Water hyacinth–derived biofilters for industrial air pollution mitigation and circular blue economy applications

A dual-scale filtration model, a bias-corrected sensing protocol and a simulation-based performance assessment

[1]Saswat Swain [2]Pravat Satpathy [3]Sonia Mohapatra
[1]Ph.D Scholar, [2,3]R & D Division
[1]KSOM - KIIT Deemed to be University ,[2,3]Tekstein Solutions,India
[1]saswat.mck@gmail.com [2]pravat@teksteinsolutions.com [3]sonia@teksteinsolutions.com

Abstract—Water hyacinth is an invasive freshwater macrophyte whose aerenchymatous stem tissue has been proposed as a biodegradable replacement for synthetic air filter media. Performance figures reported for such filters normally rest on two shortcuts: the medium is treated as a uniform mat of fibres, and removal efficiency is read off a pair of low cost gas sensors placed upstream and downstream. This work shows that both are unsafe. Aerenchyma is bimodal: open channels of the order of one hundred micrometres coexist with dense lignocellulosic walls, and the two act as parallel hydraulic paths. A dual-scale model is derived in which flow divides by the permeability of each phase, and at representative morphology about half the air bypasses the collecting fibres. The uniform description reproduces size-resolved efficiency to within roughly one percentage point, yet underestimates pressure drop by a factor of 2.4, so that the quality factor inferred from it is too high by a factor of 2.3 to 3.6. Under dust loading the walls block faster than the channels, the bypass fraction grows, and fine particle efficiency falls from 59.0 to 46.6 percent over ninety days while pressure drop rises by only 54 percent. A conventional filter behaves in the opposite sense, so a pressure-triggered replacement rule fires about four months too late. Paired metal oxide sensors are shown to amplify any multiplicative instrument error by a factor of about 2.4, so that a twenty percent mismatch in baseline resistance combined with humidity cross-sensitivity turns a true removal of eighteen percent into an apparent fifty four percent; a bypass cross-calibration protocol recovers the true value to within 2.3 percentage points. Machine learning on the resulting stream predicts remaining service life to about twenty three days against a mean life of sixty two days, and classifies saturation state at an accuracy of 0.68, well below values commonly quoted. A life cycle balance gives a net saving of seventy one kilograms of carbon dioxide equivalent per tonne of wet biomass, but the saving reverses if more than thirteen percent of the water is removed thermally. All results reported here are obtained by simulation.

Index Terms—Air filtration, aerenchyma, biomaterials, circular economy, gas sensors, invasive species, predictive maintenance, quality factor, water hyacinth.

I. Introduction

Eichhornia crassipes, commonly called water hyacinth, is among the most disruptive freshwater invasives in tropical Asia [1]. Dense floating mats shade the water column, strip dissolved oxygen, obstruct navigation and irrigation intakes, and provide breeding habitat for disease vectors. In the Indian state of Odisha alone the affected surface is of the order of fifteen thousand hectares, and mechanical clearance is funded annually without any recovery of value from the harvested material. The cleared biomass is typically dumped or burned, returning its carbon to the atmosphere and, where it decays under water, releasing methane.

At the same time, particulate air pollution in the same region imposes a large and well documented health burden [5]. Filtration media used to address it are almost entirely petrochemical: polypropylene and polyester nonwovens that are cheap, effective, and neither biodegradable nor recyclable. Closing these two loops, by converting harvested hyacinth into filter media whose spent form can be pyrolysed to biochar, has therefore attracted attention, alongside a wider interest in biopolymer filtration media [6].

The physical basis is plausible. The petiole is an aerenchyma, a gas-transport tissue in which large lacunae are separated by thin walls of cellulose, hemicellulose and lignin [2], [4]; the same walls are known to bind dissolved metals [3]. This gives a low bulk density, a high internal surface area and, in principle, a favourable ratio of capture to resistance. What is missing is a treatment that respects the structure. Published performance claims are, almost without exception, obtained by one of two routes: the medium is homogenised into an equivalent fibrous mat and classical single-fibre theory applied to it, or an efficiency figure is read from a pair of low cost metal oxide sensors on either side of the cartridge. Neither is examined critically, although the sensitivity of fibrous filters to structural inhomogeneity has been recognised for three decades [13].

This paper examines both. First, a dual-scale model is derived that treats the lacunae and the fibre walls as parallel hydraulic paths and partitions the flow between them by permeability, exposing a bypass fraction that homogenised models cannot represent. Second, the model is loaded with dust and predicts a decline of efficiency with service time, the reverse of conventional behaviour, which invalidates pressure-based maintenance triggers. Third, the inverse power law governing metal oxide sensors is shown to amplify instrument error enough to account for the removal efficiencies usually reported, and a correction protocol is given. Fourth, remaining service life prediction [18] is evaluated on a simulated cohort under a grouped and a distribution-shifted split, and bounded by an oracle given the unobservable internal state. The work is a modelling study; no laboratory measurements are reported. Parameter values are taken from published ranges for hyacinth tissue, industrial dust and the sensor family in question. The intention is to establish what should be measured, and how it should be read, before a prototype is characterised.

II. A dual-scale model of aerenchymatous filter media

A. Phase decomposition and flow partitioning

Consider a flat sheet of processed hyacinth media of thickness L presented to a face velocity U_0. The tissue is idealised as two interpenetrating phases occupying fixed area fractions of any cross-section normal to the flow, as sketched in Fig. 1. The lacunar phase consists of longitudinal channels of radius r_L and tortuosity τ, occupying an area fraction φ_L. The fibrous phase consists of the lignocellulosic walls between those channels, described as a classical fibrous filter of local packing density α_f and fibre diameter d_f, and occupying the remaining fraction.

The essential point is that the two phases are hydraulically in parallel, not in series. Both span the full thickness of the sheet, both experience the same pressure difference between the upstream and downstream faces, and the volumetric flow therefore divides between them in proportion to their permeabilities. Writing the lacunar permeability from Poiseuille flow in a bundle of tortuous capillaries and the fibrous permeability from the Davies correlation [10] gives Eq. 1 and Eq. 2. The fraction of the total volumetric flow carried by the lacunae, denoted β and referred to below as the bypass fraction, follows in Eq. 3, and the pressure drop of the composite in Eq. 4.

$$k_L = r_L^2 / (8\tau^2) \tag{1}$$

$$k_F = d_f^2 / [16\alpha_f^{3/2}(1 + 56\alpha_f^3)] \tag{2}$$

$$\beta = \varphi_L k_L / [\varphi_L k_L + (1 - \varphi_L)k_F] \tag{3}$$

$$\Delta p = \mu\, U_0 L / [\varphi_L k_L + (1 - \varphi_L)k_F] \tag{4}$$

Here μ is the dynamic viscosity of air. The bypass fraction β is the single number that a homogenised model cannot express, and it is the quantity that the remainder of this paper turns on.

B. Size-resolved penetration

Because the phases are in parallel, the downstream concentration is the flow-weighted mixture of the two streams, so penetration is additive in the flow fractions, as in Eq. 5. This is the structural difference from a homogenised description, and it has an immediate consequence: for particle sizes at which the open channels capture nothing, the overall efficiency cannot exceed $1 - \beta$ however dense the walls are made.

Within the fibrous phase, penetration follows the standard exponential form of Eq. 6, with the combined single-fibre efficiency of Eq. 7 assembled from the diffusion, interception, diffusion-interception, impaction and gravitational terms of Lee and Liu [8] and of Stechkina and Fuchs [9], evaluated at the Kuwabara hydrodynamic factor [7] for the local packing density and at the interstitial velocity of that phase. Within the lacunar phase, deposition is the product of the Gormley and Kennedy diffusional term for laminar tube flow [11], a settling term and an inertial term applied at the internal septa [12]. The quality factor of Eq. 8 [15] compares media at equal resistance.

$$P(d_p) = \beta\, P_L(d_p) + (1 - \beta)P_F(d_p) \tag{5}$$

$$P_F = \exp[- 4\alpha_f\eta_\Sigma L / \pi\, d_f(1 - \alpha_f)] \tag{6}$$

$$\eta_\Sigma = 1 - (1-\eta_D)(1-\eta_R)(1-\eta_{DR})(1-\eta_I)(1-\eta_G) \tag{7}$$

$$q_F = - \ln P / \Delta p \tag{8}$$

C. Loading dynamics

Loading is treated by transferring the captured mass into the geometry of each phase. Dust deposited in the walls raises the local packing density according to Eq. 9, while dust deposited in the channels reduces their open radius according to Eq. 10, both at a deposit bulk density reflecting the open dendritic structure of accumulated particulate [14].

The two phases do not degrade at the same rate. The walls collect far more mass per unit of their own volume than the channels, and their permeability falls as a strong power of packing density, whereas channel permeability falls only as the square of the radius. The bypass fraction therefore rises monotonically with loading, which is the mechanism behind the central prediction of Section V.

$$\mathrm{d}\alpha_f/\mathrm{d}t = (1-\beta)U_0\, \Sigma_k\, c_k[1-P_F(d_k)] / \rho_d(1-\varphi_L)L \tag{9}$$

$$r_L(t+\delta t) = r_L(t)[1 - \beta\, U_0\Sigma_k c_k(1-P_L)\delta t / \rho_d\varphi_L L]^{1/2} \tag{10}$$

In these expressions c_k is the mass concentration of dust in size bin k and ρ_d is the bulk density of the accumulated deposit.

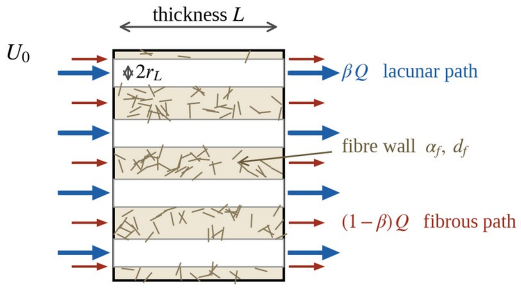

Fig. 1. Idealisation of processed hyacinth media as two parallel hydraulic phases. Open lacunae of radius r_L carry a fraction β of the volumetric flow Q; the intervening lignocellulosic walls, treated as a fibrous filter of packing density α_f and fibre diameter d_f, carry the remainder. Both phases span the full thickness L.

III. Sensor model and bias-corrected estimator

The second half of the analysis concerns how such a filter is instrumented. Metal oxide semiconductor sensors of the kind used in low cost air quality nodes respond to a reducing gas through a power law in concentration, with a resistance that also depends on humidity and temperature [16]. Writing the baseline resistance as R_0 and the environmental factor as g gives Eq. 11 and Eq. 12, the linearised coefficients being taken from datasheet families for this class.

Efficiency is normally obtained by inverting this law at both ports with a nominal calibration and taking the ratio. The inversion carries the exponent $-1/b$, so any multiplicative error in resistance, whether from the environmental factor or from a mismatch between the baseline resistances of the two devices, enters the inferred concentration raised to that power, as in Eq. 13. For the exponent typical of this family the amplification is about 2.4, so a twenty percent error in baseline resistance becomes a thirty five percent error in concentration.

Collecting both effects gives Eq. 14, in which κ is the ratio of the true baseline resistances of the downstream and upstream sensors. Neither term cancels. The environmental term survives because a dry hygroscopic bed lowers the humidity of the air passing through it, placing the two ports at different points on the humidity curve; the baseline term survives, and is the larger, because unit to unit spread in R_0 for this class runs to tens of percent [17].

The correction has two stages. Co-located humidity and temperature probes at both ports allow the environmental factor to be divided out. The baseline ratio is then estimated in situ by diverting the airflow around the cartridge for a short window each night, during which both ports sample the same air and the observed resistance ratio is, by construction, an estimate of κ alone, as in Eq. 15. Neither stage requires laboratory recalibration or sensor matching at manufacture.

$$R_s = R_0 a\, C^{-b} g(\mathrm{RH}, T) \tag{11}$$

$$g(\mathrm{RH}, T) = 1 + c_1(\mathrm{RH} - \mathrm{RH}_{ref}) + c_2(T - T_{ref}) \tag{12}$$

$$C_{est} = C\, (R_0/R_{0,nom})^{-1/b}\, g^{-1/b} \tag{13}$$

$$1 - \eta_{naive} = (1 - \eta_{true})\, \kappa^{-1/b}\, (g_{out}/g_{in})^{-1/b} \tag{14}$$

$$\kappa_{est} = \mathrm{median}_{t\in B}\, [(R_{s,out}/g_{out}) / (R_{s,in}/g_{in})] \tag{15}$$

IV. Simulation setup

Nominal parameters are listed in Table I and correspond to a moderately compressed cartridge cut from petiole tissue. The face velocity of 0.05 m s⁻¹ is 180 m³ h⁻¹ m⁻², within the range quoted for fan-assisted panels. The challenge aerosol is a bimodal industrial dust discretised into forty eight logarithmically spaced bins between twenty nanometres and twenty micrometres.

For the sensor study a thirty day deployment was synthesised at five minute resolution with diurnal humidity and temperature cycles, twin-peaked traffic-like gas loading, a true removal efficiency decaying as chemisorption sites on the fresh lignocellulose saturate, and a hygroscopic bed lagging the inlet humidity with a time constant of thirty six hours. The sensors were given a twenty percent baseline mismatch, and a thirty minute bypass window was inserted nightly.

For the prognostics study a cohort of two hundred and sixty virtual cartridges was generated by sampling morphology, thickness, face velocity, dust concentration, coarse mode diameter, fine mass fraction and deposit density from independent uniform or log-uniform ranges. End of life was a fifteen percent relative loss of fine particle efficiency. One hundred and eight units reached it within the two hundred day horizon and contributed 21409 observations; the remainder degraded too slowly to qualify, which is itself informative. Evaluation used a five-fold split grouped by cartridge, so no unit appears in both training and test sets, and then a distribution shift: training on the low dust half of the cohort, testing on the high dust half.

TABLE I. Nominal medium, aerosol and sensor parameters

Symbol	Quantity	Value
L	Media thickness	12 mm
d_f	Fibre diameter	18 µm
α_f	Packing density of wall	0.12
r_L	Lacuna radius	90 µm
φ_L	Lacunar area fraction	0.45
τ	Lacunar tortuosity	1.35
U_0	Face velocity	0.05 m s⁻¹
ρ_p, ρ_d	Particle, deposit bulk density	1500, 150 kg m⁻³
c	Dust mass concentration	500 µg m⁻³
—	Dust modes (CMD, GSD, mass)	0.25 µm, 1.8, 0.25; 4.0 µm, 2.0, 0.75
b	Sensor power-law exponent	0.42
c_1, c_2	Humidity, temperature cross-sens.	−3.5×10⁻³ %RH⁻¹, −8.0×10⁻³ K⁻¹

V. Results and discussion

A. Grade efficiency and quality factor

Fig. 2a compares the size-resolved efficiency of the dual-scale medium with that of the homogenised equivalent, formed by spreading the same solid volume uniformly over the same thickness at the same fibre diameter. The two curves are close. The most penetrating size shifts from 566 to 485 nanometres and the minimum efficiency changes from 37.5 to 38.8 percent. At 2.5 micrometres the homogenised model gives 98.0 percent against 92.6 percent for the dual-scale model. Judged on efficiency alone, the homogenised description would be considered adequate.

It is not. The pressure drops differ by a factor of 2.4, at 9.3 against 22.0 pascals, because concentrating the same solid into dense walls separated by channels of finite radius is hydraulically more expensive than spreading it uniformly. Fig. 2b gives the consequence: the homogenised model overstates quality factor by 2.3 at the most penetrating size and by 3.6 at 2.5 micrometres. Since quality factor propagates into fan energy, module sizing and operating cost, a model validated only against efficiency can be badly wrong about everything downstream of it. The dotted curve in Fig. 2a shows why the two efficiencies nevertheless agree: the open channels are not inert, and settling and inertial deposition at the septa strip coarse particles from the bypass stream even though submicron particles pass through it almost untouched.

The bypass fraction is strongly sensitive to morphology. Over channel radii from 50 to 160 micrometres and lacunar area fractions from 0.30 to 0.60, β ranges from 0.14 to 0.86. The upper end of that range corresponds to lightly processed tissue and represents a medium in which most of the air never meets a fibre.

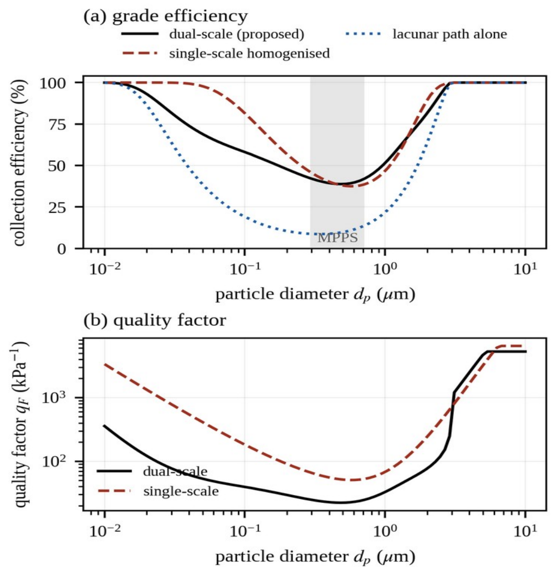

Fig. 2. (a) Size-resolved collection efficiency of the dual-scale medium, of the homogenised equivalent carrying the same solid volume, and of the lacunar path taken alone. (b) Quality factor of the same two media. The curves agree closely on efficiency and differ by a factor of 2.3 to 3.6 on quality factor.

B. Loading behaviour and maintenance triggers

Fig. 3a is the central result. Loaded at 500 µg m⁻³, the dual-scale medium accumulates 152 g m⁻² over ninety days. Over that period the bypass fraction rises from 0.506 to 0.707, fine particle efficiency falls from 59.0 to 46.6 percent, and pressure drop rises by only 54 percent. The homogenised medium, loaded identically, behaves in the familiar way: efficiency climbs from 64.8 to 94.9 percent as the growing deposit acts as an additional collector, and pressure drop reaches twice its clean value after thirty four days.

The practical consequence concerns maintenance policy. Condition monitoring for filtration modules is conventionally based on differential pressure, and the concept documents for systems of this kind propose exactly that. If a doubling of pressure drop is used as the replacement trigger, the dual-scale medium reaches it on day 194, by which time its fine particle efficiency has fallen to 39.1 percent, a third below its clean value. A trigger set instead at a fifteen percent relative loss of efficiency fires on day 61, when pressure drop has risen by only forty percent. The two rules differ by more than four months. For a medium of this class, differential pressure is not a proxy for effectiveness, and a monitoring architecture that relies on it will report a healthy filter long after the filter has stopped working. The prediction follows from the parallel-path structure and the differing rate at which the two phases lose permeability, not from the particular correlations chosen; it has not, however, been measured, and Section VI gives the experiment that would falsify it.

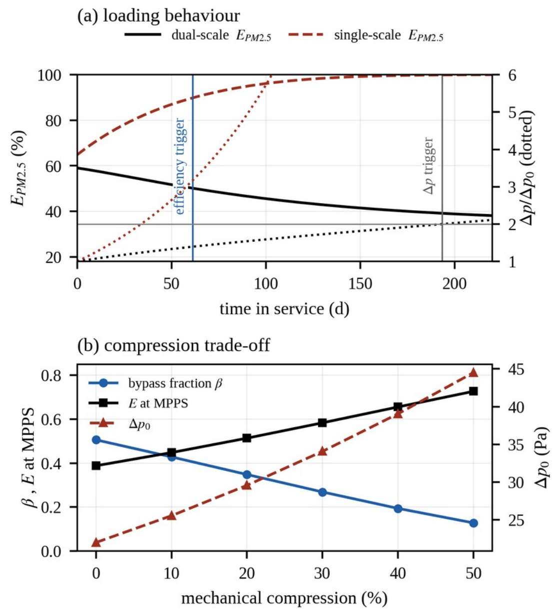

Fig. 3. (a) Fine particle efficiency (solid and dashed, left axis) and normalised pressure drop (dotted, right axis) through a loading cycle at 500 μg m^{-3}. Vertical lines mark a fifteen percent relative efficiency loss and a doubling of pressure drop for the dual-scale medium. (b) Effect of mechanical compression on bypass fraction, efficiency at the most penetrating size and clean pressure drop.

C. Compression as the principal design lever

If bypass is the problem, mechanical compression is the obvious lever, since it collapses the lacunae and raises the local packing density of the walls. Table II and Fig. 3b quantify the trade-off. Compressing the sheet by half reduces β from 0.506 to 0.128 and raises efficiency at the most penetrating size from 38.8 to 72.6 percent, at the cost of doubling clean pressure drop from 22.0 to 44.5 pascals. Expressed as quality factor the exchange is close to neutral, but the operating point is far more useful: a filter that removes seventy percent of submicron aerosol at 45 pascals is a deployable panel, whereas one that removes thirty nine percent at 22 pascals is not. Compression also delays the bypass drift described above, since it starts the medium further from the regime in which the channels dominate.

This suggests that the processing step, which in the concept literature is treated purely as a hygiene and drying operation, is in fact the principal performance control available, and should be specified against a target bypass fraction rather than against a target density.

TABLE II. Effect of mechanical compression on the nominal medium

Compression	β	Δp_0 (Pa)	E at MPPS (%)	E at 2.5 μm (%)
0 %	0.506	22.0	38.8	92.6
10 %	0.428	25.6	44.8	97.4
20 %	0.348	29.6	51.3	> 99.9
30 %	0.269	34.1	58.3	> 99.9
40 %	0.194	39.0	65.5	> 99.9
50 %	0.128	44.5	72.6	> 99.9

D. Sensor bias and its correction

Fig. 4a shows the thirty day sensor deployment. The true removal efficiency decays from 30 to 9 percent as the chemisorption sites saturate, averaging 15.9 percent over the month. The naive paired estimator reports 57.3 percent on the first day and 45.9 percent on average. This is the same order as the sixty five to seventy five percent quoted for systems of this kind, and Fig. 4b decomposes it. Humidity cross-sensitivity alone lifts an eighteen percent true removal to twenty nine percent apparent; a twenty percent baseline resistance mismatch alone lifts it to forty seven percent; together they give fifty four percent.

The two stages of the correction behave differently, and the difference is instructive. Compensating for humidity and temperature reduces the scatter of the estimate by about a third, from 0.064 to 0.042, but leaves the bias essentially untouched at 29.6 percentage points. It improves precision and not accuracy, because the dominant error is a fixed multiplicative offset between the two sensors and not a fluctuating environmental term. Adding the nightly bypass cross-calibration recovers the baseline ratio to 1.203 against a true value of 1.200 and reduces the bias to minus 0.6 percentage points, with a mean absolute error of 2.3 percentage points. Accuracy is restored only by the second stage.

In general, any efficiency obtained by differencing two uncalibrated inverse-power-law sensors is biased upward, the bias grows as the exponent falls, and environmental compensation alone will not remove it. Figures reported for biofilters instrumented this way should be read as upper bounds until a cross-calibration is documented.

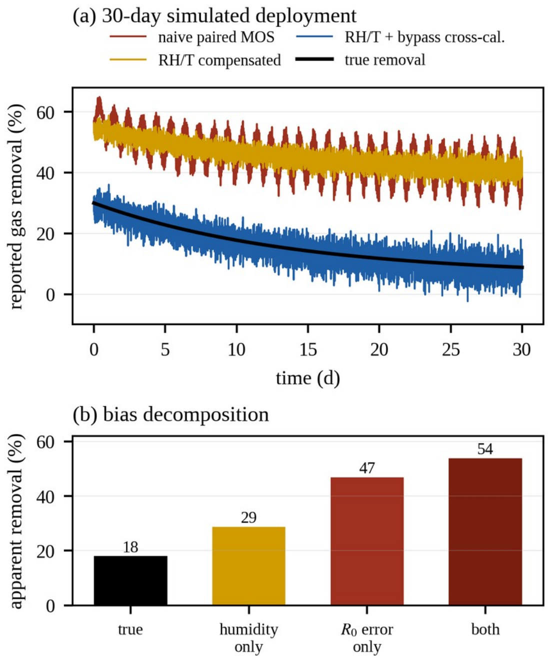

Fig. 4. (a) Reported gas removal over a thirty day simulated deployment for the naive paired estimator, for humidity and temperature compensation alone, and for the full protocol including nightly bypass cross-calibration, against the true removal. (b) Decomposition of the bias for a true removal of eighteen percent.

E. Remaining useful life and saturation state

Table III reports remaining useful life prediction. The pressure-threshold heuristic, which extrapolates linearly to a doubling of clean pressure drop, is useless, with a mean absolute error of 171 days against a mean life of 62 days; this is the quantitative form of the argument in Section V-B. Learned models do far better but are not accurate. Ridge regression on the raw stream reaches 24.0 days and gradient boosting 23.8 days; adding the dimensionless loading index and the loading-susceptibility ratio derived from Section II changes nothing distinguishable from fold-to-fold variation.

Two comparisons bound the problem. A paired particle counter, so that current efficiency is observed rather than inferred, improves the error only to 22.6 days. An oracle given the true internal state, that is the instantaneous bypass fraction and wall packing density that no instrument can measure, reaches only 20.8 days. Since perfect knowledge of the present buys thirteen percent, the residual error is dominated by the unpredictability of future dust exposure rather than by any deficiency of sensing or model class, and no additional instrumentation will close it.

Under distribution shift the ranking inverts. Trained on the low dust half of the cohort and tested on the high dust half, ridge regression degrades gracefully to 24.7 days with a positive coefficient of determination, whereas gradient boosting collapses to 45.0 days at minus 1.23, and to 65.8 days on raw features. The flexible model has fitted the loading regime it was shown rather than the mechanism, so where cartridges go to sites of unknown severity the linear model is the safer choice despite being marginally worse in-distribution.

Classifying saturation state into healthy, caution and replace categories reached an accuracy of 0.655 on the raw stream, 0.666 with the physics features and 0.683 with a paired particle counter, against a majority class baseline of 0.566 and a best macro F1 of 0.612. Accuracies above 0.85 are frequently quoted for tasks of this description; the present results suggest such figures are unlikely to survive an evaluation in which units are held out as groups and the class boundaries are set by efficiency rather than elapsed time.

TABLE III. Remaining useful life prediction, mean life 62.1 d. Unmarked columns are five-fold cross-validation grouped by cartridge; columns marked † are the distribution-shifted split (trained on the low dust half of the cohort, tested on the high dust half).

Predictor	**MAE (d)**	**R^2**	**MAE† (d)**	**R^2†**
Δp-threshold heuristic	171.3	—	—	—
Ridge, raw stream	24.0	0.569	25.6	0.094
Ridge, + physics features	24.4	0.564	24.7	0.190
Gradient boosting, raw	23.8	0.548	65.8	−3.42
Gradient boosting, + physics	23.7	0.534	45.0	−1.23
Gradient boosting, + paired PM	22.6	0.562	40.1	−0.85
Gradient boosting, oracle state	20.8	0.608	—	—

F. Circular economy balance

Table IV gives the greenhouse gas balance for one tonne of wet harvested biomass, which yields 70 kg of dry matter, 49 kg of filter media covering about 41 m² of panel, and 14.7 kg of biochar at end of life; the net figure is a saving of 71 kg of carbon dioxide equivalent. Monte Carlo propagation over moisture, dryer efficiency, grid intensity, media yield, anaerobic fraction [20], char yield, permanence [19] and haulage distance gives a mean of minus 71.5, a fifth to ninety fifth percentile range of minus 141 to minus 15, and a 0.998 probability that the balance is negative.

Two caveats qualify this. First, the avoided methane term dominates the credit side and is the least certain quantity in the inventory, since it depends on how much of the harvested material would in fact have decayed anaerobically. Excluding it entirely leaves minus 5.6 kg, essentially neutral. The climate case therefore rests principally on diverting biomass from anaerobic decay, and only secondarily on displacing polypropylene or storing carbon as char.

Second, drying is the binding constraint. Latent heat dominates the energy inventory, and if water is removed thermally the balance reverses. The break-even point is thirteen percent: if more than that share of the 930 kg of water in the tonne is driven off by fuel or grid electricity rather than by mechanical pressing and ambient drying, the process becomes a net emitter. The nominal case removes 1.8 percent thermally and sits comfortably inside the constraint, but any design substituting a continuous dryer for a press and a drying yard will not.

TABLE IV. Greenhouse gas balance per tonne of wet biomass

Term	**kg CO_2-eq**
Electricity, press and forming, 22.7 kWh	+16.1
Thermal drying and curing, 77.5 MJ	+15.3
Road transport, 60 km	+6.6
Avoided anaerobic methane	−65.1
Biochar carbon storage	−26.3
Avoided polypropylene media	−17.6
Net	**−71.0**

VI. Limitations

The principal limitation is that this is a simulation. The dual-scale model has not been validated against measured grade efficiency or measured pressure drop for real hyacinth media, and the loading prediction in particular is a structural consequence of the model rather than an observation. The experiment that would test it is straightforward: load a cartridge on a standard test rig and measure grade efficiency at intervals rather than only at the start. If efficiency at the most penetrating size falls monotonically while pressure drop rises slowly, the model is supported; if efficiency rises, the parallel-path description is wrong and the medium is better treated as an inhomogeneous but singly connected mat.

Several idealisations should be noted. The lacunae are represented as straight capillaries of a single radius with a lumped tortuosity, whereas real aerenchyma has a distribution of channel sizes and irregular septa. The deposit is assumed uniform through the thickness of each phase, which will underestimate the rate of pressure rise once a surface cake forms. Hygroscopic swelling of the wall, which alters both permeability and the flow split, is not modelled and is likely to matter in monsoon conditions. Biological growth on a moist organic substrate is a deployment risk this analysis does not address. Finally, the sensor model treats a single species with a single power law, whereas real devices respond to a mixture with species-dependent exponents, so the correction protocol is necessary rather than sufficient.

VII. Conclusion

Aerenchymatous plant tissue is not a fibrous mat, and treating it as one gives a description that looks acceptable when judged on efficiency yet is wrong about resistance, quality factor and ageing. Modelling the tissue as two parallel hydraulic phases exposes a bypass fraction near one half, predicts that efficiency falls rather than rises with dust loading, and shows that a pressure-based replacement trigger fires roughly four months after the medium has ceased to meet its specification. Mechanical compression is the principal design lever, halving the bypass fraction and nearly doubling submicron efficiency for a doubling of clean resistance.

On the instrumentation side, paired metal oxide sensors amplify multiplicative error by the reciprocal of their power law exponent, enough to convert a true removal of eighteen percent into an apparent fifty four percent. Environmental compensation improves precision only; accuracy requires an in situ cross-calibration, for which a nightly bypass window suffices. Prognostic accuracy is bounded not by sensing or model class but by the unpredictability of future load, and reported classification accuracies appear optimistic relative to a grouped, efficiency-referenced evaluation.

The circular economy case survives this scrutiny at about seventy kilograms of carbon dioxide equivalent saved per tonne of wet biomass, but rests almost entirely on avoided anaerobic decay and fails if drying is thermal. Future work should be experimental: grade efficiency measured through a loading cycle, on media prepared at several compression ratios, with the sensor chain cross-calibrated by the protocol described here.

Acknowledgment

The authors thank the Research and Development Division of Tekstein Solutions Private Limited for supporting this work. The simulation code and the scripts that generate every figure and table in this paper are available from the corresponding author.

References

[1] A. M. Villamagna and B. R. Murphy, “Ecological and socio-economic impacts of invasive water hyacinth (Eichhornia crassipes): a review,” Freshwater Biology, vol. 55, no. 2, pp. 282–298, 2010.
[2] M. B. Jackson and W. Armstrong, “Formation of aerenchyma and the processes of plant ventilation in relation to soil flooding and submergence,” Plant Biology, vol. 1, no. 3, pp. 274–287, 1999.
[3] C. Mahamadi, “Water hyacinth as a biosorbent: a review,” African Journal of Environmental Science and Technology, vol. 5, no. 13, pp. 1137–1145, 2011.
[4] J. N. Nigam, “Bioconversion of water hyacinth (Eichhornia crassipes) hemicellulose acid hydrolysate to motor fuel ethanol by xylose-fermenting yeast,” Journal of Biotechnology, vol. 97, no. 2, pp. 107–116, 2002.
[5] A. Pandey, M. Brauer, M. L. Cropper, K. Balakrishnan, P. Mathur, S. Dey, et al., “Health and economic impact of air pollution in the states of India: the Global Burden of Disease Study 2019,” The Lancet Planetary Health, vol. 5, no. 1, pp. e25–e38, 2021.
[6] H. Souzandeh, Y. Wang, A. N. Netravali, and W.-H. Zhong, “Towards sustainable and multifunctional air-filters: a review on biopolymer-based filtration materials,” Polymer Reviews, vol. 59, no. 4, pp. 651–686, 2019.
[7] S. Kuwabara, “The forces experienced by randomly distributed parallel circular cylinders or spheres in a viscous flow at small Reynolds numbers,” Journal of the Physical Society of Japan, vol. 14, no. 4, pp. 527–532, 1959.
[8] K. W. Lee and B. Y. H. Liu, “Theoretical study of aerosol filtration by fibrous filters,” Aerosol Science and Technology, vol. 1, no. 2, pp. 147–161, 1982.
[9] I. B. Stechkina and N. A. Fuchs, “Studies on fibrous aerosol filters I. Calculation of diffusional deposition of aerosols in fibrous filters,” Annals of Occupational Hygiene, vol. 9, no. 2, pp. 59–64, 1966.
[10] C. N. Davies, Air Filtration. London: Academic Press, 1973.
[11] P. G. Gormley and M. Kennedy, “Diffusion from a stream flowing through a cylindrical tube,” Proceedings of the Royal Irish Academy A, vol. 52, pp. 163–169, 1949.
[12] W. C. Hinds, Aerosol Technology: Properties, Behavior, and Measurement of Airborne Particles, 2nd ed. New York: Wiley, 1999.
[13] E. Schweers and F. Löffler, “Realistic modelling of the behaviour of fibrous filters through consideration of filter structure,” Powder Technology, vol. 80, no. 3, pp. 191–206, 1994.
[14] D. Thomas, P. Penicot, P. Contal, D. Leclerc, and J. Vendel, “Clogging of fibrous filters by solid aerosol particles: experimental and modelling study,” Chemical Engineering Science, vol. 56, no. 11, pp. 3549–3561, 2001.
[15] J. Wang, S. C. Kim, and D. Y. H. Pui, “Investigation of the figure of merit for filters with a single nanofiber layer on a substrate,” Journal of Aerosol Science, vol. 39, no. 4, pp. 323–334, 2008.
[16] C. Wang, L. Yin, L. Zhang, D. Xiang, and R. Gao, “Metal oxide gas sensors: sensitivity and influencing factors,” Sensors, vol. 10, no. 3, pp. 2088–2106, 2010.
[17] A. C. Rai, P. Kumar, F. Pilla, A. N. Skouloudis, S. Di Sabatino, C. Ratti, et al., “End-user perspective of low-cost sensors for outdoor air pollution monitoring,” Science of the Total Environment, vol. 607–608, pp. 691–705, 2017.
[18] Y. Lei, N. Li, L. Guo, N. Li, T. Yan, and J. Lin, “Machinery health prognostics: a systematic review from data acquisition to RUL prediction,” Mechanical Systems and Signal Processing, vol. 104, pp. 799–834, 2018.
[19] D. Woolf, J. E. Amonette, F. A. Street-Perrott, J. Lehmann, and S. Joseph, “Sustainable biochar to mitigate global climate change,” Nature Communications, vol. 1, art. 56, 2010.
[20] IPCC, 2006 IPCC Guidelines for National Greenhouse Gas Inventories, Volume 5: Waste. Hayama, Japan: IGES, 2006.