* control strategy is developed. This strategy exploits *random forests* (RFs) and *deep neural networks* (DNNs) to gradually learn from real experiences towards an optimal proportional–integral control action. These structures are used as surrogate models (built upon the data coming from past experiences) to converge to the optimal control parameters in a surrogate-optimisation-like manner. To manage the exploration and exploitation needs of controllers based on this approach, a learning strategy is developed and presented. Some considerations are made on the choice of the input features of the surrogate structures, which deeply affect the control strategy learning results. To assess the performances of both the control and learning strategies, one year of operations has been simulated under control settings guided by the proposed data-driven approach, showing also the potential capabilities that the adoption of RFs and DNNs has in learning, even in sea conditions with a limited number of occurrences.

1. Introduction

In the context of an increasing interest on sustainable solutions to the problem of rising energy demand, *wave energy* represents a promising resource, due to its mainly untapped potential (Terrero González et al., 2021). Devices aimed at absorbing energy from ocean waves are called *wave energy converters* (WECs). These attempt to convert the wave motion into electrical energy, exploiting different conversion principles (Czech and Bauer, 2012). In contrast to other offshore energy technologies, wave energy has not yet still reached technological maturity nor a commercialisation stage, mainly due to the absence of a definitive conversion concept (Guo and Ringwood, 2021a; Trueworthy and DuPont, 2020), able to extract energy efficiently in a wide variety of operating conditions. Among the challenges that must be overcome in the development of wave energy systems, the definition of an efficient control strategy plays a role of paramount importance (Ringwood, 2020). Suitable control strategies are responsible of maximising the energy absorption (Ringwood et al., 2014), continuously adjusting the WEC action to the wave resource whose characteristics are continuously changing (Reguero et al., 2015).

Most of the control strategies developed for WECs are *model-based*, since they exploit a previously formulated *control-oriented model* of the system to compute the associated control action. This model allows the achievement of several functionalities, like wave force estimation (Peña-Sanchez et al., 2020b), and consequent prediction (Peña-Sanchez et al., 2020a), enabling the adoption of popular predictive control approaches (Faedo et al., 2017). However, the formulation of control-oriented models often relies on several simplifications, in order

* Corresponding author.
E-mail address: edoardo.pasta@polito.it (E. Pasta).

https://doi.org/10.1016/j.apor.2023.103749
Received 8 April 2023; Received in revised form 18 September 2023; Accepted 20 September 2023
0141-1187/© 2023 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/).

to reduce the computational and representational burdens associated with their adoption, and to allow real-time operations on physical hardware. These simplifications are often associated with small motion (*i.e.*, linearised) assumptions about the zero equilibrium, neglecting the complexity associated with WEC nonlinearities (Penalba et al., 2017), like viscosity effects (Wang et al., 2018), internal machinery dynamics (Penalba and Ringwood, 2019; Xuhui et al., 2019), or mooring contributions (Palm et al., 2016; Paduano et al., 2020). Even if linear hydrodynamic modelling can be quite accurate (see, *e.g.*, Stansby et al. (2017)), the application of control actions computed on the basis of linearised models (especially whenever internal mechanical systems are involved) invalidates, in general, the linearity assumptions adopted to compute such simplified descriptions, resulting in the so-called *WEC modelling paradox* (Windt et al., 2021): Energy-maximising control indeed tends to increase the WEC motion, moving far from the linearised operational space. This phenomenon leads to non-representative synthesis of even simple control structures (Carapellese et al., 2022a). In the attempt of solving this issues, some novel approaches have been proposed to formulate WEC nonlinear models in a control-oriented fashion (Faedo et al., 2022b). However, model-based strategies are still affected by uncertainties given by the mismatch between real and modelled device (Farajvand et al., 2023), and, if this uncertainty is not taken into account during control synthesis, the resulting controllers could perform suboptimally. Moreover, WEC devices are supposed to operate for long time spans. During this period, because of ageing effects or biofouling phenomena on the hull, the WEC dynamics can change, and consequently the difference between original model and real device can increase, further leading to suboptimal performance. These considerations motivate the adoption of *data-driven* control strategies, which base their synthesis on real data, and not on WEC models (*i.e.*, they are *model-free*).

1.1. Contributions and paper positioning

Considering the discussion provided in Section 1, this paper presents a data-driven control strategy able to gradually learn from the experiences that the real WEC faces over time. To achieve this, two different approaches are presented. The first approach adopts *deep neural networks* (DNN) to describe the knowledge gathered by the device, and to continuously learn towards optimal actions. The second approach employs a *random forest* (RF) to achieve the same objective. To the best of authors' knowledge, this is the first time RFs have been exploited in the field of WEC control. Both DNN and RF are adopted for their capabilities to synthesise all the information provided by real experiences in different conditions, in an attempt to find a generalised and accurate description of the effects that the controller has on device productivity in different sea conditions. This relationship is modelled by DNNs and RFs, also called in this context *metamodels*, purely on the basis of data. It is important to highlight that the metamodels are not attempting to model in any way the device dynamics, but only its consequences on the previously mentioned relationship. In the proposed control strategy, the metamodel is used to optimise the control action for any given sea state. Furthermore, the paper introduces also a learning strategy, which suitably balances the two goals of data-driven control strategies (Pasta et al., 2022): *exploration* and *exploitation*. Finally, with the purpose of fully exploit the capabilities of DNNs and RFs, an analysis related to the selection of the *features* employed as metamodel inputs is performed, highlighting the influence that proper feature selection has on final learning performance.

To assess the performance of the proposed control strategy, an extensive numerical test is performed. More in particular, the controller is employed to maximise the power absorption of a *Pendulum Wave Energy Converter* (PeWEC) (Pozzi et al., 2018; Sirigu et al., 2020). This has been modelled with a fully-nonlinear dynamical representation for its internal mechanical system, and is subject to irregular waves taken from a validated sea states dataset from a candidate deployment site in the Mediterranean Sea obtained from ERA5 global reanalysis (Hersbach et al., 2020). In order to assess also the capabilities of the controller during the learning period, simulations of one year length are performed, considering a concatenation of wave conditions which reflects the real deployment site data. To the best of authors' knowledge, this is the first time a learning-based controller has been tested so extensively in wave energy field, demonstrating the learning capabilities in a complete set of waves which includes even sea states with low occurrences (standard analysis are done only on few conditions, without considering in the learning evaluation the sea state changes). It is important to highlight that the proposed controller design and analysis aim at more general applications that go beyond the control of PeWEC (which effectively represents only a single case study), and that the outlined procedure can be straightforwardly applied to other WECs.

1.2. Paper structure

The remainder of the paper is organised as follows. In Section 2, the basic working principles of wave energy conversion are presented, with a particular attention to PeWEC, and to the characteristics of the model adopted in the presented analysis. Section 3 is aimed at presenting the general formulation of the WEC *optimal control problem* (OCP), with a brief description of the issues that characterise this type of control application. In Section 4, the proposed data-driven control strategy is described, highlighting the role of the learning strategy in the algorithm. In Section 5, the concept of metamodel is detailed, and the basic notions related to DDNs and RFs and their training are presented (in Section 5.1 and Section 5.2, respectively). Considerations on the role of possible additional input features are formulated in Section 5.3. Section 6 is devoted to the comparison, in terms of productivity performance, of the proposed strategy (considering both DNN and RF metamodels, with and without additional input features), with a model-based counterpart. In Section 6.1, the year long simulation setup is presented, while Section 6.2 details the corresponding numerical results. Finally, in Section 7, conclusions are drawn on the learning capabilities of the proposed control strategy, providing a number of considerations on the choice of type of metamodel and input features.

2. WEC modelling

Under the umbrella of wave energy converters, devices based on different working principles have been proposed (Guo and Ringwood, 2021b). Among them, it is possible to make a distinction between *point absorbers* (Guo et al., 2022), *oscillating water columns* (Rosati et al., 2022), *terminators* (Salter, 1974), and *attenuators* (Yemm et al., 2012). Within the most popular solutions, point absorbers comprise a moored hull, whose motion is activated by action of the surrounding wave field. The energy is thus extracted by one or more device motions by means of the so-called *power take-off* (PTO) system, aimed at properly acting on the device in the attempt of absorbing the wave energy. The equation of motion, describing such type of devices, can be formulated (Falnes, 2002) as[1]:

$$M\ddot{x} = f_r(\dot{x}) + f_{hr}(x) + f_{exc}(\eta) + f_{add}(x, \dot{x}, \ddot{x}, \eta) - f_u, \tag{1}$$

where $M \in \mathbb{R}^{n\times n}$ is the generalised inertia matrix for the n degrees-of-freedom (DoFs) device, $x : \mathbb{R}^+ \to \mathbb{R}^n$ is the generalised displacement vector (which takes into account the displacement in all the DoFs), $\eta : \mathbb{R}^+ \to \mathbb{R}$ is the wave elevation, $f_r : \mathbb{R}^n \to \mathbb{R}^n$ represents radiation effects,[2] $f_{hr} : \mathbb{R}^n \to \mathbb{R}^n$ describes the hydrostatic restoring

[1] From now on, the dependence on time t is dropped when clear from the context.

[2] In the simulation model employed in this study (Eqs. (1) and (3)), the radiation terms follow the standard Cummins' equation formulation, *i.e.* these are modelled in terms of an associated convolution operator (Cummins, 1962).

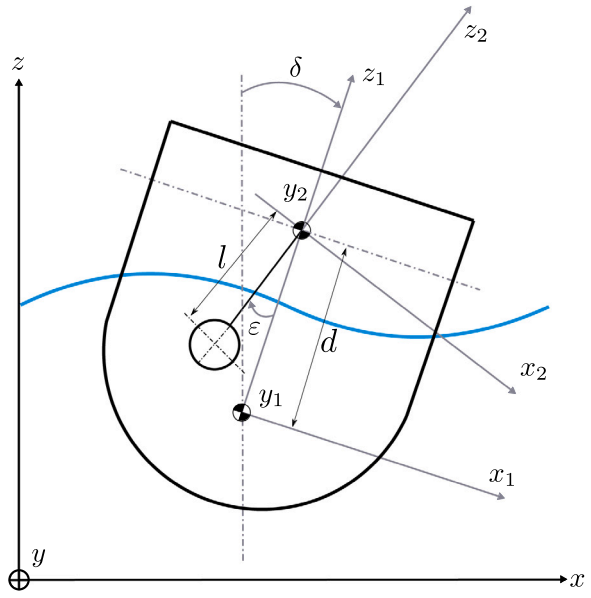

Fig. 1. Simplified schematics of the PeWEC device.

forces, $f_{exc} : \mathbb{R} \rightarrow \mathbb{R}^n$ is the excitation force contribution, while $f_{add} : \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}^n$ includes any additional effects, arising from mooring systems, mechanical transmissions, reaction forces, and PTO dynamics. Finally, the term $f_u : \mathbb{R} \rightarrow \mathbb{R}^n$ represents the control action applied to the WEC. The additional contributions in f_{add} can include nonlinear relationships. It is important to highlight also that the terms in Eq. (1), apart from the dependence from displacements, velocities, and accelerations, depend upon two external terms, which are the wave elevation η, and the control contribution f_u. For this reason, every controller aimed at maximising the absorbed energy must properly take into account the wave contribution through an estimation process (Peña-Sanchez et al., 2020b) or considering statistical synthetic information of the wave signal (Pasta et al., 2022).

In this paper the PeWEC technology is considered as a case study. PeWEC is a pitching-resonant WEC, consisting of a floating hull, moored to the seabed by a catenary mooring system (Pozzi et al., 2017; Paduano et al., 2021). Enclosed inside the hull, a pendulum system is installed and hinged to the PTO. The waves facing the hull generate a pitching motion, and the consequently induced pendulum swinging movement is damped by the PTO actuation system to extract energy. A simplified schematic of the PeWEC device is shown in Fig. 1.

In case of moorings designed to self-align with the wave direction, the floater DoFs can be reduced to surge, heave, and pitch. In this way, the PeWEC DoFs vector becomes:

$$x = \left[x_f, z_f, \delta_f, \epsilon\right]^\mathsf{T} = \left[x_{\text{tot},f}, \epsilon\right]^\mathsf{T}, \tag{2}$$

where x_f, z_f, δ_f, and ϵ are the surge, heave, pitch displacements, and the pendulum swinging angle, respectively, while $x_{\text{tot},f} = [x, z, \delta]^\mathsf{T}$. The floater motion follows the equation:

$$M_f \ddot{x}_{\text{tot},f}(t) = f_r(\dot{x}) + f_{hr}(x) + f_{exc}(\eta) + f_{pend}\left(x(t), \dot{x}(t), \ddot{x}(t), f_{u,p}\right), \tag{3}$$

where $M_f \in \mathbb{R}^{3\times3}$ is the inertia matrix of the PeWEC floater, while $f_{pend} : \mathbb{R}^3 \times \mathbb{R}^3 \times \mathbb{R}^3 \times \mathbb{R} \rightarrow \mathbb{R}^3$ are the reaction forces that the moving pendulum discharges on the hull, and $f_{u,p}$ is the PTO control action on the pendulum system. As it can be noticed, the control action f_u is not directly applied to any floater DoF, while its effects are part of the pendulum reaction forces. In this way, the overall dynamics of the WEC are controlled by means of the actuated pendulum motion. These reaction forces are a consequence of the pendulum dynamics, which can

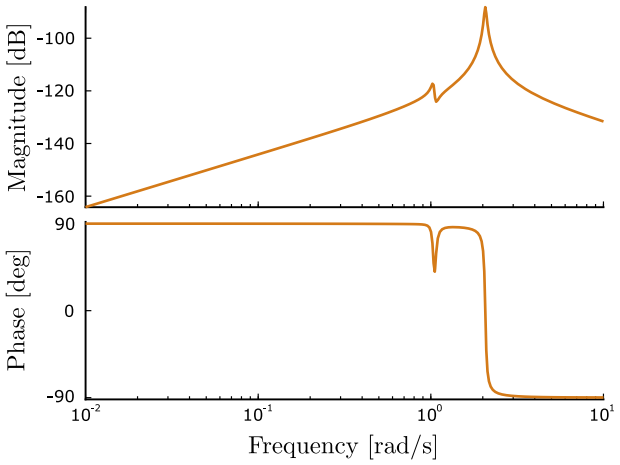

Fig. 2. Bode plot associated to the dynamics of the considered PeWEC system.

be described in terms of the following equation (Pozzi et al., 2018):

$$\begin{aligned}&\left(I_y + ml^2\right)\ddot{\varepsilon} - ml\cos(\delta_f + \varepsilon)\ddot{x}_f + ml\sin(\delta_f + \varepsilon)\ddot{z}_f \\ &+ \left(I_y + ml^2 - mdl\cos(\varepsilon)\right)\ddot{\delta}_f - mdl\sin(\varepsilon)\dot{\delta}_f^2 \\ &+ mgl\sin(\delta_f + \varepsilon) + f_{u,p} = 0,\end{aligned} \tag{4}$$

where I_y and m are the inertia around the pendulum baricentric axis, and its mass, respectively, l and d are the pendulum length and the distance between its hinge and the floater centre of mass, while g is the gravity acceleration constant. The pendulum motion is responsible of the reactions on the hull $f_{pend} = [f_{pend,x}, f_{pend,z}, f_{pend,\delta}]^\mathsf{T}$, which are exerted through its hinge and are formulated as:

$$\begin{aligned}f_{pend,x} =& - md\cos(\delta_f)\ddot{\delta}_f - ml\cos(\delta_f + \varepsilon)\left(\ddot{\delta}_f + \ddot{\varepsilon}\right) \\ &+ md\sin(\delta_f)\dot{\delta}_f^2 - ml\sin(\delta_f + \varepsilon)\left(\dot{\delta}_f + \dot{\varepsilon}\right)^2, \\ f_{pend,z} =& md\sin(\delta_f)\ddot{\delta}_f - ml\sin(\delta_f + \varepsilon)\left(\ddot{\delta}_f + \ddot{\varepsilon}\right) \\ &+ md\cos(\delta_f)\dot{\delta}_f^2 - ml\cos(\delta_f + \varepsilon)\left(\dot{\delta}_f + \dot{\varepsilon}\right)^2, \\ f_{pend,\delta} =& f_{pend,x}d\cos(\delta_f) - f_{pend,z}d\sin(\delta_f) + f_{u,p}.\end{aligned} \tag{5}$$

As it is possible to notice from Eqs. (4)–(5), the mechanical terms related to the pendulum dynamics are characterised by several non-linearities. To provide a model convenient for model-based control synthesis, the system of equations describing PeWEC is linearised under the assumptions of small oscillations about $\delta_f \approx 0$ and $\varepsilon \approx 0$ (Gioia et al., 2022). This assumption enables the adoption of linearised equations, and consequently a transfer function representation. As the main aim of these models is that of optimal control synthesis, the main input–output relationship that has to be studied is the one related to the axis of conversion (Faedo et al., 2022a), *i.e.*, from the control action ($f_{u,p}$ in the considered PeWEC case), to the related velocity ($\dot{\varepsilon}$ in PeWEC), as explained in Section 3 (in particular in Eqs. (6)–(8)). A Bode plot of such relationship, adopting the linearised equations of the considered PeWEC, is reported in Fig. 2. As it is possible to observe, the system dynamics are characterised by two main resonance conditions, one related to the hydrodynamics of the hull in pitch (first mode at 1.05 rad/s), and a second one at higher frequencies (second mode at 2.07 rad/s), which is the result of the pendulum dynamics. The represented transfer function in Fig. 2 is the one adopted in this paper to synthesise the model-based benchmark controller, following an impedance-matching approach (Faedo et al., 2022a).

3. WEC optimal control problem

In wave energy conversion, the control strategy is responsible for the proper selection of the action that has to be applied by the PTO

to maximise the power absorption. Moreover, control systems are used to suitably adjust to the marine environment, which is continuously changing, hence preventing the system from falling in possibly dangerous operational conditions. To achieve these objectives, the process of control synthesis requires the solution of an *optimal control problem* (OCP) of the *energy-maximisation* type. More specifically, if mechanical energy is considered, the performance function $\mathcal{J}$ to be maximised is the average absorbed power P_{abs} over a certain time interval $\mathcal{T} = [a, b]$, which can be formulated as [3]:

$$\mathcal{J}\left(f_u\right) = \frac{1}{T}\int_{\mathcal{T}} f_u(t)^{\mathsf{T}}\dot{x}(t)dt, \tag{6}$$

where $T = b - a$. Hence, in the general case, the WEC OCP can be fully written as[4]

$$\begin{aligned} f_u^{\text{opt}} &= \arg\max_{f_u}\ \mathcal{J}\left(f_u\right), \\ &\text{s.t.:} \\ &\text{WEC dynamics (1).} \end{aligned} \tag{7}$$

In the PeWEC case, where a single actuator (and consequent absorption DoF) is available at the pendulum hinge, Eq. (6) is rewritten as:

$$\mathcal{J}\left(f_{u,p}\right) = \frac{1}{T}\int_{\mathcal{T}} f_{u,p}(t)\dot{\epsilon}(t)dt. \tag{8}$$

Several approaches have been adopted to develop strategies able to suitably solve the WEC OCP. A classical approach is the one called *impedance-matching* (IM) (Ringwood et al., 2014; Faedo et al., 2020), based on an equivalent circuit of the excited WEC system model. Other model-based approaches have been applied to solve the WEC OCP. Among them, the most popular are *model predictive control* (Li and Belmont, 2014; Faedo et al., 2017), *pseudo-spectral* control (Garcia-Violini and Ringwood, 2021), *moment-based* control (Faedo et al., 2018, 2021), *linear quadratic Gaussian* (LQG) (Scruggs et al., 2013; Scruggs and Nie, 2015), and *LiTe-Con* (Garcia-Violini et al., 2020; Carapellese et al., 2022b). Even if some attempts have been made to 'robustify' these strategies (*e.g.*, in Faedo et al. (2019), Zhang and Li (2023)), model-based approaches are inherently sensitive to the uncertainty associated to the mismatch between the control-oriented model and the real system (Ringwood et al., 2020; Pasta et al., 2023), with respect to hydrodynamics (Farajvand et al., 2021). This possible drawback of model-based control motivates the investigation of data-driven approaches (Pasta et al., 2022; Moens de Hase et al., 2021). Among these, some of the most popular ones within the wave energy field are *extremum seeking* (Garcia-Rosa et al., 2012; Parrinello et al., 2020), *reinforcement learning* (Anderlini et al., 2016, 2018), and *maximum power point tracking* (Amon et al., 2012; Xiao et al., 2016). Some attempts at employing other strategies based on the concept of 'control by learning' have been made in Anderlini et al. (2017), exploiting neural networks, and in Gioia et al. (2022), Shi et al. (2019), employing Gaussian Process Regression. These latter approaches are also defined *surrogate optimisation-like* (Pasta et al., 2022), because of their similarity with surrogate-based optimisation algorithms in the field of global optimisation of computationally expensive functions (Forrester and Keane, 2009). In this context, this paper presents a surrogate optimisation-like WEC control strategy, which employs DNNs and RFs in the process, and adopts a predefined learning strategy to properly balance the trade-off between exploration and exploitation.

[3] The definition in (6) is presented in the general case with n available controllable actuators, one for every DoF.

[4] Other formulations of the OCP can be found, with this latter subject also to motion or control action constraints (Bacelli and Ringwood, 2013).

Table 1
Learning strategy design parameters.

Variable	Description
ΔT_e	Occurrence scatter T_e discretisation bin
ΔH_s	Occurrence scatter H_s discretisation bin.
T_{ss}	Sea state re-evaluation period
T_{min}	Minimum exploration time
$T_{eval.}$	$\mathcal{J}$ evaluation period
$T_{updt.}$	Metamodel update period
$\alpha_{gr.}$	$\epsilon_{gr.}$ slope
$\beta_{gr.}$	$\epsilon_{gr.}$ heat parameter
$\gamma_{gr.}$	$\epsilon_{gr.}$ centre horizontal position
$\delta_{gr.}$	Minimum exploration $\epsilon_{gr.}$ parameter

4. Proposed data-driven strategy

As discussed in Section 3, a data-driven control strategy based on an analogy with surrogate optimisation algorithms is proposed in this paper. In particular, the presented strategy is used to learn, in each sea state condition $S_{T_e,H_s} = \left[T_e, H_s\right]^{\mathsf{T}}$ (thus defined by the energetic period T_e, and significant height H_s), the optimal parameters of a reactive control law:

$$f_{u,p}(\theta, \epsilon, \dot{\epsilon}) = \theta^{\mathsf{T}}\begin{bmatrix}\dot{\epsilon}\\ \epsilon\end{bmatrix}, \tag{9}$$

with control law parameters $\theta = \left[\theta_1, \theta_2\right]^{\mathsf{T}} \in \mathbb{R}^2$, *i.e.*, PTO damping and stiffness, respectively. It is important to notice that the same parametrisation is often adopted in the IM approach to interpolate the optimal response of the system at a given excitation frequency. For this reason, in the performance comparison presented in Section 6, the controller employed as a model-based benchmark (*i.e.* reference) is synthesised by means of IM interpolation (Faedo et al., 2022a; Ringwood et al., 2014), which is indeed the standard practice in WEC control applications (Falnes, 2002). This choice is performed to maintain the same control parametrisation (Eq. (9)) in both the data-driven and model-based control strategies, providing a level playing field in terms of controller complexity. In the proposed data-driven approach (whose main design parameters are presented in Table 1), two main steps are present and continuously iterated:

1. The creation (and consequent periodical update) of a *metamodel* $\hat{f}_{\mathcal{M}} : \mathbb{R}^2 \times \mathbb{R}^2 \to \mathbb{R}$, $(S_{T_e,H_s}, \theta) \mapsto \hat{f}_{\mathcal{M}}(S_{T_e,H_s}, \theta)$, aimed at approximating the (unknown) relationship $\mathcal{J}_{\mathcal{M}} : \mathbb{R}^2 \times \mathbb{R}^2 \to \mathbb{R}$, $(S_{T_e,H_s}, \theta) \mapsto \mathcal{J}_{\mathcal{M}}(S_{T_e,H_s}, \theta)$. This latter tries to approximate the objective function $\mathcal{J} = P_{abs}$ (average absorbed power by the device), which, considering the control action defined as in Eq. (9) and explicitly accounting for the contribution of waves through the sea state, is consequently re-parameterised in terms of S_{T_e,H_s}, and the applied control parameters θ. The approximation $\hat{f}_{\mathcal{M}}$ obtained with the metamodel is created (and periodically updated every $T_{updt.}$) on the basis of the stored $\mathcal{J}_{past}$, S_{past}, and θ_{past}, which contain the past values of absorbed power, sea states, and applied control parameters, respectively.[5]
2. The choice of the control parameters θ to be applied in the sea state S_{T_e,H_s}. This is done balancing the actions of exploration (needed, especially in the initial stages, to build a suitable approximation $\hat{f}_{\mathcal{M}}$), and exploitation (to maximise the outcome in terms of absorbed energy). This trade-off, and consequent learning strategy, is managed by means of the *greediness function* $\epsilon_{gr.} : \mathbb{R} \to \mathbb{R}$, $N_{exp.} \mapsto \epsilon_{gr.}(N_{exp.})$, which, as a function of the number of experiences faced in a certain sea state $N_{exp.} : \mathbb{R}^2 \to \mathbb{R}$, $S_{T_e,H_s} \mapsto N_{exp.}(S_{T_e,H_s})$, defines the probability of applying an exploration action.

[5] Considered n experienced actions, $\mathcal{J}_{past}$, S_{past}, and θ_{past} are build and ordered such as the ith row of each of these matrices represents the variables associated to the ith experience, *i.e.*, $\mathcal{J}_{past} \in \mathbb{R}^{n\times 2}$, $S_{past} \in \mathbb{R}^{n\times 2}$, and $\theta_{past} \in \mathbb{R}^{n\times 2}$.

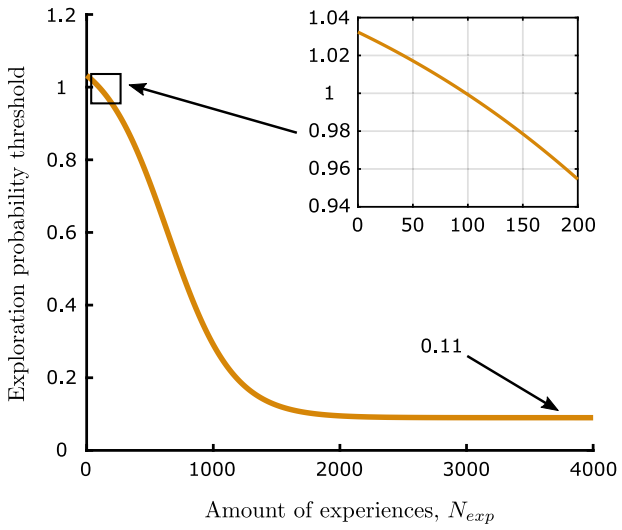

Fig. 3. Logistic sigmoid function adopted to define the greediness function $\epsilon_{gr.}$.

As it has been stressed above, apart from the metamodel approximation function, two main functions strongly influence the behaviour of the control strategy in terms of choices of control parameters. The first one is $N_{exp.}(S_{T_e,H_s})$, which counts the experiences that the strategy itself faces in a certain area of the sea state scatter. In particular, since $N_{exp.}$ plays a central role in the definition of the choices made by the greedy function, and the metamodel approximation $\hat{f}_{\mathcal{M}}(S_{T_e,H_s},\theta)$ is unique for all sea states (*i.e.*, a single function is built for all S_{T_e,H_s}, thus merging the knowledge $\hat{f}_{\mathcal{M}}$ has in contiguous sea state conditions, and exploiting consequently such information, in contrast to, *e.g.*, Gioia et al. (2022), Shi et al. (2019)), it would not be useful to count separately the experiences for each S_{T_e,H_s}. Instead, the domain of these two variables is divided in non-overlapping subsets considering bins of width ΔT_e and ΔH_s, respectively, and the number of experiences is counted for each cell of the resulting grid for S_{T_e,H_s}. In this way, ΔT_e and ΔH_s are design parameters of $N_{exp.}(S_{T_e,H_s})$, which returns the amount of past experiences in the sea state subset to which S_{T_e,H_s} belongs.

The other function of paramount importance in the learning process is $\epsilon_{gr.}(N_{exp.})$, which describes the probability of exploration that must be adopted in a sea state cell, considering the amount of past experiences around this wave condition. This greediness function is designed as a logistic sigmoid function (Cramer, 2004), given by the equation:

$$\epsilon_{gr.}(N_{exp.}) = -\frac{\alpha_{gr.}}{1+e^{-\beta_{gr.}(N_{exp.}-\gamma_{gr.})}} + \delta_{gr.}, \tag{10}$$

where $\alpha_{gr.}$ modulates the height of the sigmoid, $\beta_{gr.}$ is the so-called 'heat parameter', which changes the slope of the sigmoid (the higher $\beta_{gr.}$ is, the more the function tends to a Heaviside step function), $\gamma_{gr.}$ defines the horizontal position of the central point of the logistic function, while $\delta_{gr.}$ controls the exploration ratio value at $N_{exp.} \to \infty$. A representation of such function is reported in Fig. 3 (with $\alpha_{gr.} = 1.02$, $\beta_{gr.} = 3.9 \cdot 10^{-3}$, $\gamma_{gr.} = 740$, and $\delta_{gr.} = 1.11$, as in the numerical case study in Section 6). As it can be noticed, the design of such logistic function (10) allows to guarantee a certain minimum amount of exploration (*e.g.*, the 100 experiences required to reach a probability lower than 1), and to enforce a persistent degree of exploration on the long run, to allow the adaptation of the metamodel to the changes that could affect the system dynamics due to ageing (if properly coupled with a forgetting process in the management of past experiences $\mathcal{J}_{past}$, S_{past}, and θ_{past}).

A schematic representation of the control algorithm general workflow is shown in Fig. 4.[6] After an initialisation stage, in which the parameters related to the past experiences are set up, the strategy chooses, in the measured sea state condition, if an action aimed at exploring or exploiting has to be pursued. It is important to notice that, in the very first stages after deployment (when the current time t is less than a suitably defined minimum time T_{min}), the only allowed action is exploration. Whenever an exploration action is applied, the parameters $\theta \in \Theta$ are chosen *quasi-randomly* (following a Sobol's sequence (Bratley and Fox, 1988)) from a uniform distribution U_{Θ}, where $\Theta = [\underline{\theta}, \bar{\theta}]$ represents the limited set over which θ could be chosen (with $\underline{\theta} = [\underline{\theta}_1, \underline{\theta}_2]^{\mathsf{T}}$ lower bounds, and $\bar{\theta} = [\bar{\theta}_1, \bar{\theta}_2]^{\mathsf{T}}$ upper bounds), *i.e.*,:

$$\theta = \text{Sobol}(\mathsf{U}_{\Theta}). \tag{11}$$

The choice of a Sobol's sequence to define the quasi-random action is motivated by the benefits that a low-discrepancy exploration potentially has in building a metamodel with a 'well-distributed' training set (Tuffin, 2004). Once the control is defined, $\mathcal{J}$ is evaluated. As mentioned in Section 4, in this paper, the adopted performance function is the average absorbed power over the time a single θ configuration is applied. This time length $T_{eval.}$, in irregular wave conditions, has a strong influence over the variability of the measured absorbed power (Merigaud and Ringwood, 2018; Pasta et al., 2021), and consequently on the data adopted by the metamodel. We decided to consider a period of 10 min in the evaluation of $\mathcal{J}$ (*i.e.*, $T_{eval.} = 10$ minutes). After the control action is applied, and the performance function has been evaluated, $\mathcal{J}_{past}$, S_{past}, and θ_{past} are updated. Here, it is possible to define a maximum amount of experiences (each one represented by a row of $\mathcal{J}_{past}$, S_{past}, and θ_{past}), after which, the strategy starts to forget the eldest ones. In the assessments presented in Section 6, this value has been set to 100 experiences for each sea state condition. As it can be seen in Fig. 4, the defined control parameters are not applied for the entire (supposed) time length of the faced sea state (which is re-measured every T_{ss}). A shorter evaluation time can be considered, and a time counter variable t_{ss} is adopted to check whenever the time between a sea state measurement and the subsequent elapses. In this way, during the same sea state time window T_{ss}, multiple experiences can be achieved. Whenever instead, after the evaluation of $\epsilon_{gr.}$, an exploitation action has to be chosen, the control parameters are defined as:

$$\theta^{\star} = \underset{\theta\in\Theta}{\arg\max}\, \hat{f}_{\mathcal{M}}(S_{T_e,H_s},\theta), \tag{12}$$

where, in this way, $\theta^{\star}$ corresponds to the optimal value of θ found by maximising the approximated value $\hat{f}_{\mathcal{M}}$ given by the metamodel. The optimisation process is performed by means of a genetic algorithm (Thakur et al., 2014). It is important to highlight that Eq. (12) defines only the first exploitation action adopted in the time window between two sea state evaluations, as it can be appreciated by the algorithm work-flow scheme in Fig. 4. In fact, after the first exploitation action, for the remaining time within the measured sea state, an 'exploration around $\theta^{\star}$' is performed. This results in an exploration action over the bounded set $\Theta^{\star} = [(1-\frac{\epsilon}{2})\theta^{\star}, (1+\frac{\epsilon}{2})\theta^{\star}]$ which contracts as the amount of experiences grows:

$$\theta = \text{Sobol}(\mathsf{U}_{\Theta^{\star}}). \tag{13}$$

As it can be seen in Eq. (13), the same assumptions made for the explorative actions in Eq. (11), are here employed. Moreover, the set

[6] From there on, the notation $\text{Sobol}(\mathsf{D}_{\mathcal{A}})$ is adopted to refer to a value chosen from a distribution D over the set $\mathcal{A}$ and employing a Sobol's sequence quasi-random choice (Tuffin, 2004). Considering the variable $\alpha \in \mathbb{R}^n$, the notation $\mathcal{A} = [\underline{\alpha}, \bar{\alpha}]$ has been employed to refer to a bounded box set defined by the set of constraints $\alpha \geq \underline{\alpha}$, and $\alpha \leq \bar{\alpha}$, where $\underline{\alpha}$ and $\bar{\alpha}$ are lower and upper bounds of α, respectively.

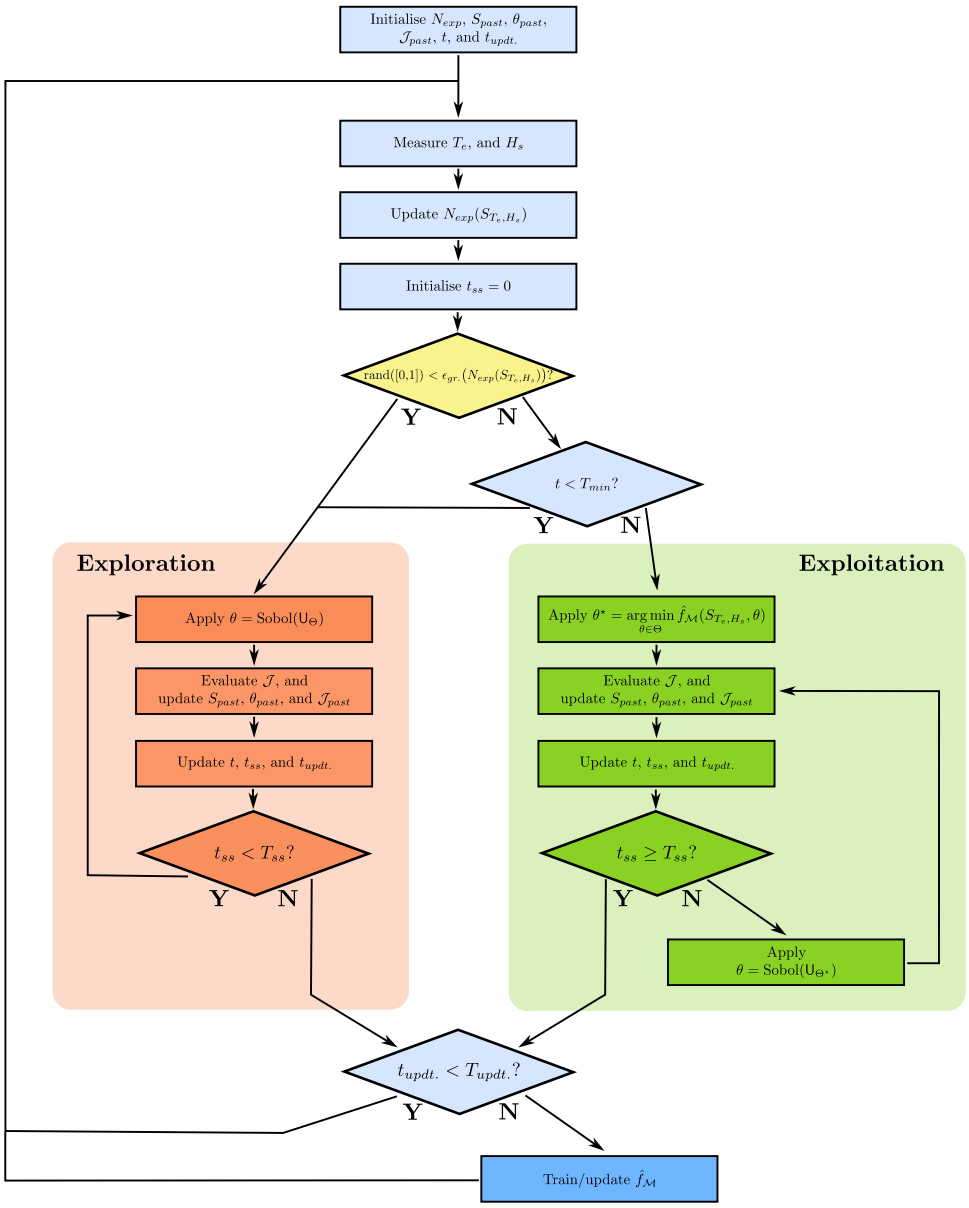

Fig. 4. Schematic representation of the work-flow of the proposed control algorithm. Following the choice of the action to be adopted (yellow node in the algorithm scheme) by means of the greediness function, the control parameters are defined either in an exploration (orange branch) or exploitation (green branch) manner, respectively. When a sufficient amount of time $T_{updt.}$ has passed, the metamodel is updated accordingly.

over which the control can be chosen is centred in the previously found $\theta^\star$, and its extent depends upon the value of $\epsilon_{gr.}(N_{exp.})$. This is done in the attempt of having an higher exploration of the set of parameters around the optimum, whenever the amount of experiences upon which the metamodel adopted to compute $\theta^\star$ is trained is low (to better approximate the area around the supposed optimum).

5. Metamodel structures

As stressed in Section 4, one of the key elements in the proposed control strategy is the metamodel. The definition of the structure adopted to approximate the function to optimise is of paramount importance for the convergence of the algorithm to the effective optima, and for the full exploitation of the information given by the collected data. In practice, a structure is used to map the input features $\mathcal{X} \in \mathbb{R}^{n_\mathcal{X}}$ into the approximated value $\hat{f}_\mathcal{M}$. To do that, in this paper, two different structures are adopted: DNN and RF. Sections 5.1 and 5.2 are devoted to provide a brief introduction to the formulation of these two types of models, and their training process. The mapping $(S_{T_e,H_s}, \theta) \mapsto \mathcal{X} = g_\mathcal{X}(S_{T_e,H_s}, \theta)$ between sea state, control parameters, and the corresponding input features is described in Section 5.3, together with its relevance and relationship with the metamodel structures.

5.1. Deep neural networks

DNNs are computational models with a structure inspired by biological considerations on animal brains. They are constituted by a combination of lower level entities, *neurons*, connected by weights, and are able to transform a set of inputs, generating a corresponding set of outputs, modified, eventually, through a non-linear relationship. The DNN goal is to approximate a target function. For this reason, the DNN itself can be seen as a map $f_{\text{DNN}}(\mathcal{X}, \gamma_{w,b})$, whose output y_{DNN} is aimed at approximating the target function. DNNs used in an approximation context are said to have an optimal configuration of their parameters $\gamma_{w,b}$ (weights, and possible biases) once they are able to reproduce with minimum error the shape of the target function. The process of

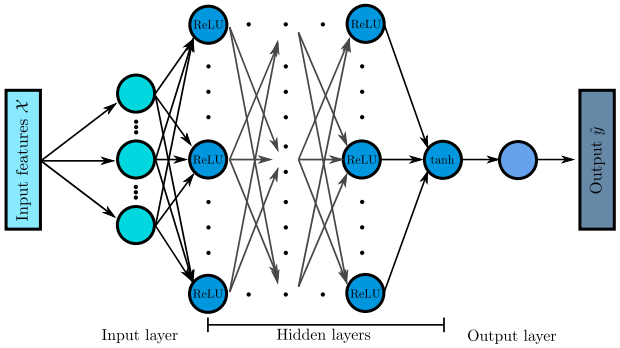

Fig. 5. Schematics of an example of DNN structure of feed-forward type. Architecture and activation functions (*Rectified Linear Unit*, ReLU, *hyperbolic tangent*, tanh) reflect the ones adopted in this study.

optimisation of such parameters is called training of the network. A graphical representation of an example of DNN structure is given in Fig. 5.

As it is possible to see from such representation, from a macroscopic point of view, neural networks consist of different layers. The input layer is that containing the neurons related to the input features $\mathcal{X}$. It is made by a number of neurons equal to the number of inputs. The output layer, instead, is related to the targets (or outputs). It comprises a number of neurons that correspond to the number of outputs. The output neurons are aimed at finalising the relationship connecting previous layers and inputs, shaping the outputs into a function as close as possible to the desired one. Eventually, it is possible to incorporate other layers, different from those devoted to inputs and outputs, which act as a link between these two layers. These are called hidden layers. The main characteristic of hidden layer neurons is that their input is the output of another neuron, and, at the same time, their output is the input for another neuron entity. The DNN fundamental entities (*i.e.*, the neurons) internally implements the relationship:

$$\begin{cases} d_i = w_{i-1}^{\mathsf{T}} \hat{y}_{\mathrm{DNN},i-1} + b_{i-1}, \\ \hat{y}_{\mathrm{DNN},i} = g_{act,i}(d_i), \end{cases} \tag{14}$$

where $d_i \in \mathbb{R}^{n_{\mathrm{DNN},i}}$ represents the partial output of the ith layer of the DNN (with $n_{\mathrm{DNN},i}$ number of neurons of the ith layer), $w_i \in \mathbb{R}^{n_{\mathrm{DNN},i-1} \times n_{\mathrm{DNN},i}}$ the ith layer weight matrix, $b_i \in \mathbb{R}^{n_{\mathrm{DNN},i}}$ represents the biases[7] at ith layer, $\hat{y}_{\mathrm{DNN},i} \in \mathbb{R}^{n_{\mathrm{DNN},i}}$ is the output of ith layer, and the function $g_{act,i} : \mathbb{R}^{n_{\mathrm{DNN},i}} \to \mathbb{R}^{n_{\mathrm{DNN},i}}$ is the activation function of the ith layer (here, for the sake of simplicity of exposition, it is assumed that all neurons in a layer have the same activation function as in standard feed-forward DNNs). The choice of the activation function determines the type of nonlinearity that is implemented by the DNN itself, hence being a design parameter.

In this study, a neural network with 5 hidden layers are adopted. Among these layers, the first 4 have 32 neurons each, and the last has only one neuron. Regarding the activation functions, the first 4 hidden layers adopt the Rectified Linear Unit (ReLU) activation function, while the last hidden layer implements a hyperbolic tangent (tanh) function. The output layer linearly scales the resulting output. The training is carried out adopting a mini-batch approach coupled with the *Adaptive Moment Estimation* (Adam) optimiser (Kingma and Lei Ba, 2015), using a mini-batch size of 64, to increase the generalisation capabilities of the DNN metamodel (Keskar et al., 2017). The hyperparameters defining the DNN structure (number of layers and neurons, batch size), have been defined through an extensive analysis of the performance of a trained network with an set of data coming from sample pilot simulations, in the attempt to minimise the approximation error, while reducing the network complexity. In order to deal properly with the different scales that the input features $\mathcal{X}$ can have, the input layer implements a *z-score* normalisation (Milligan and Cooper, 1988). During the training process, the available data are divided in training, validation, and test set (with percentages of 70%, 20%, and 10% of the full available dataset, respectively).

[7] The weights w_i and the biases b_i, together constitute the network parameters γ_i for ith layer. All parameters of all the layers together, constitute the network parameters $\gamma_{w,b}$.

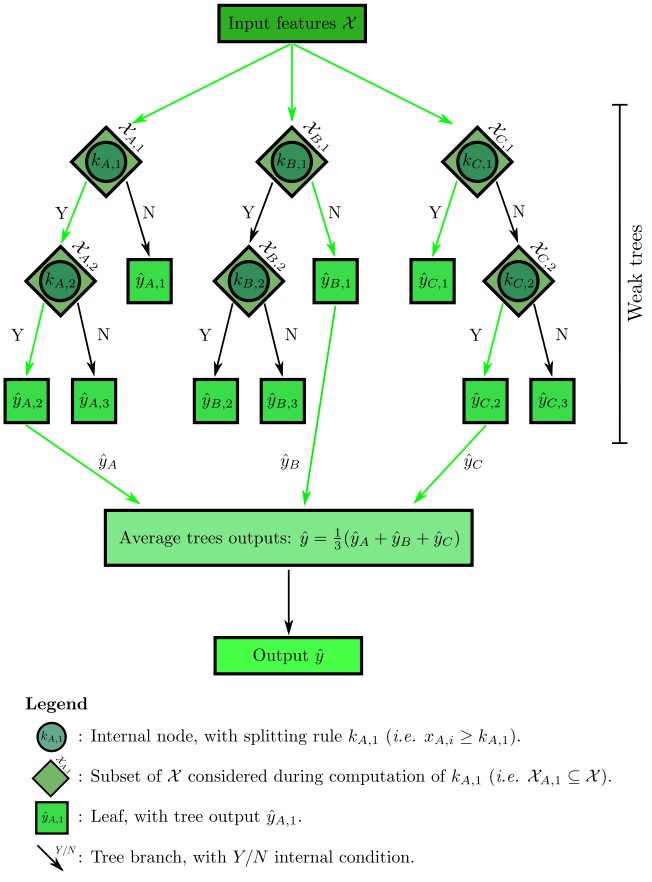

Fig. 6. Schematics of an example of regression trees constituting a RF.

5.2. Random forests

RFs are a particular version of the *bootstrap aggregation* technique, which employs a large ensemble of not correlated *regression trees*, and then averages them to obtain an estimated prediction function (Hastie et al., 2009, Chapter 15). The fundamental elements, *i.e.*, the regression trees, are structures which, starting from the input features, generate a series of branches, creating different paths dependent on the input values that all terminate into different target outputs (Breiman et al., 2017). An example of aggregation of regression trees, aggregated to constitute a RF, is shown in Fig. 6. As it can be noticed, the single regression tree comprises two main elements: leaves and internal nodes. The leaves (also called terminal nodes) partition the output space in regions, and their output value is the average of the data employed in the training stage that fall in that part of the output space. The conditions that split the output space are defined by the internal nodes. Each of them is defined in terms of an inequality (*i.e.*, the splitting rule), whose validity guides the 'output flow' in the proper leaf direction. The splitting rule is formulated as:

$$x_{A,i} \geq k_{A,i}, \tag{15}$$

where $x_{A,i}$ is the element of $\mathcal{X}$ considered in the ith internal node of the Ath tree of the forest, while $k_{A,i}$ is the parameter defining the splitting rule of such node.

The training stage of RFs is the combination of the concepts of bootstrap aggregation and non-correlationship of the internal regressors (*i.e.*, the trees). Each tree is trained separately, following a top-down greedy recursive binary splitting for the definition of the splitting rules. During the training of a single tree, the first step is the random definition of training datasets. This is done by randomly sampling with replacement from the full training dataset (*i.e.*, bootstrap sampling), creating a training dataset which has less data and, possibly, less information about the process to be approximated. For this reason, the single trees built on the basis of such sampled training datasets are also called *weak trees* (or weak learners). Training a regression tree consists of the definition of the splitting rules, which define the tree branches. Due to the high complexity that such choices can have, a single optimisation of the entire tree structure is not possible. For this reason, each node is optimised separately, starting from the first splitting definition and progressing towards the leaves. In RFs, the optimisation of the splitting rule of each node is done following three steps:

1. Random definition of the pool of elements of $\mathcal{X}$ that are candidates in the definition of the splitting rule. Adopting a random approach makes the trees obtained in this way uncorrelated (Hastie et al., 2009, Chapter 15), avoiding the same features to be considered in the first branches in all the trees (which are the more significant ones). Considering, for example, the ith node of the Ath tree of a RF with input features $\mathcal{X}$, the candidate input features $\mathcal{X}_{A,i}$ considered in the computation of the splitting rule $x_{A,i} \geq k_{A,i}$ are usually composed by a number of features $n_{\mathcal{X}_{tree}}$ (Hastie et al., 2009, Chapter 15):

$$n_{\mathcal{X}_{tree}} = \text{ceil}(\sqrt{n_{\mathcal{X}}}), \tag{16}$$

   where $n_{\mathcal{X}}$ is the amount of features in $\mathcal{X}$, and $n_{\mathcal{X}_{tree}}$ is rounded up to the next integer by the ceiling function.
2. Having selected the set of candidate features $\mathcal{X}_{A,i}$ in the definition of the splitting rule, elements $x_{A,i}$ and $k_{A,i}$ are chosen to minimise the mean squared error that the tree would have on a validation dataset, with the such node definition.
3. The process is repeated adding new nodes, until each leaf describes an output space which includes a number of training data below a certain threshold. At that point, the regression tree is fully trained.

In this study, random forests made up of 50 trees are employed, together with a maximum leaf dimension of 10 data points. Similarly to the DNNs case, these hyper-parameters have been chosen after an extensive analysis of the performance of a trained RF with a set of sample pilot data, trying to minimise the approximation error, while reducing the forest complexity. The available dataset has been divided into training and validation sets, with a ratio of 70% and 30%, respectively, to maintain the same training set percentage employed with DNNs. As the uncorrelationship of the single trees of the forest is given by the random choice of the candidate selection, and the number of considered features is regulated by Eq. (16), the definition of the input features (and, possibly, of additional ones) is of paramount importance in the random forest performance.

5.3. Choice of the input features

As discussed in Sections 5.1 and 5.2, the input features $\mathcal{X}$ are inputs fed to the metamodels, upon which the approximation $\hat{f}_{\mathcal{M}}(S_{T_e,H_s}, \theta)$ of the target function is obtained. The mapping $g_{\mathcal{X}}$ linking S_{T_e,H_s} and θ with $\mathcal{X}$ (as introduced in Section 5) is employed to add additional features to the ones available in the training of the metamodels employed by the control strategy. This has different benefits. The first one is related to the nonlinearities that can be included by such operation in the final metamodel structure. Indeed, through $g_{\mathcal{X}}$, is possible to explicitly add additional nonlinear terms inside the structure (*e.g.*, nonlinear functions or products between the inputs), which, in the end, can consequently increase the degree of complexity the resulting metamodel has, potentially enhancing its capabilities in describing the unknown function $\mathcal{J}_{\mathcal{M}}$ (Kuhn and Johnson, 2019; Heaton, 2016). Moreover, as discussed in Section 5.2, RFs strongly benefit from the presence of additional available inputs, since these increase the number of features available to uncorrelate the weak trees (the bigger the pool of available information, the more different and 'informed' are the single weak trees constituting the RF). Indeed, the number of input features employed by each single weak tree $n_{\mathcal{X}_{tree}}$, as shown in Eq. (16), depends upon the number of forest input features $n_{\mathcal{X}}$. Similarly, DNNs performance benefits from feature construction and the consequent additional features availability (Piramuthu, 1996; Piramuthu et al., 1998).

In this study, we compare the effect of adding additional input features has on the capability of approximating $\mathcal{J}_{\mathcal{M}}$, and in the consequent energy production. To do so, two different input feature definitions are employed[8]:

$$\mathcal{X}_{\text{no Add. Feat}}(S_{T_e,H_s}, \theta) = \begin{bmatrix} T_e \\ H_s \\ \theta_1 \\ \theta_2 \end{bmatrix}, \tag{17}$$

$$\mathcal{X}_{\text{with Add. Feat}}(S_{T_e,H_s}, \theta) = \begin{bmatrix} T_e \\ H_s \\ \theta_1 \\ \theta_2 \\ T_e^2 \\ H_s^2 \\ \theta_1^2 \\ \theta_2^2 \\ \log_{10}(T_e)\log_{10}(H_s) \\ \log_{10}(T_e)\log_{10}(\theta_1) \\ \log_{10}(T_e)\log_{10}(-\theta_2) \\ \log_{10}(H_s)\log_{10}(\theta_1) \\ \log_{10}(H_s)\log_{10}(-\theta_2) \\ \log_{10}(\theta_1)\log_{10}(-\theta_2) \end{bmatrix}. \tag{18}$$

As it can be seen in Eq. (17)–(18) the number of features $n_{\mathcal{X}}$ considered changes from 4 to 14. Moreover, the additional features present in $\mathcal{X}_{\text{with Add. Feat}}$ include the square of each single feature, and add information on the product between these, by considering the combinations the product between logarithms. This is done to avoid an excessive contribution of θ on the product, by scaling everything logarithmically. Indeed, the values adopted for the control parameters span[9] from $\underline{\theta} = [0.2 \cdot 10^5, -9 \cdot 10^5]^\mathsf{T}$ to $\bar{\theta} = [1.2 \cdot 10^5, -200]^\mathsf{T}$, and, for this reason, have a different order of magnitude with respect to T_e and H_s, whose values go from 0.5 s and 0 m, to 10 s and 5 m respectively. Inside the logarithms, θ_2 has been considered changed in sign, since it assumes only negative values (this term is used to reduce the inherent stiffness of the system due to the presence of the pendulum).

[8] From now on, the metamodels employing $\mathcal{X} = \mathcal{X}_{\text{no Add. Feat}}$ are called 'no Add. Feat', while those adopting $\mathcal{X} = \mathcal{X}_{\text{with Add. Feat}}$ are 'with Add. Feat'.

[9] The boundaries $\underline{\theta}$ and $\bar{\theta}$ of the control parameters have been chosen after a preliminary analysis with a simplified version of the PeWEC model, to avoid an excessive (or too tight) exploration space of θ.

6. Performance comparison

To assess the learning capabilities and the performance of the proposed strategy, we perform a simulation of one year of learning, and computing the *Expected Annual Energy Production* (EAEP) trend over the year. To do that, following the choice of a deployment site, a validated sequence of sea states (obtained from ERA5 global reanalysis (Hersbach et al., 2020)) is considered, and the corresponding wave forces applied to the WEC devices are computed and simulated, applying every 10 min a control action according to the proposed strategy. Then, every two days of simulation of the learning year:

1. The learning simulation is stopped, and the optimal control parameter $\theta^\star$ for each sea state is computed, forcing an exploitation action which follows Eq. (12).
2. For the purpose of better evaluating the learning performance, once $\theta^\star$ is computed for each sea state, this is applied in 50 simulations per sea state, each one with a different wave realisation of the wave contribution. This has been done to evaluate the variability of the energy production due to the stochastic components characterising the wave (Merigaud and Ringwood, 2018). The results of these simulations are not considered in the experience datasets $\mathcal{J}_{past}$, $\mathcal{S}_{past}$, and θ_{past} employed in the metamodel training.
3. Considering the 50 realisations per sea state, 50 values of the EAEP are computed as:

$$\text{EAEP}_i = \frac{3600 \cdot 24 \cdot 365}{100} \sum_{w=1}^{N_w} Occ_{\%_w} P_{abs_i}\left(S_{T_e,H_{s_w}}, \theta^\star(S_{T_e,H_{s_w}})\right), \tag{19}$$

   where EAEP_i is the EAEP computed using the ith wave realisations, w is the index spanning the N_w sea states with non-zero occurrence percentage $Occ_{\%_w}$, while $P_{abs_i}\left(S_{T_e,H_{s_w}}, \theta^\star(S_{T_e,H_{s_w}})\right)$ is the average absorbed power in $S_{T_e,H_{s_w}}$ wave conditions when the control configuration $\theta^\star(S_{T_e,H_{s_w}})$ is applied and the ith realisation of such wave simulated. This index represents the annual productivity that would be reached by the device if, from the instant in which is computed, the learning process is stopped, and only exploitation actions are applied.
4. The learning simulation is restarted from the point in which it was stopped.

The year simulation is performed adopting the proposed control strategy and employing four different types of metamodels:

- DNN no Add. Feat.: DNN metamodel employing $\mathcal{X} = \mathcal{X}_{\text{no Add. Feat}}$,
- DNN with Add. Feat.: DNN metamodel employing $\mathcal{X} = \mathcal{X}_{\text{with Add. Feat}}$,
- RF no Add. Feat.: RF metamodel employing $\mathcal{X} = \mathcal{X}_{\text{no Add. Feat}}$,
- RF with Add. Feat.: RF metamodel employing $\mathcal{X} = \mathcal{X}_{\text{with Add. Feat}}$.

The results of the EAEP over the learning year are compared with the EAEP obtained, employing in each sea state, control parameters θ computed following an IM approach to interpolate the optimal response in the frequency corresponding to T_e and adopting the linearised model of PeWEC (Gioia et al., 2022). This approach, as discussed within Section 4, is considered as the model-based reference controller.

6.1. Simulation setup

In the performance assessment, to replicate realistic wave conditions, a target deployment site has been chosen in the Mediterranean Sea. In particular, the coast in front of Alghero, in the Northwestern side of Sardinia, Italy, is considered. To perform one year of simulation, the validated sea state sequence provided by ERA5 (Hersbach et al., 2020) over 2018 has been adopted. The sea states occurrence scatter generated by such sequence is shown in Fig. 7. This figure also graphically shows the bins employed by the learning strategy to define $N_{exp.}(S_{T_e,H_s})$, which are $\Delta T_e = 0.5$s, and $\Delta H_s = 0.25$ m respectively, chosen to guarantee that most of the corresponding cells a number of annual occurrences that allows also exploitation actions (and not only the initial exploration). The remainder of the design parameters defining the learning strategy are presented in Table 2. The simulation model employed is the one developed by the nonlinear equations described in Section 2, and the design parameters of the PeWEC device are the ones presented in Carapellese et al. (2022b). It is important to highlight that this model acts as the real unknown system, which is different from the control-oriented linearised model employed to synthesise the IM model-based benchmark control. In the simulation, the wave input contributions have been generated considering a Pierson–Moskowitz spectral density function (SDF) (Pierson and Moskowitz, 1964) and following the *Random Amplitude Scheme* described in Merigaud and Ringwood (2018).

Table 2
Design parameters adopted in the learning strategy.

Design parameter	Value
ΔT_e	0.5 s
ΔH_s	0.25 m
T_{ss}	1 h
T_{min}	4 days
$T_{eval.}$	10 min
$T_{updt.}$	2 days
$\alpha_{gr.}$	1.02
$\beta_{gr.}$	$3.9 \cdot 10^{-3}$
$\gamma_{gr.}$	740
$\delta_{gr.}$	1.11

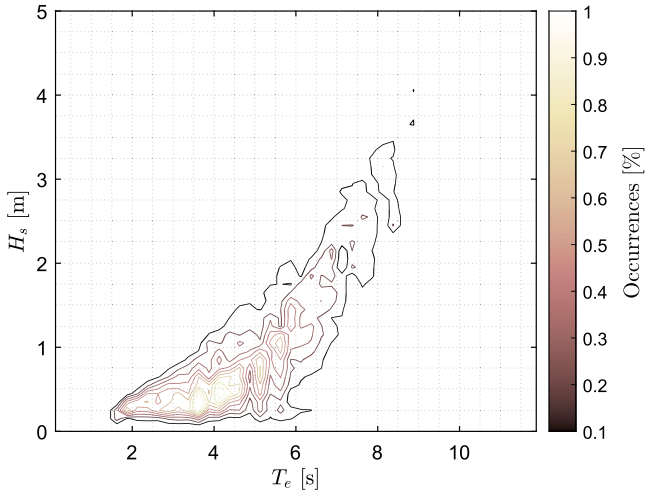

Fig. 7. Occurrence scatter of the sea states (as a function of significant height H_s and energy period T_e) characterising the site of Alghero, Italy, obtained from ERA5 data (Hersbach et al., 2020).

6.2. Results

The first analysis of the results is related to the trends that the EAEP has with the different kinds of metamodels and considered input features. To do that, every two days of simulation, the average, highest and lowest EAEP values among the 50 ones computed considering the different wave realisations are considered. These are presented in Fig. 8, where the coloured solid lines represent the average trends adopting the proposed strategy with the different metamodels and set of input features, compared with the average productivity result obtained with the parameters coming from the model-based IM approach (in dotted black line). The shaded area represents the EAEP values included

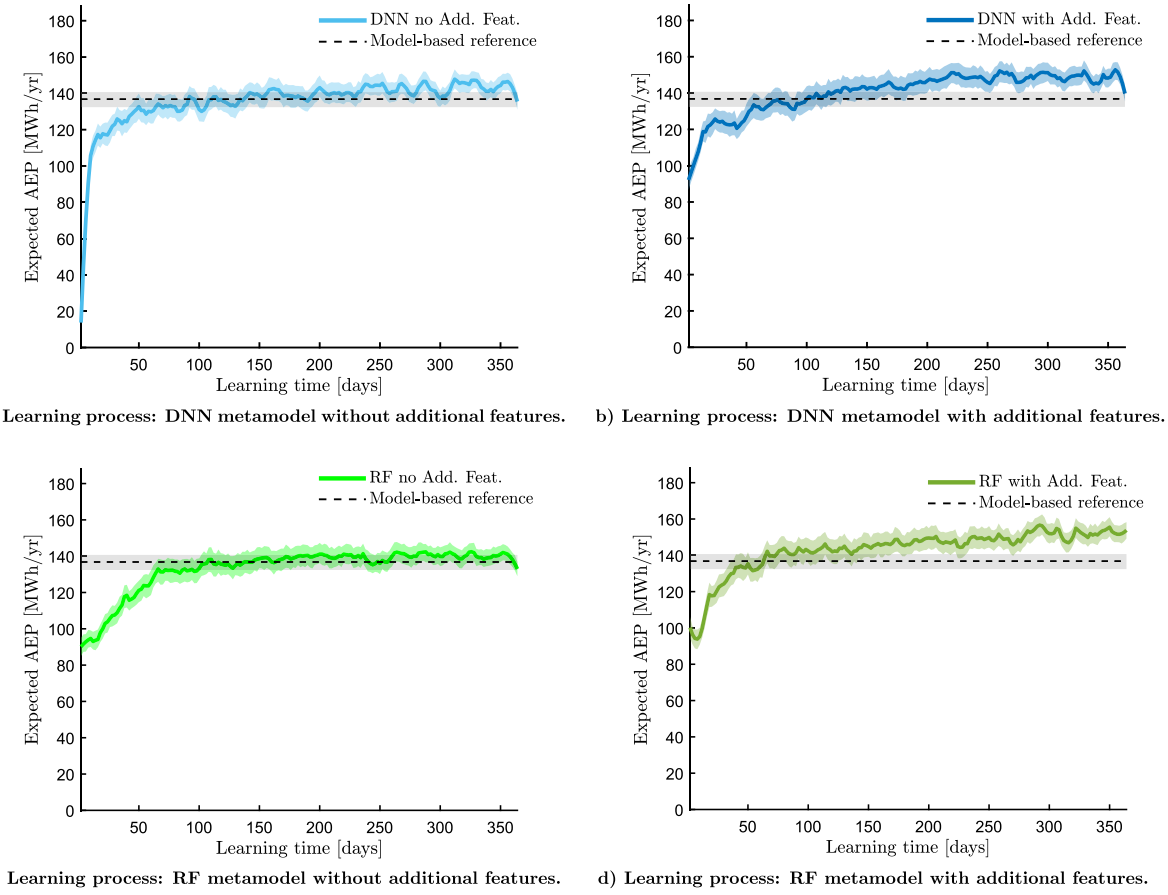

a) Learning process: DNN metamodel without additional features.

b) Learning process: DNN metamodel with additional features.

c) Learning process: RF metamodel without additional features.

d) Learning process: RF metamodel with additional features.

Fig. 8. One year learning simulation: Expected annual energy production EAEP results given by the adoption of different types of metamodels (deep neural networks DNN, and random forests RF) coupled with input features considering or not additional features.

between the highest and lowest EAEP realisation values. As it can be appreciated from such graphs, apart from the strategy employing the RF without added features, all the metamodels enable EAEPs that are consistently above their model-based counterpart. Moreover, the adoption of additional features makes the proposed strategy able to reach a worst case productivity which is higher than the best one obtained with the model-based IM, as Fig. 8a and 8b highlight with both DNN and RF metamodels. At the end of the simulations, the best EAEP obtained by the proposed strategy adopting the different metamodels are 158 MWh/yr (+16.2%), 155 MWh/yr (+14.8%), 147 MWh/yr (+8.1%), and 141 MWh/yr (+3.7%) for the RF with additional features, DNN with additional features, DNN without additional features, and RF without additional features, respectively. Another consideration that can be made is related to the production obtained adopting the proposed control strategy in early learning stages. As it can be seen, DNN metamodels without added features are not able to perform well in the early stages, since the initial expected production is lower than 20 MWh/yr, while the other configurations all begin with productivities around 85 MWh/yr. The choice of the metamodel coupled with the employed features influences also the 'learning speed' of the algorithm. As highlighted by Fig. 9, which summarises the learning process of all the presented approaches, RFs with additional features are the fastest to outperform the model-based productivity (after around 60 days), while the rest of the strategies take about 100 days to reach enough knowledge achieve such benchmark value. The effect of the additional features in speeding up the learning process is more evident whenever RF metamodels are adopted, since the difference in slope of the EAEP trend is milder without the additional features. These have also a positive impact on the long run, as, with both DNN and RF, allow convergence to higher values of EAEP. These results are all

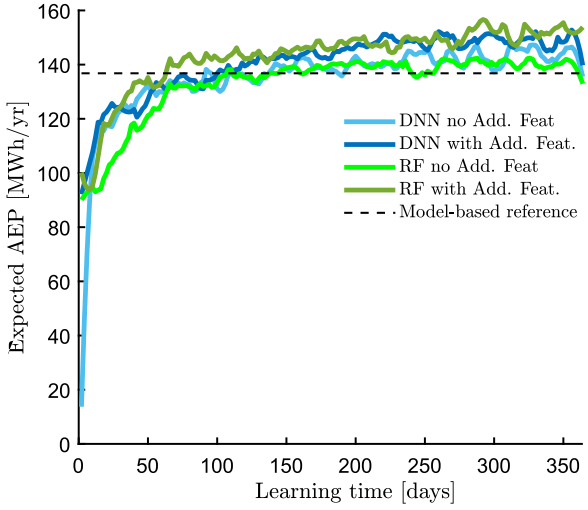

Fig. 9. One year learning simulation: Comparison between the proposed strategy with different metamodel and features combination and the model-based equivalent obtained through IM approach and linearised model.

consequences of the different capabilities in approximating $\mathcal{J}_{\mathcal{M}}$ that the metamodels have.

The metamodel approximation $\hat{f}_{\mathcal{M}}$ depends upon the ability of the metamodel in exploiting the past experience data, and correctly describing the shape of the mapping between S_{T_e,H_s}, θ an average absorbed power P_{abs}. Figs. 10, 11, 12 and 13 show the evolution of the

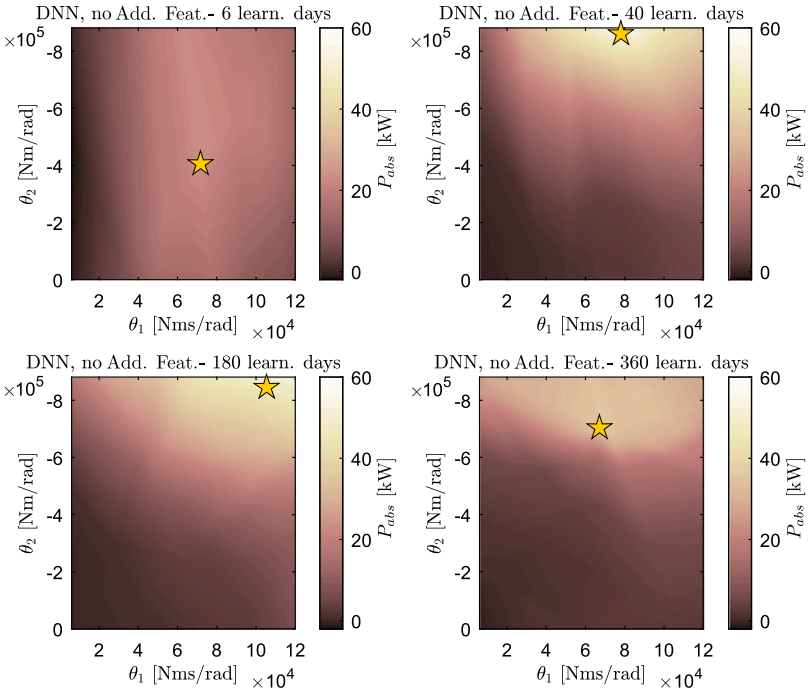

Fig. 10. Metamodel approximation of P_{abs} over the learning year: DNN with no additional features case (T_e = 6 s, H_s = 1 m).

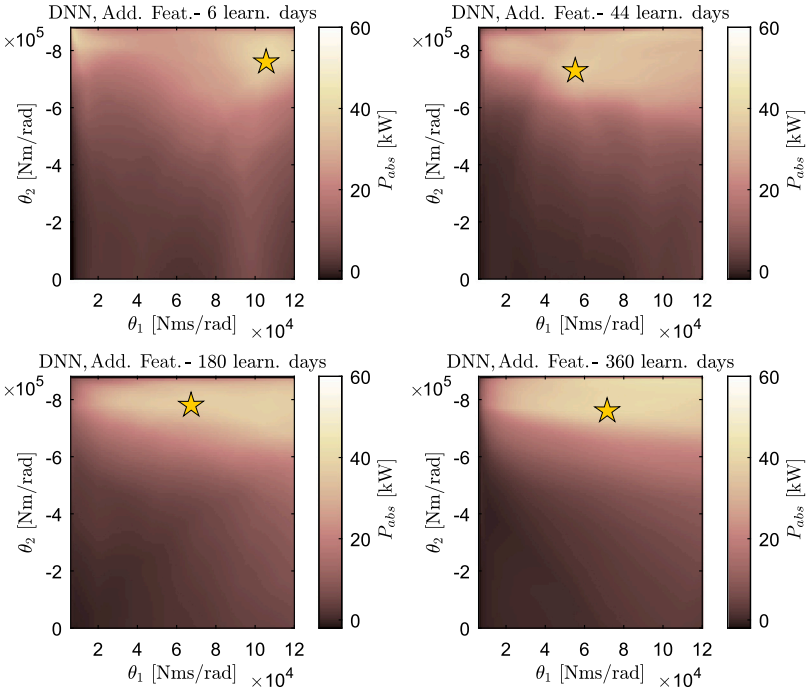

Fig. 11. Metamodel approximation of P_{abs} over the learning year: DNN with additional features case (T_e = 6 s, H_s = 1 m).

metamodel approximation of the average absorbed power P_{abs} over the year of simulation. In these figures, the sea state has been fixed to T_e = 6 s, and H_s = 1 m, which is an average sea state in terms of occurrences in the Alghero deployment site, and that also has an energetic period around the PeWEC hull resonance one. In these figures, the optimal value (*i.e.*, $\theta^\star$ in that sea state at that point of the learning simulation) is represented by a yellow star. As it can be seen from Fig. 10, the adoption of DNN metamodels without additional features is not able to describe well the map to which the other metamodels converge. In fact, it starts 'converging' towards such a map after 40 days, in contrast to what achieved by the other configurations. Adoption of additional features in $\mathcal{X}$ results in a better description of the target relationship, especially after the 40th day. Before, indeed, as Fig. 11 shows, at the 6th day, the metamodel describes a second peak (in the left-upper side of the graph) which is not present whenever the learning process converges to better approximations. Regarding the RF cases, as it is possible to see, the adoption of additional features enhances the approximation of the metamodel from early stages. Comparing the 6th day approximation, it is possible to notice that the effect of the additional features on the metamodel approximation in case of less

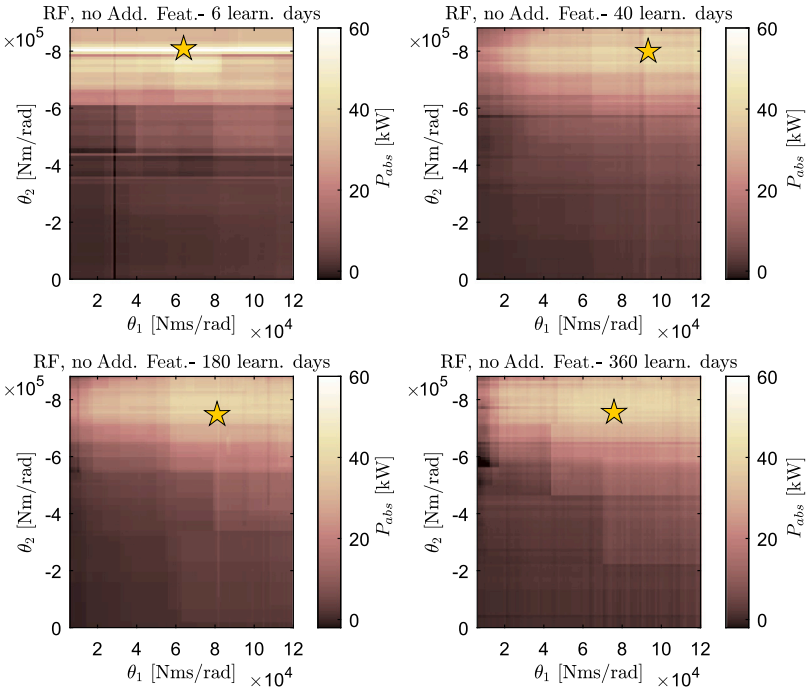

Fig. 12. Metamodel approximation of P_{abs} over the learning year: RF with no additional features case (T_e = 6 s, H_s = 1 m).

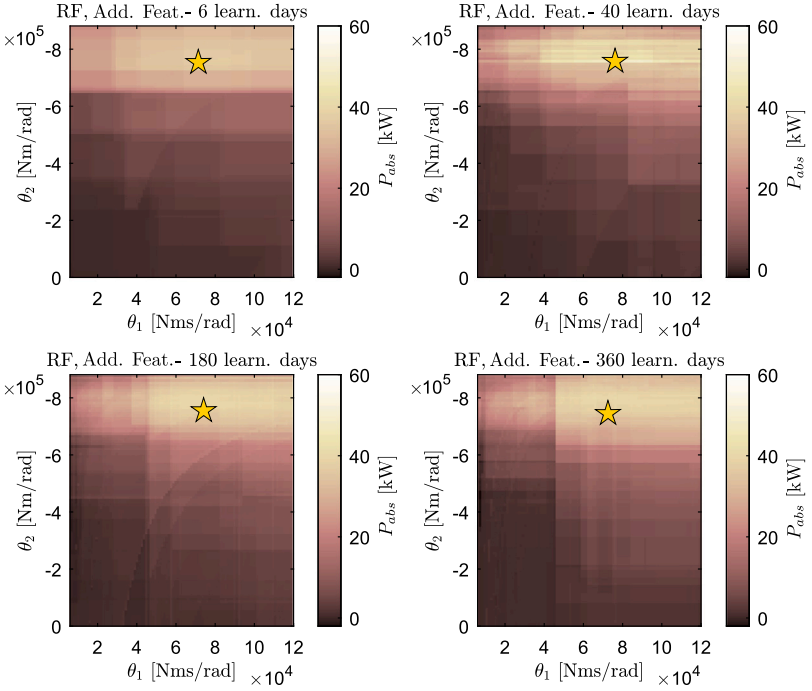

Fig. 13. Metamodel approximation of P_{abs} over the learning year: RF with additional features case (T_e = 6 s, H_s = 1 m).

available data is in 'smoothing' the map described by the approximation itself, and consequently enhancing the obtained result. This is particularly evident comparing the 6th day map with the 40th day one in the case of RF metamodel without additional features. At the 6th day, the expected absorbed power changes abruptly in the neighbourhood of $\theta_2 = -8 \cdot 10^5$ [Nm/rad], while this steep change is mitigated at the 40th day. All these phenomena affect the final energy production over the learning year, making the two configurations (DNN and RF) with additional features the best candidates to be applied in the proposed control strategy.

The final analysis is related to the comparison between the optimal $\theta^\star$ computed with the proposed strategy (with the two best configurations) and the parameters computed following the model-based IM approach. The maps describing the control parameters at the end of the learning year and the IM ones are presented in Fig. 14. As it can be seen from this figure, the control parameters with the proposed approach are almost equal whenever DNN and RF are employed (especially in the area around T_e = 6 s). On the contrary, they differ significantly from those computed employing the linearised model of the device. This can be explained by the presence of nonlinear effects in the model adopted

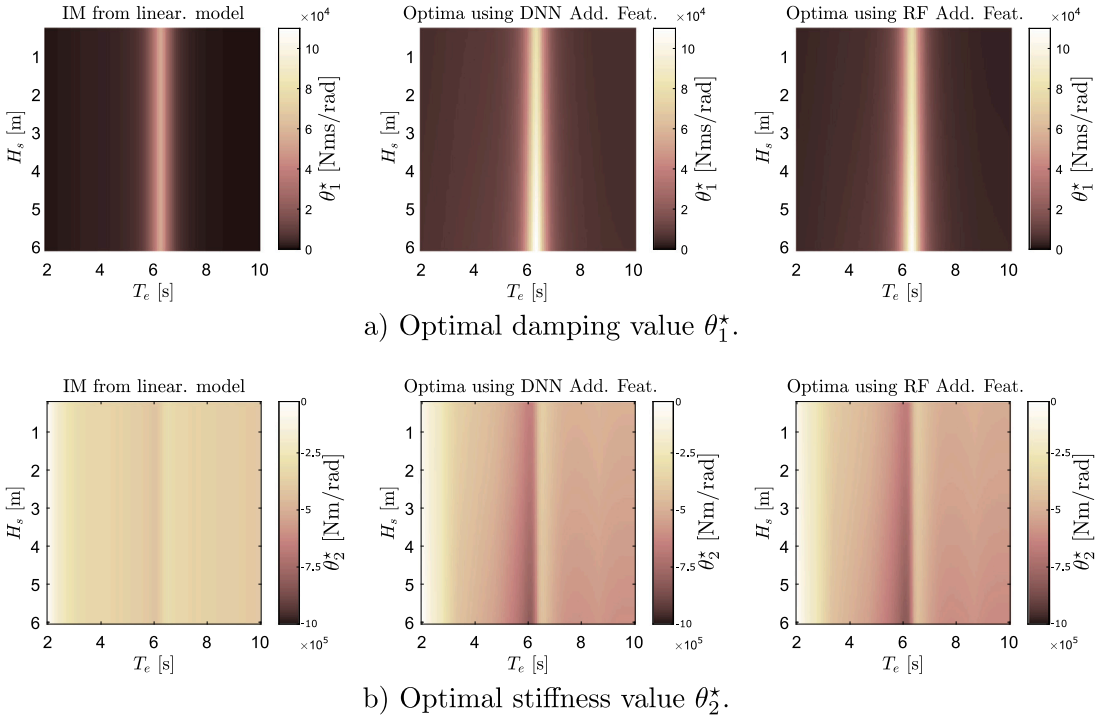

Fig. 14. Comparison of the optimal control parameters $\theta^\star$ for the impedance-matching-based (IM) controller, and data-driven strategies based on deep neural networks (DNN) and random forests (RF), after a year of learning.

in the learning stage, which acts as an additional viscous damping (as also seen in the dynamical analysis of a similar device in Carapellese et al. (2022a), the absence of this sort of damping in a linearised model leads to the overestimation of device motions). As a consequence, the control computed on the basis of the nonlinear system employs higher values of damping and stiffness control parameters.

7. Conclusions

Aiming to ameliorate the effect of modelling uncertainty in the performance of model-based WEC control procedures, in this paper, a data-driven control strategy which adopts RFs and DNNs in its learning process has been proposed. This strategy, inspired by surrogate optimisation algorithms, employs such structures (called metamodels) to approximate the relationship between sea state and control parameters with average absorbed power on the basis of measured data only. Through this approximation, and a developed learning strategy able to manage the trade-off between exploration and exploitation, this model-free controller is able to converge to the optimal parameters of a reactive control law, enhancing the device energy absorption. We analysed the effects that the choice of input features has in the approximation capabilities that these metamodels have, and on the consequent device productivity. To assess the performance of the proposed strategy employing different combinations of metamodels and input features, simulations of one learning year are performed on a nonlinear simulation model of PeWEC (employed here to simulate the real unknown system, different from the linearised one employed as the control-oriented counterpart), computing the trend that the EAEP has over the learning process. The results show average productivity improvements up to +16.2% and +14.8% (RF-based and DNN-based approaches, respectively) with respect to the model-based counterpart employing an IM-based control synthesised on the basis of a linearised model, being able to converge to more effective control parameters for the real WEC system. Further studies will include performance assessment of the proposed technique in experimental environments.

CRediT authorship contribution statement

Edoardo Pasta: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing – original draft, Writing – review & editing, Visualization. **Fabio Carapellese:** Methodology, Software, Validation, Visualization, Writing – review & editing. **Nicolás Faedo:** Methodology, Formal analysis, Writing – review & editing, Visualization, Supervision, Project administration, Funding acquisition. **Paolo Brandimarte:** Conceptualization, Methodology, Formal analysis, Writing – review & editing, Supervision, Project administration.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

EP acknowledges *Fondazione Cecilia Gilardi*, Italy for the support received for this project.

References

Amon, E.A., Brekken, T.K.A., Schacher, A.A., 2012. Maximum power point tracking for ocean wave energy conversion. IEEE Trans. Ind. Appl. 48 (3), 1079–1086. http://dx.doi.org/10.1109/TIA.2012.2190255.

Anderlini, E., Forehand, D., Bannon, E., Abusara, M., 2017. Reactive control of a wave energy converter using artificial neural networks. Int. J. Mar. Energy 19, 207–220. http://dx.doi.org/10.1016/j.ijome.2017.08.001.

Anderlini, E., Forehand, D., Bannon, E., Xiao, Q., Abusara, M., 2018. Reactive control of a two-body point absorber using reinforcement learning. Ocean Eng. 148, 650–658. http://dx.doi.org/10.1016/j.oceaneng.2017.08.017.

Anderlini, E., Forehand, D.I.M., Stansell, P., Xiao, Q., Abusara, M., 2016. Control of a point absorber using reinforcement learning. IEEE Trans. Sustain. Energy 7 (4), 1681–1690. http://dx.doi.org/10.1109/TSTE.2016.2568754.

Bacelli, G., Ringwood, J.V., 2013. A geometric tool for the analysis of position and force constraints in wave energy converters. Ocean Eng. 65, 10–18. http://dx.doi.org/10.1016/j.oceaneng.2013.03.011.

Bratley, P., Fox, B.L., 1988. Algorithm 659 - implementing sobol's quasirandom sequence generator. ACM Trans. Math. Software 14 (1), 88–100. http://dx.doi.org/10.1145/42288.214372.

Breiman, L., Friedman, J.H., Olshen, R.A., Stone, C.J., 2017. Classification And Regression Trees. Routledge, http://dx.doi.org/10.1201/9781315139470.

Carapellese, F., Pasta, E., Faedo, N., Giorgi, G., 2022a. Dynamic analysis and performance assessment of the inertial sea wave energy converter (ISWEC) device via harmonic balance. In: IFAC (Ed.), IFAC-PapersOnLine 55 (31), 439–444. http://dx.doi.org/10.1016/j.ifacol.2022.10.467.

Carapellese, F., Pasta, E., Paduano, B., Faedo, N., Mattiazzo, G., 2022b. Intuitive LTI energy-maximising control for multi-degree of freedom wave energy converters: The PeWEC case. Ocean Eng. 256, 111444. http://dx.doi.org/10.1016/j.oceaneng.2022.111444.

Cramer, J., 2004. The early origins of the logit model. Stud. Hist. Philos. Sci. Part C: Stud. Hist. Philos. Biol. Biomed. Sci. 35 (4), 613–626. http://dx.doi.org/10.1016/j.shpsc.2004.09.003.

Cummins, W., 1962. The Impulse Response Function and Ship Motions. Technical Report 1661. Technical Report, Department of the Navy, David Taylor model basin, Washington DC.

Czech, B., Bauer, P., 2012. Wave energy converter concepts: Design challenges and classification. IEEE Ind. Electron. Mag. 6 (2), 4–16. http://dx.doi.org/10.1109/MIE.2012.2193290.

Faedo, N., Carapellese, F., Pasta, E., Mattiazzo, G., 2022a. On the principle of impedance-matching for underactuated wave energy harvesting systems. Appl. Ocean Res. 118, 102958. http://dx.doi.org/10.1016/j.apor.2021.102958.

Faedo, N., Dores Piuma, F.J., Giorgi, G., Ringwood, J.V., 2020. Nonlinear model reduction for wave energy systems: A moment-matching-based approach. Nonlinear Dynam. 102 (3), 1215–1237. http://dx.doi.org/10.1007/s11071-020-06028-0.

Faedo, N., Garcia-Violini, D., Scarciotti, G., Astolfi, A., Ringwood, J.V., 2019. Robust moment-based energy-maximising optimal control of wave energy converters. In: 2019 IEEE 58th Conference on Decision and Control. (CDC), IEEE, pp. 4286–4291. http://dx.doi.org/10.1109/CDC40024.2019.9029578.

Faedo, N., Giorgi, G., Ringwood, J.V., Mattiazzo, G., 2022b. Optimal control of wave energy systems considering nonlinear Froude–Krylov effects: control-oriented modelling and moment-based control. Nonlinear Dynam. 109 (3), 1777–1804. http://dx.doi.org/10.1007/s11071-022-07530-3.

Faedo, N., Olaya, S., Ringwood, J.V., 2017. Optimal control, MPC and MPC-like algorithms for wave energy systems: An overview. IFAC J. Syst. Control 1, 37–56. http://dx.doi.org/10.1016/j.ifacsc.2017.07.001.

Faedo, N., Scarciotti, G., Astolfi, A., Ringwood, J.V., 2018. Energy-maximising control of wave energy converters using a moment-domain representation. Control Eng. Pract. 81, 85–96. http://dx.doi.org/10.1016/j.conengprac.2018.08.010.

Faedo, N., Scarciotti, G., Astolfi, A., Ringwood, J.V., 2021. Nonlinear energy-maximizing optimal control of wave energy systems: A moment-based approach. IEEE Trans. Control Syst. Technol. 1–15. http://dx.doi.org/10.1109/tcst.2020.3047229.

Falnes, J., 2002. Ocean Waves and Oscillating Systems. Cambridge University Press, http://dx.doi.org/10.1017/CBO9780511754630.

Farajvand, M., Garcia-Violini, D., Windt, C., Grazioso, V., Ringwood, J.V., 2021. Quantifying hydrodynamic model uncertainty for robust control of wave energy devices. In: Proceedings of the 14th European Wave and Tidal Energy Conference. EWTEC 2021, European Wave and Tidal Energy Conference Series, Plymouth, UK, 2251–1–2251–10.

Farajvand, M., Grazioso, V., García-Violini, D., Ringwood, J.V., 2023. Uncertainty estimation in wave energy systems with applications in robust energy maximising control. Renew. Energy 203, 194–204. http://dx.doi.org/10.1016/j.renene.2022.12.054.

Forrester, A.I., Keane, A.J., 2009. Recent advances in surrogate-based optimization. Prog. Aerosp. Sci. 45 (1–3), 50–79. http://dx.doi.org/10.1016/j.paerosci.2008.11.001.

Garcia-Rosa, P.B., Lizarralde, F., Estefen, S.F., 2012. Optimization of the wave energy absorption in oscillating-body systems using extremum seeking approach. In: 2012 American Control Conference. (ACC), pp. 1011–1016. http://dx.doi.org/10.1109/ACC.2012.6314858.

Garcia-Violini, D., Peña-Sanchez, Y., Faedo, N., Ringwood, J.V., 2020. An energy-maximising linear time invariant controller (LiTe-Con) for wave energy devices. IEEE Trans. Sustain. Energy 11 (4), 2713–2721. http://dx.doi.org/10.1109/TSTE.2020.2971392.

Garcia-Violini, D., Ringwood, J.V., 2021. Energy maximising robust control for spectral and pseudospectral methods with application to wave energy systems. Internat. J. Control 94 (4), 1102–1113. http://dx.doi.org/10.1080/00207179.2019.1632491.

Gioia, D.G., Pasta, E., Brandimarte, P., Mattiazzo, G., 2022. Data-driven control of a pendulum wave energy converter: A Gaussian process regression approach. Ocean Eng. 253, 111191. http://dx.doi.org/10.1016/j.oceaneng.2022.111191.

Guo, B., Ringwood, J.V., 2021a. A review of wave energy technology from a research and commercial perspective. IET Renew. Power Gener. 15 (14), 3065–3090. http://dx.doi.org/10.1049/rpg2.12302.

Guo, B., Ringwood, J.V., 2021b. Geometric optimisation of wave energy conversion devices: A survey. Appl. Energy 297, 117100. http://dx.doi.org/10.1016/j.apenergy.2021.117100.

Guo, B., Wang, T., Jin, S., Duan, S., Yang, K., Zhao, Y., 2022. A review of point absorber wave energy converters. J. Mar. Sci. Eng. 10 (10), 1534. http://dx.doi.org/10.3390/jmse10101534.

Hastie, T., Tibshirani, R., Friedman, J., 2009. The Elements of Statistical Learning: Data Mining, Inference, and Prediction, second ed. Springer.

Heaton, J., 2016. An empirical analysis of feature engineering for predictive modeling. In: SoutheastCon 2016. IEEE, pp. 1–6. http://dx.doi.org/10.1109/SECON.2016.7506650.

Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., De Chiara, G., Dahlgren, P., Dee, D., Diamantakis, M., Dragani, R., Flemming, J., Forbes, R., Fuentes, M., Geer, A., Haimberger, L., Healy, S., Hogan, R.J., Hólm, E., Janisková, M., Keeley, S., Laloyaux, P., Lopez, P., Lupu, C., Radnoti, G., de Rosnay, P., Rozum, I., Vamborg, F., Villaume, S., Thépaut, J.-N., 2020. The ERA5 global reanalysis. Q. J. R. Meteorol. Soc. 146 (730), 1999–2049. http://dx.doi.org/10.1002/qj.3803.

Keskar, N., Nocedal, J., Tang, P., Mudigere, D., Smelyanskiy, M., 2017. On large-batch training for deep learning: Generalization gap and sharp minima. In: 5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings.

Kingma, D.P., Lei Ba, J., 2015. Adam: A method for stochastic optimization. In: 3rd International Conference for Learning Representations. San Diego, US, arXiv:1412.6980v9.

Kuhn, M., Johnson, K., 2019. Feature Engineering and Selection. Chapman and Hall/CRC, http://dx.doi.org/10.1201/9781315108230.

Li, G., Belmont, M.R., 2014. Model predictive control of sea wave energy converters – Part I: A convex approach for the case of a single device. Renew. Energy 69, 453–463. http://dx.doi.org/10.1016/j.renene.2014.03.070.

Merigaud, A., Ringwood, J.V., 2018. Free-surface time-series generation for wave energy applications. IEEE J. Ocean. Eng. 43 (1), 19–35. http://dx.doi.org/10.1109/JOE.2017.2691199.

Milligan, G.W., Cooper, M.C., 1988. A study of standardization of variables in cluster analysis. J. Classification 5 (2), 181–204. http://dx.doi.org/10.1007/BF01897163.

Moens de Hase, D., Pasta, E., Faedo, N., Ringwood, J.V., 2021. Towards efficient extremum-seeking control of wave energy systems: possibilities and pitfalls. In: 14th European Wave and Tidal Energy Conference (EWTEC). Plymouth, UK.

Paduano, B., Giorgi, G., Gomes, R.P., Pasta, E., Henriques, J.C., Gato, L.M., Mattiazzo, G., 2020. Experimental validation and comparison of numerical models for the mooring system of a floating wave energy converter. J. Mar. Sci. Eng. 8 (8), http://dx.doi.org/10.3390/JMSE8080565.

Paduano, B., Pasta, E., Papini, G., Carapellese, F., Bracco, G., 2021. Mooring influence on the productivity of a pitching wave energy converter. In: OCEANS 2021: San Diego – Porto. IEEE, pp. 1–6. http://dx.doi.org/10.23919/OCEANS44145.2021.9706108.

Palm, J., Eskilsson, C., Paredes, G.M., Bergdahl, L., 2016. Coupled mooring analysis for floating wave energy converters using CFD: Formulation and validation. Int. J. Mar. Energy 16, 83–99. http://dx.doi.org/10.1016/j.ijome.2016.05.003.

Parrinello, L., Dafnakis, P., Pasta, E., Bracco, G., Naseradinmousavi, P., Mattiazzo, G., Bhalla, A.P.S., 2020. An adaptive and energy-maximizing control optimization of wave energy converters using an extremum-seeking approach. Phys. Fluids 32 (11), 113307. http://dx.doi.org/10.1063/5.0028500.

Pasta, E., Carapellese, F., Brandimarte, P., Parrinello, L., Mattiazzo, G., 2021. A model-free control strategy based on artificial neural networks for pewec. In: 14th European Wave and Tidal Energy Conference. (EWTEC), Plymouth, UK.

Pasta, E., Paduano, B., Mattiazzo, G., Faedo, N., Ringwood, J.V., 2023. On data-based control-oriented modelling applications in wave energy systems. In: Proceedings of the 15th European Wave and Tidal Energy Conference. EWTEC 2023, Bilbao, Spain, p. 409. http://dx.doi.org/10.36688/ewtec-2023-409.

Pasta, E., Papini, G., Faedo, N., Mattiazzo, G., Ringwood, J., 2022. On optimization-based strategies in data-driven control of wave energy systems. In: Trends in Renewable Energies Offshore. CRC Press, London, pp. 401–409. http://dx.doi.org/10.1201/9781003360773-46.

Peña-Sanchez, Y., Merigaud, A., Ringwood, J.V., 2020a. Short-term forecasting of sea surface elevation for wave energy applications: The autoregressive model revisited. IEEE J. Ocean. Eng. 45 (2), 462–471. http://dx.doi.org/10.1109/JOE.2018.2875575.

Peña-Sanchez, Y., Windt, C., Davidson, J., Ringwood, J.V., 2020b. A critical comparison of excitation force estimators for wave-energy devices. IEEE Trans. Control Syst. Technol. 28 (6), 2263–2275. http://dx.doi.org/10.1109/TCST.2019.2939092.

Penalba, M., Giorgi, G., Ringwood, J.V., 2017. Mathematical modelling of wave energy converters: A review of nonlinear approaches. Renew. Sustain. Energy Rev. 78, 1188–1207. http://dx.doi.org/10.1016/j.rser.2016.11.137.

Penalba, M., Ringwood, J.V., 2019. A high-fidelity wave-to-wire model for wave energy converters. Renew. Energy 134, 367–378. http://dx.doi.org/10.1016/j.renene.2018.11.040.

Pierson, W.J., Moskowitz, L., 1964. A proposed spectral form for fully developed wind seas based on the similarity theory of S. A. Kitaigorodskii. J. Geophys. Res. 69 (24), 5181–5190. http://dx.doi.org/10.1029/JZ069i024p05181.

Piramuthu, S., 1996. Feed-forward neural networks and feature construction with correlation information: An integrated framework. European J. Oper. Res. 93 (2), 418–427. http://dx.doi.org/10.1016/0377-2217(96)83599-5.

Piramuthu, S., Ragavan, H., Shaw, M.J., 1998. Using feature construction to improve the performance of neural networks. Manage. Sci. 44 (3), 416–430. http://dx.doi.org/10.1287/mnsc.44.3.416.

Pozzi, N., Bracco, G., Passione, B., Sirigu, S.A., Mattiazzo, G., 2018. PeWEC: Experimental validation of wave to PTO numerical model. Ocean Eng. 167, 114–129. http://dx.doi.org/10.1016/j.oceaneng.2018.08.028.

Pozzi, N., Bracco, G., Passione, B., Sirigu, S.A., Vissio, G., Mattiazzo, G., Sannino, G., 2017. Wave tank testing of a pendulum wave energy converter 1:12 scale model. Int. J. Appl. Mech. 9 (2), http://dx.doi.org/10.1142/S1758825117500247.

Reguero, B., Losada, I., Méndez, F., 2015. A global wave power resource and its seasonal, interannual and long-term variability. Appl. Energy 148, 366–380. http://dx.doi.org/10.1016/j.apenergy.2015.03.114.

Ringwood, J.V., 2020. Wave energy control: status and perspectives 2020. IFAC-PapersOnLine 53 (2), 12271–12282. http://dx.doi.org/10.1016/j.ifacol.2020.12.1162.

Ringwood, J.V., Bacelli, G., Fusco, F., 2014. Energy-maximizing control of wave-energy converters: The development of control system technology to optimize their operation. IEEE Control Syst. 34 (5), 30–55. http://dx.doi.org/10.1109/MCS.2014.2333253.

Ringwood, J.V., Mérigaud, A., Faedo, N., Fusco, F., 2020. An analytical and numerical sensitivity and robustness analysis of wave energy control systems. IEEE Trans. Control Syst. Technol. 28 (4), 1337–1348. http://dx.doi.org/10.1109/TCST.2019.2909719.

Rosati, M., Henriques, J., Ringwood, J., 2022. Oscillating-water-column wave energy converters: A critical review of numerical modelling and control. Energy Convers. Manage.: X 16, 100322. http://dx.doi.org/10.1016/j.ecmx.2022.100322.

Salter, S.H., 1974. Wave power. Nature 249 (5459), 720–724. http://dx.doi.org/10.1038/249720a0.

Scruggs, J., Lattanzio, S., Taflanidis, A., Cassidy, I., 2013. Optimal causal control of a wave energy converter in a random sea. Appl. Ocean Res. 42, 1–15. http://dx.doi.org/10.1016/j.apor.2013.03.004.

Scruggs, J., Nie, R., 2015. Disturbance-adaptive stochastic optimal control of energy harvesters, with application to ocean wave energy conversion. Annu. Rev. Control 40, 102–115. http://dx.doi.org/10.1016/j.arcontrol.2015.09.017.

Shi, S., Patton, R.J., Abdelrahman, M., Liu, Y., 2019. Learning a predictionless resonating controller for wave energy converters. In: Volume 10: Ocean Renewable Energy. American Society of Mechanical Engineers, http://dx.doi.org/10.1115/OMAE2019-95619.

Sirigu, S.A., Foglietta, L., Giorgi, G., Bonfanti, M., Cervelli, G., Bracco, G., Mattiazzo, G., 2020. Techno-economic optimisation for a wave energy converter via genetic algorithm. J. Mar. Sci. Eng. 8 (8), http://dx.doi.org/10.3390/JMSE8070482.

Stansby, P., Carpintero Moreno, E., Stallard, T., 2017. Large capacity multi-float configurations for the wave energy converter M4 using a time-domain linear diffraction model. Appl. Ocean Res. 68, 53–64. http://dx.doi.org/10.1016/j.apor.2017.07.018.

Terrero González, A., Dunning, P., Howard, I., McKee, K., Wiercigroch, M., 2021. Is wave energy untapped potential? Int. J. Mech. Sci. 205, 106544. http://dx.doi.org/10.1016/j.ijmecsci.2021.106544.

Thakur, M., Meghwani, S.S., Jalota, H., 2014. A modified real coded genetic algorithm for constrained optimization. Appl. Math. Comput. 235, 292–317. http://dx.doi.org/10.1016/j.amc.2014.02.093.

Trueworthy, A., DuPont, B., 2020. The wave energy converter design process: Methods applied in industry and shortcomings of current practices. J. Mar. Sci. Eng. 8 (11), 932. http://dx.doi.org/10.3390/jmse8110932.

Tuffin, B., 2004. Randomization of quasi-Monte Carlo methods for error estimation: Survey and normal approximation. Monte Carlo Methods Appl. 10 (3–4), http://dx.doi.org/10.1515/mcma.2004.10.3-4.617.

Wang, R.q., Ning, D.z., Zhang, C.w., Zou, Q.p., Liu, Z., 2018. Nonlinear and viscous effects on the hydrodynamic performance of a fixed OWC wave energy converter. Coast. Eng. 131, 42–50. http://dx.doi.org/10.1016/j.coastaleng.2017.10.012.

Windt, C., Faedo, N., Penalba, M., Dias, F., Ringwood, J.V., 2021. Reactive control of wave energy devices – the modelling paradox. Appl. Ocean Res. 109, 102574. http://dx.doi.org/10.1016/j.apor.2021.102574.

Xiao, X., Huang, X., Kang, Q., 2016. A hill-climbing-method-based maximum-power-point-tracking strategy for direct-drive wave energy converters. IEEE Trans. Ind. Electron. 63 (1), 257–267. http://dx.doi.org/10.1109/TIE.2015.2465964.

Xuhui, Y., Qijuan, C., Zenghui, W., Dazhou, G., Donglin, Y., Wen, J., Weiyu, W., 2019. A novel nonlinear state space model for the hydraulic power take-off of a wave energy converter. Energy 180, 465–479. http://dx.doi.org/10.1016/j.energy.2019.05.095.

Yemm, R., Pizer, D., Retzler, C., Henderson, R., 2012. Pelamis: experience from concept to connection. Phil. Trans. R. Soc. A 370 (1959), 365–380. http://dx.doi.org/10.1098/rsta.2011.0312.

Zhang, Y., Li, G., 2023. Robust tube-based model predictive control for wave energy converters. IEEE Trans. Sustain. Energy 14 (1), 65–74. http://dx.doi.org/10.1109/TSTE.2022.3201771.