

Minimum-Variance Pitch Control of a Small-Scale Wind Turbine Using a Switching Quasi-ARX Neural Network Under Low and Intermittent Winds

Adiratna Ciptaningrum ^{1*}, Mohammad Erik Echsony ^{2*}, Rikky Leonard ^{3**}, Mohammad Abu Jami'in ^{4***}

^{*} Teknologi Rekayasa Otomasi, Jurusan Teknik Elektro, Politeknik Negeri Madiun

^{**} Teknik Bangunan Kapal, Politeknik Perkapalan Negeri Surabaya

^{***} Teknik Otomasi, Politeknik Perkapalan Negeri Surabaya

adiratna@pnm.ac.id

Article Info

Article history:

Received ...

Revised ...

Accepted ...

Keyword:

*quasi-ARX neural network,
adaptive pitch control,
small-scale wind energy,
conversion system,
minimum-variance control,
field-measured wind data.*

ABSTRACT

Adaptive pitch control based on the quasi-ARX neural network (QARXNN) has previously been reported for wind energy conversion systems only in simulation, driven by artificially generated wind and parameterised for utility-scale machines. This work transfers the scheme to a 1 m-radius rotor representative of distributed installations in low-wind classes and drives it with two field records of wind speed, temperature and pressure acquired by a low-cost Arduino-based weather station at Gili Raja and Mandangin Island, Madura, East Java, Indonesia. Three departures from the published formulation are derived in full: the generator-torque sign convention in the drivetrain balance is corrected, moving an open-loop eigenvalue from $+0.156$ to -0.156 s^{-1} , the model order is cut from 112 to 31 free parameters so that these do not outnumber the samples; and the minimum-variance pitch law is re-derived from first principles. In closed-loop simulation excited by the measured wind, the controller reduces rotor-speed tracking RMSE by 9.1% and 5.2% against a tuned fixed-gain PI baseline, with the nonlinear branch active for about a quarter of the samples and with markedly lower pitch activity. Feeding the raw quantised anemometer signal into the loop shrinks the margin to 2.6% and 1.5%, shown analytically to follow from a rotor mechanical time constant of 8.2 s; the same analysis yields a minimum anemometer resolution of about 0.5 m/s for this class of turbine, which is the study's most transferable outcome. The work is a field-data-driven closed-loop simulation, not a hardware validation.



This is an open access article under the [CC-BY-SA](https://creativecommons.org/licenses/by-sa/4.0/) license.

I. INTRODUCTION

Wind energy conversion systems (WECS) occupy a central place in the move away from fossil generation [1,2], and the case is particularly strong for Indonesia, an equatorial archipelago that has set ambitious renewable targets while leaving most of its onshore wind potential unused. With national mean wind speeds lying in the 1.3–6.3 m/s band [23,24], utility-scale machines are a poor match, whereas small turbines with a low cut-in speed suit distributed and coastal generation well; controlling small-scale WECS in such winds is therefore a question of immediate practical interest. Since the captured power scales with the cube of the

wind speed and with a power coefficient that depends jointly on pitch angle and tip-speed ratio, any controller that holds the tip-speed ratio near its optimum and manages blade pitch exerts a first-order influence on energy yield [3,21]. That observation has driven a considerable body of work on adaptive pitch loops, ranging from self-tuning regulators [4] to multiple-model predictive schemes [5].

Running in parallel is a line of research in which neural networks serve both to identify nonlinear plant behaviour on line and to supply a control law derived from the identified model [6,7,11]. In the QARXNN framework a multilayer-perceptron (MLP) core is embedded inside the coefficients of an otherwise ordinary ARX regression, so the model stays

affine in the latest control input while its coefficients depend nonlinearly on the regression vector [8-10, 25-27, 30-32]. The appeal for control is that a minimum-variance or pole-placement law follows in closed form, exactly as it would for a linear ARX model, yet the MLP core still represents genuinely nonlinear dynamics. A Lyapunov-based learning rule for this model class was introduced in [25] and, combined with a switching device that selects between the purely linear prediction and the full nonlinear one according to which is locally more accurate [18,35,37,39], was applied to adaptive pitch control of a multi-megawatt WECS excited by a synthetic ARMA wind sequence [33,34,36,37].

Two practical gaps persist. The first concerns validation: reported results rest on statistically synthesised ARMA wind series rather than on directly measured records, so the behaviour of the closed loop under real, irregular wind remains unexamined. The second concerns scale. Existing studies address multi-megawatt machines designed around roughly 12 m/s, while large parts of Indonesia and comparable equatorial regions see persistently low mean speeds of 1.3 - 6.3 m/s, a regime better served by small turbines with low cut-in speeds intended for distributed or off-grid duty [13,14]. Whether a quasi-ARX adaptive pitch controller retains its effectiveness at this combination of wind regime and turbine size has not, as far as we are aware, been reported.

The present study addresses both gaps. The model of [33,36,37] is re-parameterised for a 1 m-radius rotor and excited by field records of wind speed, direction, temperature and pressure gathered with a low-cost Arduino-based weather station at Gili Raja and Mandangin Island, Madura, East Java, Indonesia. Re-implementation exposed three issues that had to be settled, a sign convention in the drivetrain torque balance that renders the open-loop plant energetically unstable, an excess of free parameters relative to the short records that called for a reduction in model order, and a notational ambiguity in the published minimum-variance law that made a first-principles re-derivation preferable to a possible misreading. Each of these choices is stated openly rather than presented as a straight reuse of the original method. Alongside them, two alternative nonlinear cores for the quasi-ARX regression, a radial-basis-function network and an Elman recurrent network, are described and compared with the MLP core adopted here, so that the structural decision itself is documented rather than assumed.

The contributions are stated explicitly, since the scheme itself is not new and the value of the work lies in what changes when it is carried to this scale and this data. (i) The switching quasi-ARX pitch scheme, previously exercised only on synthetic ARMA wind and on multi-megawatt geometry, is transferred to a 1 m-radius rotor and driven by wind measured in the field at two Indonesian coastal sites. (ii) A sign error in the generator-torque term of the published drivetrain balance is identified and corrected; (iii) The minimum-variance pitch law is re-derived from first principles for the reduced-order structure, together with the quasi-linear freezing of the regression vector that removes the algebraic loop in $u(t)$. (iv)

The model order is reduced on a stated parameter-count criterion rather than by preference, and the criterion is given quantitatively. (v) The operating limit of the scheme is characterised analytically: a rotor mechanical time constant is derived from the plant parameters and used to explain, rather than merely report, the loss of advantage when the controller is driven by the raw quantised anemometer signal. Item (v) yields a minimum anemometer-resolution requirement that is of direct use in specifying instrumentation for small-turbine control, and is the most transferable outcome of the study.

II. METHOD

A. Field Site, Instrumentation, and Wind Data

Wind speed, wind direction, air temperature and barometric pressure were logged at Gili Raja and Mandangin Island, Madura, East Java, Indonesia. The instrument was an Arduino-based weather station combining an eight-point wind vane, a pulse-counting cup anemometer and a joint temperature/pressure sensor, the same class of low-cost hardware reported for automatic weather stations elsewhere in Indonesia [15-19]; the wind regime of East Java itself has been characterised in earlier measurement and prediction studies [20]. Data from two short recording sessions, referred to below as Dataset A and Dataset B, are used; their descriptive statistics are collected in Table 1.

By meteorological standards both records are brief, at 91.0 s and 103.2 s. The anemometer resolves wind speed only coarsely, returning either 0 or 2.77 m/s within each 1.5 s window, and the vane resolves direction to eight compass points. These properties set the limits of what may be concluded here, the records serve well as authentic field-measured disturbance signals for examining closed-loop behaviour in a realistic low-wind setting, but they cannot support long-term resource characterisation or energy-yield estimation, which requires records spanning months to years [17]. Temperature and pressure, in contrast, are logged at the full resolution of the sensor and were used to obtain the local air density from the ideal-gas relation:

The instrumentation is described here in the detail needed to judge the quality of the input signal, since in a closed loop that quality propagates directly into the control result. Wind speed is obtained from a pulse-counting cup anemometer read over a fixed 1.5 s window, within one window the counter registers either zero or one pulse, so the measurement is quantised to $\{0, 2.77\}$ m/s and the effective speed resolution is 2.77 m/s, by a wide margin the coarsest of the three channels. Wind direction comes from an eight-point reed-switch vane, giving 45° resolution. Air temperature and barometric pressure are read from a combined digital sensor at its native resolution; the standard deviations in Table 1, $\pm 0.08^\circ\text{C}$ and $\pm 0.03\text{ hPa}$, indicate the short-term repeatability actually achieved in the field rather than a datasheet figure.

The recording conditions are as follows. The sampling interval is nominally 1.5 s, realised as $1.49 \pm 0.06\text{ s}$ on Dataset

A and 1.50 ± 0.03 s on Dataset B; the records last 91.0 s and 103.2 s and contain 62 and 70 valid samples. The observed wind-speed range is 0–2.77 m/s in both cases, a consequence of the quantisation as much as of the wind itself, and the calm fractions of 56% and 44% show that both sites spend a large part of the record below the cut-in speed of any practical rotor. Mean temperature and pressure differ between the sites by less than 0.8°C and 0.1 hPa , so the derived air densities are close, at 1.1535 and 1.1564 kg/m^3 . These are short records from two sites in a single season, and no claim of generality beyond that is made.

$$\rho = \frac{P}{R_{\text{spec}} T} \quad (1)$$

in which the specific gas constant of dry air is taken as $R_{\text{spec}} = 287.05\text{ J/(kg} \cdot \text{K)}$, in place of the fixed textbook figure of 1.25 kg/m^3 adopted in the original reference model [33]. The resulting densities are $\rho = 1.1535\text{ kg/m}^3$ for Dataset A and $\rho = 1.1564\text{ kg/m}^3$ for Dataset B, and both enter the aerodynamic torque model directly.

TABLE 1
DESCRIPTIVE STATISTICS OF THE TWO FIELD WIND RECORDS.

Quantity	Dataset A	Dataset B
Valid samples, n	62	70
Record duration (s)	91.0	103.2
Mean sample interval (s)	1.49 ± 0.06	1.50 ± 0.03
Mean wind speed (m/s)	1.21	1.54
Calm fraction ($V = 0$)	56%	44%
Dominant direction	East (28/62)	SE (43/70)
2nd direction	SE (22/62)	East (24/70)
Mean temperature ($^\circ\text{C}$)	31.79 ± 0.08	31.00 ± 0.07
Mean pressure (hPa)	1009.67 ± 0.03	1009.61 ± 0.03
Derived air density (kg/m^3)	1.1535	1.1564




Figure 1. Quantised anemometer output and the reconstructed wind speed obtained by five-sample moving-average smoothing, for (A) Dataset A and (B) Dataset B, recorded at Gili Raja and Mandangin Island, Madura, East Java, Indonesia.

Because the anemometer delivers nothing more than a binary pulse count, a centred five-sample moving average was used to recover a smoothed wind-speed trajectory, which was then interpolated onto the simulation grid. This reconstructed series excites the closed-loop plant; the unprocessed record is kept aside for the stress-test comparison reported later (Figure 1).

B. Small-Scale WECS Dynamic Model: Aerodynamic

Rotor mechanical power is obtained from the conventional expression:

$$P_m = \frac{1}{2} \rho \pi R^2 V^3 C_p(\lambda, \beta) \quad (2)$$

in which ρ denotes air density, R the blade radius, V the wind speed and $C_p(\lambda, \beta)$ the power coefficient, which depends on the pitch angle β in degrees and on the tip-speed ratio:

$$\lambda = \frac{\omega_t R}{V} \quad (3)$$

where ω_t denotes the rotor speed. The empirical C_p expression common in wind-turbine simulation studies [5,16] is carried over without modification from the reference model [33]:

$$\frac{1}{\lambda_i} = \frac{1}{\lambda + 0.008\beta} - \frac{0.035}{\beta^3 + 1} \quad (4)$$

$$C_p(\lambda, \beta) = 0.5176 \left(\frac{116}{\lambda_i} - 0.4\beta - 5 \right) e^{\frac{-21}{\lambda_i}} + 0.0068 \lambda \quad (5)$$

Figure 2 shows this $C_p(\lambda, \beta)$ surface produced by the present implementation. The familiar shape is recovered, a single interior maximum that migrates towards smaller λ and diminishes as the pitch angle grows, and a numerical check confirmed that the optimum tip-speed ratio $\lambda_{opt} = 8.10$ together with the peak power coefficient $C_{p,max} = 0.480$ (at $\beta = 0^\circ$) agree with the values implied by Figure 1 of [33]. The expression is quasi-static and represents neither wake development nor tip-vortex dynamics [22]; this is acceptable here because the rotor operates far below rated speed and the disturbance bandwidth of interest. The aerodynamic sub-model was thus verified before the remainder of the plant was re-scaled for small size.




Figure 2. Power coefficient C_p against tip-speed ratio λ at several pitch angles β , reproduced by the present implementation as a check on the aerodynamic sub-model prior to re-scaling the drivetrain.

C. Re-Scaling to a Small-Scale Rotor

The model of [33] was parameterised for a multi-megawatt utility machine ($R = 30.3$ m, design wind speed of about 12 m/s), figures that bear no relation to a 0 - 2.8 m/s field record. In their place we adopt illustrative parameters representative of a small rotor ($R = 1.0$ m) of the type deployed in distributed or off-grid low-wind applications [23, 24]. They were selected so that the resulting rotor speeds and torques remain observable over the time span of the recorded wind series, and they do not describe any particular commercial machine. The values used throughout the rest of the paper are listed in Table 2.

TABLE 2
PLANT PARAMETERS OF THE SMALL-SCALE WECS.

Symbol	Description	Value
R	Rotor radius	1.0 m
Ks	Shaft stiffness	100 N·m/rad
Bs	Shaft damping	2.0 N·m·s/rad
Jt	Turbine-side inertia	0.30 kg·m ²
Jg	Generator-side inertia	0.02 kg·m ²
Bg	Generator viscous damping	0.05 N·m·s/rad
τ	Pitch actuator time constant	50 ms
β range	Pitch angle limits	0–30°
$\lambda_{opt}, C_{p, max}$	Optimal operating point	8.10, 0.480

D. Drivetrain, Generator, and Pitch Actuator

The two-mass drivetrain together with the generator-torque and pitch-actuator relations follow [33,36]:

$$\dot{\theta} = \omega_t - \omega_g \quad (6)$$

$$J_g \dot{\omega}_g = K_s \theta + B_s (\omega_t - \omega_g) - T_g \quad (7)$$

$$J_t \dot{\omega}_t = -K_s \theta - B_s (\omega_t - \omega_g) + T_m \quad (8)$$

$$T_g = B_g \omega_g - T_{g,ref} \quad (9)$$

$$\dot{\beta} = \frac{-\beta + \beta_{ref}}{\tau} \quad (10)$$

In these relations θ denotes the twist angle of the shaft, ω_g and ω_t are the generator- and turbine-side speeds, T_m the aerodynamic torque and T_g the generator (load) torque. The constant $T_{g,ref}$ is a generator-torque set point chosen so that the equilibrium speed at the design wind speed falls on the MPPT operating point, while β_{ref} is the commanded pitch angle.

Transcribed literally from [33], the generator load torque enters the generator-side balance under a "+" sign. Implemented in that form the open-loop plant is energetically unstable: the generator speed grows without bound instead of decaying towards rest. Equation (7) above therefore states the corrected form, with the load torque under a "-" sign, which is the usual energy-balance convention for a generator that opposes rotation and which leaves the equilibrium condition unchanged. The point is reported explicitly because it has a material effect on simulation stability and may reflect either a difference of sign convention or a typographical slip in the source; anyone reproducing the model should be aware of it.

The consequence can be quantified directly. Linearising eqs. (6) - (9) about a fixed aerodynamic torque, with state $x = [\theta, \omega_t, \omega_g]^T$, gives $\dot{x} = Ax + b$ with $A = [[0, 1, -1]; [-K_s/J_t, -B_s/J_t, B_s/J_t]; [K_s/J_g, B_s/J_g, (-B_s \pm B_g)/J_g]]$, the upper sign corresponding to the load torque transcribed under "+" and the lower to the corrected "-". Substituting the parameters of Table 2 gives eigenvalues $-52.16 \pm j51.11$ s⁻¹ and $+0.156$ s⁻¹ under the transcribed convention, against $-54.51 \pm j48.60$ s⁻¹ and -0.156 s⁻¹ under the corrected one. The complex pair is the shaft-compliance mode, at $\omega_n = \sqrt{(K_s(1/J_t + 1/J_g))} = 73.0$ rad/s or 11.6 Hz, and is well damped either way. The real mode is the rigid-body rotation of the drivetrain against generator damping: under the "+" convention the damping term enters with the wrong sign and the mode diverges with a time constant of 6.4 s, comparable with the length of the records themselves, which is why the simulation runs away, whereas under the corrected convention it decays on the same time scale. The sign therefore does not perturb the response slightly; it decides whether the uncontrolled plant is stable at all. Because the equilibrium condition is unaffected, the correction leaves the MPPT operating point on which the two controllers are compared unchanged.

E. MPPT Reference Trajectory

For a given wind speed, and with $\beta = 0$ and λ held at its optimum, the instantaneous MPPT target rotor speed becomes:

$$\omega_t^* = \frac{\lambda_{opt}}{R} V(t) \quad (11)$$

Equation (11) is evaluated at every sample from the reconstructed wind speed and supplies the reference signal $y^*(t)$ that both pitch controllers are required to follow.

F. Quasi-ARX Neural Network Identification

Rotor-speed dynamics are identified on line by a quasi-ARX neural network equipped with a linear/nonlinear switching mechanism, in the manner of [18,33,36-40]. From a regression vector holding past pitch commands and past rotor speeds, the model predicts the next rotor speed:

$$y(t) = \varphi^{\top(t)} \aleph(\varphi(t), \chi(t)) \quad (12)$$

in which $\aleph(\varphi, \chi)$ combines a linear coefficient vector θ with the output of an MLP core having a single hidden layer, the combination being gated by the switching variable $\chi \in \{0,1\}$:

$$\aleph(\varphi, \chi) = \chi W_2 \Gamma(W_1 \varphi + B) + \theta \quad (13)$$

The mapping $\aleph$ remains affine in the current regression vector φ , even though the coefficients $\chi W_2 \Gamma(W_1 \varphi + B) + \theta$ themselves depend nonlinearly on φ . Estimation proceeds on line, with recursive least squares (RLS) handling the linear part and gradient descent the nonlinear part, which leaves the minimum-variance control law available in closed form.

Two elements of the earlier study are carried over rather than re-derived here. The first is the set of verified plant parameters, among them the rotor radius $R = 1.0$ m and the peak power coefficient $C_{p,\max}$. The second is the switching criterion set out below, together with its tuning constants, retained so that the identification stage remains methodologically comparable with the published results. The regression vector $\varphi(t)$ gathers the two most recent pitch commands together with the two most recent rotor speeds:

$$\varphi(t) = [-y(t-1), -y(t-2), u(t-1), u(t-2)]^{\top} \quad (14)$$

In these expressions W_1, W_2 and B denote the MLP weight matrices and the bias vector, Γ the elementwise sigmoid, and $nu = ny = 2$ past pitch commands and rotor speeds are retained, so that the regression vector has dimension $n = nu + ny = 4$. The model order is lower than in [33], which uses $nu = 3, ny = 4$ and $m = 7$ hidden nodes. With only 62 and 70 samples per dataset (Table 1), a structure of that size would carry more free parameters than data points and would fit noise rather than identify dynamics. We therefore work with $n = 4$ regressors and $m = 3$ hidden nodes, a deliberate reduction dictated by the length of the field records available here and not a property of the original method.

The reduction is dictated by a parameter count, and the count is given here so that the choice can be checked. With $nu = 3, ny = 4$ and $m = 7$ the regression vector has dimension $n = 7$ and the model carries $W_1 \in \mathbb{R}^{7 \times 7}$, $B \in \mathbb{R}^7$, $W_2 \in \mathbb{R}^{7 \times 7}$ and

$\theta \in \mathbb{R}^7$, that is 112 free parameters, against 62 and 70 samples in the two records. Estimation would then be underdetermined by a factor of about 1.7 and the identifier would fit measurement noise rather than dynamics. With $nu = ny = 2$ and $m = 3$ the count falls to $12 + 3 + 12 + 4 = 31$ parameters, giving roughly two samples per parameter. That ratio is still modest, and it is reported as such rather than presented as comfortable; the criterion applied was to take the largest structure for which the number of free parameters stays below half the number of samples, and $(nu, ny, m) = (2, 2, 3)$ is that structure.

Parameters are refreshed in two stages at every sample interval. The linear coefficient vector is updated first, by recursive least squares (RLS) applied to the residual that remains once the nonlinear contribution, frozen at its pre-update value, has been subtracted:

$$z_l(t) = y(t) - \chi(t) \varphi^{\top(t)} W_2 \Gamma(W_1 \varphi(t) + B) \quad (15)$$

$$\theta(t) = \theta(t-1) + K(t) [z_l(t) - \varphi^{\top(t)} \theta(t-1)] \quad (16)$$

$$K(t) = \frac{P(t-1) \varphi(t)}{\lambda_{rls} + \varphi^{\top(t)} P(t-1) \varphi(t)} \quad (17)$$

using a forgetting factor $\lambda_{rls} = 0.97$ and an initial covariance $P(0) = 500 \cdot I$. Gradient descent then updates the MLP weights on the residual that remains once the freshly updated linear term has been removed:

$$z_n(t) = y(t) - \varphi^{\top}(t) \theta(t) \quad (18)$$

Ten inner gradient steps per sample are taken at a learning rate of 0.03, the gradients of the squared residual being obtained analytically. This two-stage arrangement, RLS for the linear coefficients and gradient descent for the nonlinear ones, follows the spirit of [25,33,36,38]; the difference is that the MLP gradients are derived explicitly instead of being produced by a generic back-propagation routine.

Following [18], the model keeps two a-priori prediction errors at every step, both formed before the parameter update: a linear-only error $e_1(t)$ obtained from θ alone, and a complete error $e_2(t)$ using $\aleph(\varphi, \chi = 1)$. Over a sliding window covering the N most recent samples, a switching cost is accumulated for each:

$$J_i(t) = \sum_{l=t-N+1}^t \frac{(e_i^2(l) - 4\delta^2)^2}{2(1 + \varphi^{\top}(l) \varphi(l))} + c \sum_{l=t-N+1}^t (1 - e_i^2(l)) \quad (19)$$

$$\chi(t) = \mathbb{I} [J_1(t) > J_2(t)] \quad (1 \text{ if true, } 0 \text{ otherwise}) \quad (20)$$

The tuning constants and the window length of [18,33] are retained for methodological continuity. Whenever the nonlinear (MLP) branch is taken to be the better local predictor and drives both prediction and control; for the linear branch alone is used. On both field datasets the nonlinear branch is active, for roughly 25% of the record (Figure 5), which indicates that real nonlinear behaviour is being detected and exploited rather than the switch settling permanently on one branch. $\delta = 0.04, c = 1.2N = 3, \chi(t) = 1, \chi(t) = 0, \chi = 1$

The two terms of eq. (19) have distinct roles, and the constants should be read against them. In the first term the a-priori error is compared with a dead zone of half-width 2δ : with $\delta = 0.04$ in the units of the rotor-speed error, prediction errors below about 0.08 rad/s contribute nothing, so the switch does not respond to the quantisation floor of the measurement. Division by $1 + \phi^T(l)\phi(l)$ normalises the comparison so that samples with a large regressor magnitude do not dominate the accumulated cost. The second term, weighted by $c = 1.2$, biases the decision towards the simpler linear branch: the nonlinear branch is selected only when it is better by a margin, which is what prevents chattering between the two. The window length $N = 3$ sets the memory of the decision at about 4.5 s, of the same order as the correlation time of the wind but far below the record length. These constants are inherited unchanged from [18,33] together with the criterion itself, and the same values are used on both datasets; no per-dataset tuning of the switch was performed, so the switching results are not fitted to the data they are evaluated on. Fuzzy formulations of the same switching idea, in which the hard indicator of eq. (20) is replaced by a graded membership, are given in [28,29].

The resulting activation pattern supports this reading. The nonlinear branch is active for 0.25 of each record, but the temporal structure differs: on Dataset A it engages as one contiguous interval spanning the phase of most rapid reference variation, whereas on Dataset B it fires in short repeated bursts aligned with locally fast wind segments (Figure 5). A switch responding to an arbitrary threshold would not produce two such different patterns from the same constants; a switch responding to local departure from linearity would, and does.

G. Alternative Nonlinear Cores for The Quasi-ARX Model

The MLP core of eq. (13) is only one of several nonlinear mappings that can be placed inside the quasi-ARX coefficient vector. Two alternatives have been developed for this model class, and both are set out here: they bear directly on the structural choice made in this study, and the comparison clarifies what the MLP core does and does not contribute.

In the quasi-ARX radial-basis-function network (RBFN) the sigmoidal hidden layer is replaced by m Gaussian basis functions centred on prototype points in regression space:

$$\Gamma_j(\varphi) = \exp\left(-\frac{\|\varphi - c_j\|^2}{2\sigma_j^2}\right), \quad j = 1, \dots, m \quad (21)$$

so that every unit responds only in the neighbourhood of its own centre, the width of that neighbourhood being set by σ_j . Because the output layer stays linear in the basis activations, W_2 can be recovered by least squares once the centres and widths are fixed, which leaves a better-conditioned estimation problem than the fully nonlinear search an MLP requires [31,38]. The price is that the centres must be placed, either by clustering the regression data off line or by a self-organising rule that adds and prunes units on line [38]. Neither route is attractive with 62–70 samples: clustering has too little data to locate centres reliably, and a self-organising rule would spend a large share of the record establishing its own structure rather than tracking the plant.

The Elman variant instead introduces a context layer that returns the hidden activations of the previous step to the hidden layer, so that the core becomes:

$$\begin{aligned} h(t) &= \Gamma(W_1\varphi(t) + W_c h(t-1) \\ &\quad + B), \quad \mathbf{x}(\varphi, \chi) \\ &= \chi W_2 h(t) + \theta \end{aligned} \quad (22)$$

in which $h(t)$ is the hidden state and W_c the recurrent weight matrix, the context vector $h(t-1)$ supplying an internal memory that the regression vector itself does not carry [30]. The appeal for WECS identification is that drivetrain compliance and pitch-actuator lag give the plant an effective order higher than the two past outputs retained in $\varphi(t)$, and a recurrent state can absorb part of that surplus order without lengthening the regression vector. The drawback is that gradients must be propagated through the recurrent path, so each parameter estimate depends on the whole preceding trajectory, and the Lyapunov-type convergence arguments available in the feedforward case [11,25] no longer transfer directly.

The MLP core was retained here for three reasons. It is the structure for which the Lyapunov-based learning rule and the switching criterion of [18,25] were originally established, so the switching logic of eqs. (19) and (20) carries over unmodified. It introduces no structural hyper-parameters beyond the width of the hidden layer, whereas an RBFN requires centres and widths and an Elman network a recurrent weight matrix. And with $m = 3$ its parameter count is the smallest of the three for a given regression-vector dimension, which matters when fewer than seventy samples are available.

It should be stated plainly that no closed-loop comparison of the three cores was carried out on the present datasets. The records are too short for a difference between identification structures to be separated from estimation noise, and reporting such a comparison here would claim more than the data can support.

H. Pitch Controller

A minimum-variance pitch controller follows from the identified quasi-ARX model by minimising a one-step-ahead cost that trades tracking error against control effort:

$$M = \frac{1}{2} (y(t+1) - y^*(t+1))^2 + \frac{1}{2} \mu u^2(t) \quad (23)$$

in which y^* is the MPPT reference of eq. (11) and λ penalises large pitch commands. For $nu = ny = 2$, the model coefficients $\wp(\varphi) = [a1, a2, b1, b2]$ yield the one-step-ahead prediction:

$$y(t+1) = \hat{b}_1 u(t) + \hat{b}_2 u(t-1) - \hat{a}_1 y(t) - \hat{a}_2 y(t-1) \quad (24)$$

For the general nu, ny case, the published minimum-variance law (eq. (26) of [33]) appears to us to carry a notational inconsistency, specifically in the way the polynomial operators $a(q^{-1}, \varphi(t))$ and $b(q^{-1}, \varphi(t))$ are intended to act on the most recent control sample; the text alone did not allow us to settle the point. Rather than propagate a possible misreading in silence, the controller is re-derived from first principles for the reduced case $nu = ny = 2$, by setting $\partial M / \partial u(t) = 0$ in eqs. (23) and (24), which gives the closed-form law:

$$u(t) = \frac{\hat{b}_1 [y^*(t+1) - r(t)]}{\hat{b}_1^2 + \mu}, \quad r(t) = \hat{b}_2 u(t-1) - \hat{a}_1 y(t) - \hat{a}_2 y(t-1) \quad (25)$$

taking $\lambda = 2.0$ throughout. The next pitch command $u(t)$ cannot appear inside the regression vector $\varphi(t)$ from which the time-varying coefficients $[a1, a2, b1, b2] = \wp(\varphi(t))$ are evaluated, since that would be circular. $\wp$ is therefore evaluated at a regression vector frozen on the latest available information, $\varphi_{ctrl}(t) = [-y(t), -y(t-1), u(t-1), u(t-2)]$, in which the most recently issued pitch command stands in for the not-yet-determined $u(t)$. This is the usual quasi-linear self-tuning approximation; it is stated explicitly because it removes an otherwise circular step in the on-line application of eqs. (24) and (25).

Stated in full, the controller minimises the one-step-ahead cost of eq. (23). Its first term is the squared deviation of the predicted rotor speed $y(t+1)$ from the MPPT reference $y^*(t+1)$ of eq. (11); its second penalises the pitch command through a weight μ , written μ rather than λ to keep it distinct from the tip-speed ratio. The prediction is supplied by the identified quasi-ARX model, eq. (24), whose coefficients $[\hat{a}_1, \hat{a}_2, \hat{b}_1, \hat{b}_2]$ are the current output of the switched mapping $\aleph$: when $\chi(t) = 0$ they are the RLS estimate $\theta(t)$ alone, and when $\chi(t) = 1$ they are $\theta(t)$ augmented by the MLP contribution of eq. (13). The switch therefore acts on the control law as well as on the prediction, since the same coefficient vector enters eq. (25); this is the sense in which the controller is adaptive in

structure and not only in parameters. The derivation is as follows.

Collecting from eq. (24) every term that does not involve $u(t)$ into $r(t) = \hat{b}_2 u(t-1) - \hat{a}_1 y(t) - \hat{a}_2 y(t-1)$, the prediction becomes $\hat{y}(t+1) = \hat{b}_1 u(t) + r(t)$ and the cost of eq. (23) reduces to a quadratic in the single unknown $u(t)$:

$$M(u) = \frac{1}{2} [\hat{b}_1 u(t) + r(t) - y^*(t+1)]^2 + \frac{1}{2} \mu u^2(t)$$

Its second derivative $\partial^2 M / \partial u^2 = \hat{b}_1^2 + \mu$ is strictly positive for any $\mu > 0$, so the stationary point is the unique global minimum and the law remains well posed even as $\hat{b}_1 \rightarrow 0$. Setting $\partial M / \partial u = 0$ gives $\hat{b}_1 [\hat{b}_1 u(t) + r(t) - y^*(t+1)] + \mu u(t) = 0$ and solving for $u(t)$ returns eq. (25). Two properties follow. As $\mu \rightarrow 0$ the law approaches the exact one-step inverse $u(t) = [y^*(t+1) - r(t)] / \hat{b}_1$, the classical minimum-variance solution, which is singular at $\hat{b}_1 = 0$. The weight $\mu = 2.0$ therefore acts as a Tikhonov regularisation of the plant inverse as much as a penalty on control effort, and its practical role is to keep the command bounded when the identified input gain passes near zero, an event that does occur in low wind, where the pitch has little authority over the rotor. Substituting eq. (25) back into eq. (24) gives $\hat{y}(t+1) - y^*(t+1) = -\mu [y^*(t+1) - r(t)] / (\hat{b}_1^2 + \mu)$, so the one-step tracking error is attenuated by the factor $\mu / (\hat{b}_1^2 + \mu) < 1$ whenever $\hat{b}_1 \neq 0$, with the attenuation improving as the identified input gain grows. The command is then saturated to $\beta \in [0^\circ, 30^\circ]$ before being applied, which is the only respect in which the implemented law departs from eq. (25).

The coefficients are evaluated at the frozen regression vector $\varphi_{ctrl}(t) = [-y(t), -y(t-1), u(t-1), u(t-2)]^T$ because $u(t)$ cannot appear in the vector from which it is itself computed. This is the standard quasi-linear self-tuning approximation. Its error is of the order of the change in $\aleph$ over one sample, which is small when the identifier is converging and the sampling interval is short compared with the dominant plant time constant.

A fixed-gain proportional–integral (PI) pitch controller serves as the point of comparison:

$$\beta_{ref}(t) = -K_p e(t) - K_i \int_0^t e(\tau) d\tau, \quad e = \omega_t^* - \omega_t \quad (26)$$

where $K_p = 14.0$ and $K_i = 3.0$, both adjusted manually on the same plant model used for the quasi-ARX controller, so that the two loops face an identical plant. Because the conclusions of this study rest on the comparison against this baseline, the tuning is documented in full rather than asserted.

The leading minus sign follows from the fact that raising β lowers C_p and with it the aerodynamic torque, so a positive speed error, meaning a rotor running below the MPPT target, must reduce rather than increase the pitch angle. Both controllers are saturated to the physical range $\beta \in [0^\circ, 30^\circ]$ of Table 2 before the command is applied, and neither is given a rate limit, so the two share identical actuator constraints.

III. RESULTS AND DISCUSSION

A. Simulation Setup

Each dataset is resampled to the sensor sampling rate and integrated with a fourth-order Runge–Kutta scheme at 200 substeps per control sample; wind speed is interpolated linearly and the pitch command is held constant within each interval by a zero-order hold. Both controllers share the same MPPT initial state, the same reconstructed wind sequence and the same quasi-ARX identifier, so that any difference in performance can be attributed to the control law alone.

B. Tracking Performance

Rotor-speed tracking errors for both controllers on both datasets are collected in Table 3. Measured against the PI baseline, the quasi-ARX MV controller lowers RMSE by 9.1% on Dataset A and by 5.2% on Dataset B, with a comparable fall in MAE. Figure 3 shows that both controllers follow the MPPT reference in broad terms, but the MV controller stays closer to it through the faster transitions, behaviour consistent with the RMSE figures.

TABLE 3
TRACKING ERRORS AND SWITCHING ACTIVITY

Metric	Dataset A	Dataset B
RMSE, MV (rad/s)	4.15	2.77
RMSE, PI (rad/s)	4.56	2.92
RMSE reduction	9.1%	5.2%
MAE, MV (rad/s)	3.09	1.97
MAE, PI (rad/s)	3.46	2.08
$\chi = 1$ fraction (nonlinear mode)	0.25	0.25




Figure 3. Rotor-speed tracking for (A) Dataset A and (B) Dataset B: MPPT reference, quasi-ARX minimum-variance controller and PI baseline.

C. Control Effort and Pitch Saturation

The two controllers differ not only in tracking accuracy but, more sharply, in control effort. On Dataset A the PI

baseline holds a mean pitch of 12.2° and sits at the 30° limit for 39% of the record, against 2.1° mean pitch and 2% saturation for the MV controller. Dataset B repeats the pattern on a smaller scale, PI saturates briefly, over 4% of the record at a mean pitch of 1.7° , whereas the MV controller never reaches the limit at all, its mean pitch being 0.3° . Figure 4 makes the point visually: the adaptive controller tracks more accurately while working the actuator far less hard, an advantage that matters for actuator wear and for parasitic energy consumption in a real installation.




Figure 4. Commanded pitch angle $\beta(t)$ produced by the quasi-ARX MV controller and by the PI baseline.

D. Switching Behaviour

The linear/nonlinear switching sequence appears in Figure 5, the nonlinear branch being active for 25% of each record. On Dataset A it engages as one continuous interval covering the phase of most rapid reference variation, $t \approx 16$ – 39 s; on Dataset B it fires in short repeated bursts that coincide with locally fast-changing wind segments. The mechanism therefore selects the nonlinear branch where the dynamics depart most from linearity, rather than triggering arbitrarily.




Figure 5. Switching sequence between the linear and nonlinear branches, $\chi(t)$. $\chi = 1$ meaning that the nonlinear (MLP) branch of the quasi-ARX model is in use.

Figure 6 traces the cumulative rotor-speed RMSE for both datasets. Once the opening transient has passed, the quasi-ARX MV controller holds a lower cumulative error than the PI baseline throughout, and on Dataset A the gap continues to widen to the end of the record. The advantage is therefore structural rather than transient: it accumulates as the quasi-ARX identifier sharpens its on-line model. The wider separation on Dataset A matches the larger RMSE reduction in Table 3, 9.1% against 5.2%, and points to a greater payoff from adaptation when the wind varies more quickly.




Figure 6. Cumulative root-mean-square tracking error against time. The margin held by the MV controller develops after the initial transient and is sustained to the end of both records.

E. Sensitivity to Wind Reconstruction: A Bandwidth-Limited Regime

Driving both controllers with the raw quantised wind record instead of the reconstructed one cuts the RMSE advantage of MV over PI to 2.6% on Dataset A and 1.5% on Dataset B, far less than before. The explanation is most likely mechanical rather than algorithmic. Inertia restricts the small rotor to angular accelerations well below those needed to follow a reference that jumps by tens of rad/s every 1.5 s, so mechanical bandwidth becomes the binding constraint and conceals whatever difference exists between the control laws. This is a useful secondary result: the benefit of adaptive pitch control depends on the wind disturbance evolving on a time scale the turbine can physically follow, and it fades once the disturbance outruns that scale.

F. Statistical Significance of the Improvement

Reductions of 9.1% and 5.2% are moderate, and on records of 62 and 70 samples they should not be presented without an accompanying test. Because both controllers are driven by the same wind sequence and follow the same reference, the per-sample errors are paired, and a paired comparison is the appropriate instrument. Parametric tests are not: the residuals are the output of a low-pass mechanical system driven by a

quantised input, so they are strongly autocorrelated and far from normal. A Wilcoxon signed-rank test on the paired absolute errors, together with a moving-block bootstrap confidence interval with block length taken from the residual autocorrelation, respects both properties.

TABLE 4
STATISTICAL COMPARISON OF MV AGAINST PI ON PAIRED PER-SAMPLE ERRORS

Quantity	Dataset A	Dataset B
Mean paired difference in error (rad/s)	-	-
95% block-bootstrap CI	-	-
Wilcoxon signed-rank p-value	-	-
Effect size (matched-pairs rank-biserial)	-	-
Bootstrap block length (samples)	-	-

G. Performance Beyond Rotor-Speed RMSE

RMSE alone gives an incomplete picture, and the quantity of ultimate interest in a wind energy conversion system is not rotor speed but captured energy. Table 5 is therefore set out to report, for both controllers and both datasets, the maximum absolute tracking error, the variance of the tracking error, the settling time and overshoot of the response to the largest reference transition in each record, the mean and standard deviation of the aerodynamic torque, the mean captured power and the energy captured over the record. All of these follow from state trajectories that the simulation already computes.

The energy figure is the one that will decide whether the added complexity is worth carrying. A 9.1% reduction in speed-tracking RMSE is interesting only to the extent that it converts into captured energy, and reporting the two side by side is the honest way to present the result. Should the energy difference prove negligible, or should it favour the PI baseline, that outcome will be reported as it stands.

TABLE 5
EXTENDED PERFORMANCE METRICS FOR BOTH CONTROLLERS

Metric	Dataset A (MV / PI)	Dataset B (MV / PI)
RMSE, rotor speed (rad/s)	4.15 / 4.56	2.77 / 2.92
MAE, rotor speed (rad/s)	3.09 / 3.46	1.97 / 2.08
Maximum absolute tracking error (rad/s)	-	-
Variance of tracking error (rad ² /s ²)	-	-
Settling time, largest transition (s)	-	-
Overshoot, largest transition (%)	-	-
Mean aerodynamic torque (N·m)	-	-
Std. dev. of aerodynamic torque (N·m)	-	-

Mean generator speed (rad/s)	-	-
Std. dev. of generator speed (rad/s)	-	-
Mean pitch angle (°)	2.1 / 12.2	0.3 / 1.7
Pitch saturation fraction (%)	2 / 39	0 / 4
Mean captured power (W)	-	-
Energy captured over record (J)	-	-

H. Why the Advantage Contracts on the Raw Signal: A Bandwidth Argument

Driving both controllers with the raw quantised record instead of the reconstructed one cuts the RMSE advantage of MV over PI to 2.6% on Dataset A and 1.5% on Dataset B. The explanation offered above, that the mechanical bandwidth of the rotor becomes the binding constraint, can be made quantitative from the parameters of Table 2 alone, and is set out here because Reviewer A is right that the claim was asserted rather than shown.

The MPPT reference of eq. (11) maps wind speed to rotor speed with gain $\lambda_{opt}/R = 8.10 \text{ rad} \cdot \text{s}^{-1} \text{ per m/s}$. A single anemometer quantum of 2.77 m/s therefore commands a step of $\Delta\omega^* = 22.4 \text{ rad/s}$ in the reference, and the raw signal delivers such steps at the 1.5 s sampling interval, so following the reference would require a mean angular acceleration of $22.4/1.5 = 15.0 \text{ rad/s}^2$.

The acceleration actually available is bounded by the aerodynamic torque. At $V = 2.77 \text{ m/s}$, with $\beta = 0$ and the rotor at its MPPT speed of 22.4 rad/s , eq. (2) gives $P_m = 18.5 \text{ W}$ and hence $T_m = P_m/\omega_t = 0.824 \text{ N} \cdot \text{m}$; with $J_t = 0.30 \text{ kg} \cdot \text{m}^2$ the largest attainable acceleration is 2.75 rad/s^2 . The demand exceeds the supply by a factor of 5.4, and the rotor would need 8.2 s to reach the commanded speed from rest, more than five sampling intervals. This figure is computed at the highest wind speed present in the records and is therefore an upper bound: at the record mean speeds of 1.21 and 1.54 m/s the available acceleration falls to 0.52 and 0.85 rad/s^2 and the mismatch is correspondingly worse.

The same quantities give an effective mechanical time constant $\tau_{mech} = J_t\omega^*/T_m \approx 8.2 \text{ s}$ and a cut-off frequency of $1/(2\pi\tau_{mech}) \approx 0.019 \text{ Hz}$. The raw reference varies at up to 0.33 Hz, the Nyquist frequency of the 1.5 s sampling, roughly seventeen times above that cut-off. The five-sample moving average used for reconstruction has its first null at $1/(5 \times 1.5) = 0.13 \text{ Hz}$ and so removes most of that excess. For completeness, the drivetrain compliance mode at $\omega_n = \sqrt{K_s(1/J_t + 1/J_g))} = 73.0 \text{ rad/s}$, or 11.6 Hz, lies three orders of magnitude above the disturbance and plays no part in the argument: the limit here is inertial, not structural.

The mechanism is therefore not algorithmic. When the reference moves an order of magnitude faster than the plant can respond, both control laws run against the same inertial bound, the tracking error is dominated by that bound rather than by the law, and the difference between the laws is

compressed to the 2.6% and 1.5% that survive. The practical reading is that adaptive pitch control repays its complexity only while the wind disturbance evolves on a time scale the turbine can physically follow.

The argument also yields a sensor specification, which is the more useful product. For the reference to remain followable, a wind-speed quantum ΔV must satisfy $\lambda_{opt}\Delta V/R \leq a_{max}\Delta t$, that is $\Delta V \leq R \cdot a_{max} \cdot \Delta t / \lambda_{opt}$. With $a_{max} = 2.75 \text{ rad/s}^2$ and $\Delta t = 1.5 \text{ s}$ this gives $\Delta V \leq 0.51 \text{ m/s}$. The anemometer used here, with its 2.77 m/s quantum, is about five times too coarse for the rotor it drives, and an instrument resolving 0.5 m/s or better at the same update rate is the minimum acceptable for this class of turbine. Equivalently, holding the resolution fixed, the update interval would have to satisfy $\Delta t \geq \lambda_{opt}\Delta V/(R \cdot a_{max}) = 8.2 \text{ s}$, which is the mechanical time constant again.

I. Computational Cost and Embedded Feasibility

Because the target is a small turbine in a distributed installation, the controller must run on an embedded processor rather than on a workstation. The cost is estimated here by operation count, and is presented as an estimate rather than a measurement.

Per control sample the identifier and controller perform: the forward pass of the model, comprising $W1\phi$ ($m \times n = 12$ multiply-accumulates), three sigmoid evaluations, $W2\Gamma$ ($n \times m = 12$), the addition of θ (4) and the inner product that forms the prediction (4), about 32 operations in total; the RLS update of eqs. (16) and (17) with $n = 4$, dominated by the covariance update, about 60 operations and one division; ten inner gradient steps on the MLP weights, each a forward and a backward sweep of roughly 70 operations, about 700; and the control law of eq. (25), 6 operations and one division. The total is of order 800 multiply-accumulates and 30 sigmoid evaluations per sample, against about 5 operations for the PI law of eq. (26).

TABLE 6
COMPUTATIONAL COST PER CONTROL SAMPLE

Quantity	QARXNN MV	Fixed-gain PI
Multiply-accumulate operations per sample (estimated)	≈ 800	≈ 5
Transcendental (sigmoid) evaluations per sample	≈ 30	0
Divisions per sample	2	0
Stored parameters (words)	31	2
Total working memory (bytes, single precision)	< 300	< 32
Estimated execution time, 168 MHz Cortex-M4 (μs)	< 100	< 1
Measured execution time (μs)	-	-
Processor duty at 1.5 s control interval (%)	< 0.01	≈ 0

The ratio is two orders of magnitude, but the absolute figure is small. On a 168 MHz Cortex-M4 with a single-precision floating-point unit, of the class of the STM32F407 used in the authors' laboratory hardware, 800 multiply-accumulates execute in the order of 10 μ s, and with the sigmoid evaluated by table lookup the complete update stays well under 100 μ s. Against a control interval of 1.5 s the processor duty is below 0.01%. Working memory is likewise modest: 31 model parameters, 16 words for the RLS covariance and about 20 for the regression and error buffers, some 70 single-precision words or under 300 bytes, against the 192 kB of on-chip SRAM.

J. Position Relative to Other Adaptive Pitch-Control Schemes

The comparison reported above is against a fixed-gain PI baseline, which is the standard of practice on small turbines but is not the only alternative, and the manuscript should not imply otherwise. Gain-scheduled pitch control [3] adapts the gains to a measured operating point, but the schedule must be designed off line from a model that the present turbine does not have. Adaptive PI and self-tuning regulators [4] adapt on line while retaining a linear plant description, so they cannot represent the joint dependence of C_p on λ and β that the quasi-ARX coefficients carry. Model predictive control [5] optimises over a horizon and handles the pitch saturation of Table 2 explicitly, at a computational cost several orders above the figures of the preceding subsection. Direct neural-network control [8,11] dispenses with the linear structure altogether but forfeits the closed-form control law that makes the quasi-ARX formulation attractive on an embedded target. The scheme studied here sits deliberately between these: a nonlinear model is identified on line, yet the control law stays explicit.

Two adaptive baselines have been added so that the comparison separates two questions the fixed-gain baseline alone conflates: whether the improvement comes from adapting at all, or from the nonlinear quasi-ARX model in particular. The first is a self-tuning PI whose gains are recomputed at every sample from the same recursive least-squares estimate that drives the linear branch of eq. (13), so it carries no neural core and almost no additional computation. Writing the identified model as $A(q^{-1})y = B(q^{-1})u$ with $A = 1 + \hat{a}_1q^{-1} + \hat{a}_2q^{-2}$ and $B = \hat{b}_1q^{-1} + \hat{b}_2q^{-2}$, and the controller in velocity form as $(1 - q^{-1})u = (s_0 + s_1q^{-1})e$, the closed-loop characteristic polynomial is $A(q^{-1})(1 - q^{-1}) + B(q^{-1})S(q^{-1})$. That polynomial is of third degree while S supplies only two free parameters, so exact placement is not available and s_0 and s_1 are obtained by least squares against the target $(1 - pq^{-1})^3$ with $p = \exp(-\Delta t/\tau_{cl})$ and $\tau_{cl} = 8.2$ s, the rotor mechanical time constant. Because eq. (26) carries an explicit leading minus, the gains follow as $Kp = s_1$ and $Ki = -(s_0 + s_1)/\Delta t$.

Four safeguards, all standard in self-tuning regulation, are applied so that the baseline is not weakened by estimator transients on records this short, and they are stated here so that

its fairness can be judged: no adaptation during the first five samples while the estimator settles; rejection of any gain pair that is not finite and positive, in which case the previous pair is held; exponential blending of an accepted pair with the previous one at a rate of 0.3; and a hard clamp at $Kp = 100$ and $Ki = 50$. The fraction of samples on which the gains had to be held is reported below, since it measures how often the input gain $\hat{b}_1$ passes close enough to zero for the plant to become unidentifiable, which is the same low-wind condition approach.

The second baseline is a gain-scheduled PI in which the nominal gains are scaled by $(V_d/V)^3$ and clipped to one decade in each direction, following the cubic dependence of aerodynamic power on wind speed in eq. (2). It adapts to the operating point but uses no on-line identification at all, and so separates the value of measurement-driven scheduling from the value of estimation. Table 7 reports all four controllers on both records.

TABLE 7
COMPARISON AGAINST ADAPTIVE BASELINES: TRACKING RMSE (RAD/S)
AND IMPROVEMENT OVER THE FIXED-GAIN PI (%)

Controller	Dataset A	Dataset B
Fixed-gain PI	4.56 / 0.0	2.92 / 0.0
Gain-scheduled PI	- / -	- / -
Self-tuning PI (RLS)	- / -	- / -
QARXNN minimum-variance	4.15 / 9.1	2.77 / 5.2
Gains held, %	-	-

The reading of Table 7 is set out in advance so that it does not become a matter of interpretation after the fact. If the self-tuning PI recovers most of the margin that the minimum-variance controller holds over the fixed-gain baseline, then what the data support is a claim about adaptation rather than about the quasi-ARX model, and this paper will say so. If it does not, the nonlinear model is doing work that on-line adaptation alone does not, and that is the strongest form the central claim of this study can take. Until those results are available the claim made here is confined to what has been tested, namely an improvement over a tuned fixed-gain PI on two field records.

K. Scope of the Claims and Limitations

Four limits should be stated plainly rather than left to be inferred. The first concerns terminology. This study is a closed-loop simulation driven by field-measured wind, not a field validation of the controller. The wind data are genuine measurements; the turbine, drivetrain, generator and pitch actuator are models, and the control law has never been executed on hardware in real time. The phrase "field validation" has accordingly been removed from the abstract, keywords and conclusion of this revision and replaced by "field-data-driven closed-loop simulation". Real-time hardware-in-the-loop testing, in which the controller runs on

the embedded target against a real-time plant model, is the natural next step and is not claimed here.

The second concerns coverage. The evidence base is two records of 91 s and 103 s from two sites in one season. Robustness across wind conditions, sustained low wind, intermittency, sudden gusts, elevated turbulence intensity, and changes of direction, has not been tested, and the improvements reported cannot be assumed to hold outside these two records.

The third concerns the nonlinear cores. The MLP, RBFN and Elman variants are compared structurally and not experimentally. The abstract of the previous version could be read as promising a closed-loop comparison of the three; it does not deliver one, and has been reworded so that it no longer implies otherwise.

The fourth concerns the plant. The parameters of Table 2 are illustrative values chosen to represent a small rotor, not measurements from a machine that has been built. Statements about absolute performance are therefore not warranted. What remains meaningful is the comparison between the two controllers, which face an identical plant, an identical reference and an identical wind sequence.

IV. CONCLUSION

A switching quasi-ARX neural network adaptive pitch-control scheme, previously exercised only against synthetic ARMA wind and multi-megawatt geometry, has been transferred to a 1 m-radius rotor and driven by wind measured in the field at Gili Raja and Mandangin Island, Madura, East Java, Indonesia. Three implementation issues were identified and settled: the sign convention of the generator torque in the drivetrain balance, whose correction moves an open-loop eigenvalue from $+0.156 \text{ s}^{-1}$ to -0.156 s^{-1} ; a reduction of model order from 112 to 31 free parameters, justified against the 62 and 70 available samples; and a first-principles re-derivation of the minimum-variance pitch law for the reduced structure.

In closed-loop simulation the controller reduced rotor-speed tracking RMSE by 9.1% and 5.2% relative to a tuned fixed-gain PI baseline while commanding markedly less pitch activity, and the linear/nonlinear switch selected the nonlinear branch for about a quarter of each record. Driven by the raw quantised anemometer signal the margin fell to 2.6% and 1.5%. That contraction was shown to follow from a rotor mechanical time constant of 8.2 s derived from the plant parameters, and the same argument yields a minimum anemometer resolution of about 0.5 m/s at a 1.5 s update rate for this class of turbine, which is the most directly usable outcome of the study.

These results support a narrower statement than a general superiority of the scheme on small wind turbines. What has been shown is that it retains a measurable advantage over a tuned PI baseline on two short field records from one wind regime, in simulation, with illustrative plant parameters, and that the advantage is conditional on the wind disturbance evolving on a time scale the rotor can physically follow. Generalisation is limited by the number and length of the

datasets, by the coarse anemometer quantisation, by the absence of hardware testing, and by the absence of comparison against adaptive or predictive baselines.

Further work follows directly from those limits: longer deployments across seasons and sites; hardware-in-the-loop and then bench testing on the embedded target; comparison against adaptive PI, gain-scheduled and model-predictive baselines; an experimental comparison of the MLP, RBFN and Elman cores once records are long enough to separate them; and a return to the higher-order structure ($n_u = 3$, $n_y = 4$) once the sample count supports it.

ACKNOWLEDGEMENT

The authors gratefully acknowledge the financial support of the Kementerian Pendidikan Tinggi, Sains, dan Teknologi Republik Indonesia (Kemendikristek) under the BIMA Penelitian Dosen Pemula programme. The work was carried out under Contract No. 194/C3/DT.05.00/PL-BARU/2026 and No. 111/PL33.310/LT/2026.

REFERENCES

- [1] Shahbaz, M. and Lean, H.H.: The dynamics of electricity consumption and economic growth: A revisit study of their causality in Pakistan. *Energy*, 39(1): 146-153, 2012. DOI: 10.1016/j.energy.2012.01.048
- [2] Aguilera, R.F., Narayanan, S.S. and Aguilera, R.: Environmental footprint of the energy transition: A critical review. *Renewable and Sustainable Energy Reviews*, 171: 113009, 2023. DOI: 10.1016/j.rser.2022.113009
- [3] Bianchi, F.D., De Battista, H. and Mantz, R.J.: *Wind Turbine Control Systems: Principles, Modelling and Gain Scheduling Design*. Springer, London, UK, 2007
- [4] Åström, K.J. and Wittenmark, B.: *Adaptive Control*. 2nd Ed., Addison-Wesley, Reading, MA, USA, 1995
- [5] Shi, L., Dai, J. and Zhao, Q.: Multiple model predictive control for wind turbine with uncertainties. *Renewable Energy*, 168: 1-14, 2021. DOI: 10.1016/j.renene.2020.12.045
- [6] Oyague, F.: Gearbox Modeling and Load Simulation of a Baseline 750-kW Wind Turbine Using State-of-the-Art Simulation Codes. *NREL Technical Report TP-500-41160*, National Renewable Energy Laboratory, Golden, CO, USA, 2009
- [7] Chen, Z. and Blaabjerg, F.: Wind power – a review and research directions. *IEEE Industrial Electronics Magazine*, 3(2): 10-22, 2009. DOI: 10.1109/MIE.2009.932491
- [8] Narendra, K.S. and Parthasarathy, K.: Identification and control of dynamical systems using neural networks. *IEEE Transactions on Neural Networks*, 1(1): 4-27, 1990. DOI: 10.1109/72.80202
- [9] Ljung, L.: *System Identification: Theory for the User*. 2nd Ed., Prentice-Hall, Upper Saddle River, NJ, USA, 1999
- [10] Kukreja, S.L., Löfberg, J. and Brenner, M.J.: A least absolute shrinkage and selection operator (LASSO) for nonlinear system identification. *IFAC Proceedings Volumes*, 39(1): 814-819, 2006. DOI: 10.3182/20060329-3-AU-2901.00131
- [11] Jafarian, M. and Ranjbar, A.M.: Fuzzy modeling techniques and artificial neural networks to estimate annual energy output of a wind turbine. *Renewable Energy*, 35(9): 2008-2014, 2010. DOI: 10.1016/j.renene.2010.02.001
- [12] Shen, X., Zhu, X. and Du, Z.: Wind turbine aerodynamics and loads control in wind shear flow. *Energy*, 36(3): 1424-1434, 2011. DOI: 10.1016/j.energy.2011.01.028
- [13] Leung, D.Y.C. and Yang, Y.: Wind energy development and its environmental impact: A review. *Renewable and Sustainable Energy Reviews*, 16(1): 1031-1039, 2012. DOI: 10.1016/j.rser.2011.09.024

- [14] Bukala, J., Damaziak, K., Kroszczynski, K., Krzeszowiec, M. and Malachowski, J.: Investigation of parameters influencing the efficiency of small wind turbines. *Journal of Wind Engineering and Industrial Aerodynamics*, 146: 29-38, 2015. DOI: 10.1016/j.jweia.2015.06.017
- [15] Sugiyono, R.A. and Santoso, B.: Development of automatic weather station (AWS) based on Arduino microcontroller in Indonesia. *Journal of Physics: Conference Series*, 1528: 012047, 2020. DOI: 10.1088/1742-6596/1528/1/012047
- [16] Hau, E.: *Wind Turbines: Fundamentals, Technologies, Application, Economics*. 3rd Ed., Springer, Berlin, Germany, 2013
- [17] Malvoni, M., Leggieri, A., Maggiotto, G., Congedo, P.M. and De Giorgi, M.G.: Long-term performance, losses and efficiency analysis of a 960 kWp photovoltaic system in the Mediterranean climate, Italy. *Energy Conversion and Management*, 145: 169-181, 2017. DOI: 10.1016/j.enconman.2017.04.062
- [18] Gu, H., Li, S. and Zhao, H.: A switching quasi-ARX neural network identification scheme for nonlinear systems with a linear/nonlinear switching mechanism. *Neurocomputing*, 262: 20-32, 2017. DOI: 10.1016/j.neucom.2017.03.069
- [19] Wahyudi, I., Cahyo, N. and Setiawan, E.: Design and implementation of automatic weather station using Arduino-based sensor for coastal monitoring in Indonesia. *IOP Conference Series: Earth and Environmental Science*, 389: 012024, 2019. DOI: 10.1088/1755-1315/389/1/012024
- [20] Handoko, S.U., Setijadi, E. and Affandi, A.: Comparison of wind speed prediction methods in East Java using neural network and ARIMA. *EMITTER International Journal of Engineering Technology*, 7(1): 1-14, 2019. DOI: 10.24003/emitter.v7i1.370
- [21] Muljadi, E. and Butterfield, C.P.: Pitch-controlled variable-speed wind turbine generation. *IEEE Transactions on Industry Applications*, 37(1): 240-246, 2001. DOI: 10.1109/28.903156
- [22] Okulov, V.L. and Sorensen, J.N.: Stability of helical tip vortices in a rotor far wake. *Journal of Fluid Mechanics*, 576: 1-25, 2007. DOI: 10.1017/S0022112006004228
- [23] Baktiar, A., Syukri, M. and Munir, A.: Assessment of wind energy potential in Indonesia for small wind turbines using NASA meteorological data. *International Journal of Renewable Energy Research*, 11(2): 668-678, 2021
- [24] Purba, N.P., Kelvin, J., Annisaa, S., Faizal, C.K., Daulat, A., Martasuganda, M.K., Subhan, M.A. and Pranowo, W.S.: Suitable locations of ocean renewable energy (ORE) in Indonesia: Assessment of wind and wave energy. *Energy Reports*, 4: 70-78, 2018. DOI: 10.1016/j.egyr.2017.12.003
- [25] Abu Jami'in, M., Sutrisno, I. and Hu, J.: Lyapunov learning algorithm for quasi-ARX neural network to identification of nonlinear dynamical system. *IEEE International Conference on Systems, Man, Cybernetics*, pp. 14-17, 2012
- [26] Sutrisno, I., Abu Jami'in, M. and Hu, J.: Neural predictive controller of nonlinear systems based on quasi-ARX neural network. *Proceedings of the 18th International Conference on Automation & Computing*, Loughborough University, Leicestershire, UK, 2012
- [27] Abu Jami'in, M., Sutrisno, I. and Hu, J.: Deep searching for parameter estimation of the linear time invariant (LTI) system by using quasi-ARX neural network. *Proceedings of International Joint Conference on Neural Networks*, Dallas, Texas, USA, 2013
- [28] Sutrisno, I., Abu Jami'in, M. and Hu, J.: Implementation of Lyapunov learning algorithm for fuzzy switching adaptive controller modeled under quasi-ARX neural network. *2nd International Conference on Measurement, Information and Control*, ISSN 978-1-4799-1390-9, 2013
- [29] Sutrisno, I., Abu Jami'in, M. and Hu, J.: Modified fuzzy adaptive controller applied to nonlinear systems modeled under quasi-ARX neural network. *Artificial Life and Robotics*, 19: 22-26, 2014. ISSN 1015-013-0137-6
- [30] Sutrisno, I., Abu Jami'in, M. and Hu, J.: An improved Elman neural network controller based on quasi-ARX neural network for nonlinear systems. *IEEJ Transactions on Electrical and Electronic Engineering*, 9: 494-501, 2014
- [31] Sutrisno, I., Abu Jami'in, M., Hu, J. and Hamiruce Marhaban, M.: Nonlinear model-predictive control based on quasi-ARX radial-basis function-neural-network. *8th Asia Modelling Symposium*, ISSN 978-1-4799-6487-1, 2014
- [32] Abu Jami'in, M., Sutrisno, I., Hu, J., Bin Mariun, N. and Hamiruce Marhaban, M.: An adaptive predictive control based on a quasi-ARX neural network model. *13th International Conference on Control, Automation, Robotics and Vision (ICARCV)*, Marina Bay Sands, ISSN 978-1-4799-5199-4, 2014
- [33] Abu Jami'in, M., Sutrisno, I. and Hu, J.: Maximum power tracking control for a wind energy conversion system based on a quasi-ARX neural network model. *IEEJ Transactions on Electrical and Electronic Engineering*, 10: 368-375, 2015
- [34] Abu Jami'in, M., Sutrisno, I. and Hu, J.: Nonlinear adaptive control for wind energy conversion systems based on quasi-ARX neural networks model. *Proceedings of the International MultiConference of Engineers and Computer Scientists*, 1: 12-14, 2015
- [35] Abu Jami'in, M., Sutrisno, I. and Hu, J.: The state-dynamic-error-based switching control under quasi-ARX neural network model. *The Twentieth International Symposium on Artificial Life and Robotics*, 2015
- [36] Abu Jami'in, M., Hu, J., Hamiruce Marhaban, M., Sutrisno, I. and Bin Mariun, N.: Quasi-ARX neural network based adaptive predictive control for nonlinear systems. *IEEJ Transactions on Electrical and Electronic Engineering*, 11: 83-90, 2016
- [37] Abu Jami'in, M., Hu, J. and Julianto, E.: A Lyapunov based switching control to track maximum power point of WECS. *International Joint Conference on Neural Networks (IJCNN)*, ISSN 978-1-5090-0620-5, 2016
- [38] Sutrisno, I., Abu Jami'in, M., Hu, J. and Hamiruce Marhaban, M.: A self-organizing quasi-linear ARX RBFN model for nonlinear dynamical systems identification. *SICE Journal of Control, Measurement, and System Integration*, 9(2): 070-077, 2016
- [39] Abu Jami'in, M., Yuyun and Julianto, E.: Hierarchical algorithms of quasi-linear ARX neural networks for identification of nonlinear systems. *Engineering Letters*, 25(3), 2017
- [40] Abu Jami'in, M., Anam, K., Rulaningtyas, R. and Echsony, M.: Hierarchical algorithm for the identification of parameter estimation of linear system. *3rd International Conference on Sustainable Information Engineering and Technology (SIET)*, 2018