

Article

Neural Calibration of the Resistance Prediction for Slender Ship Hulls

Davor Mimica , Ines Bezić , Martina Bašić , Branko Blagojević  and Josip Bašić * 

Department of Naval Architecture, Faculty of Electrical Engineering, Mechanical Engineering and Naval Architecture, University of Split, 21000 Split, Croatia; damimica@fesb.hr (D.M.); ines.bezic@fesb.hr (I.B.); martina.basic.00@fesb.hr (M.B.); bblag@fesb.hr (B.B.)

* Correspondence: jobasic@fesb.hr

Abstract

Fast and accurate resistance prediction is critical in early-stage ship design. While Michell's thin-ship theory provides rapid evaluations, its linear assumptions limit accuracy, particularly as hull forms deviate from ideal slenderness. This paper introduces a physics-preserving neural calibration method that improves Michell's theory without replacing the underlying solver. We train a two-dimensional convolutional encoder–decoder, conditioned on Froude numbers via global FiLM modulation, to predict a bounded correction to the geometric effective-slope field. Because the solver remains unchanged, the learned correction acts as an interpretable spatial perturbation rather than a black-box resistance map. Evaluated under a strict leave-one-family-out (LOFO) protocol on a fleet of five slender hull families (DTMB, NPL-4A, Wide-Light Canoe, Wigley, and Delft 372), the neural calibration achieves a mean absolute percentage error (MAPE) of 0.0741. This represents a 24% improvement over a reproduced 2020 baseline and a 7.9% improvement over the uncorrected Michell solver. The 2020 baseline is the rigid boundary-layer and phase-deflection correction of an earlier study by the present group, re-evaluated here on the present hulls at their measured attitudes. Ablation studies show that much of this aggregate gain is captured by a bounded global slope offset, indicating that a spatially uniform displacement correction accounts for most of the improvement on slender hulls, while the spatially varying field mainly adds per-family headroom. Finally, we map the physical boundaries of this approach. Dedicated recovery campaigns on fuller forms (KCS and Series 60) show that the model regresses compared to baselines. This confirms that while the correction successfully refines the linear source distribution for slender hulls, it cannot synthesize missing physics, such as stagnation pressure, separated flow, or wave interference, for fuller or unrelated geometries.

Keywords: Michell theory; wave resistance; neural calibration; convolutional neural network; thin-ship theory; slender hulls; leave-one-family-out validation; FiLM conditioning; effective-slope correction; physics-preserving AI; solver-in-the-loop training



Academic Editor: Hwan-Sik Yoon

Received: 20 May 2026

Revised: 29 June 2026

Accepted: 1 July 2026

Published: 3 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Accurate resistance prediction is needed early in the design process, when the geometry is still changing and many candidate hulls must be screened. The most useful tools at that stage are not the highest-fidelity ones; they are the tools that are fast enough to be used repeatedly, stable enough to support design decisions, and interpretable enough to show why one hull form is better than another. Michell's thin-ship theory [1] remains attractive

for exactly this reason: it reduces wave resistance to a geometry-driven potential-flow integral and evaluates a hull far faster than viscous CFD or model-scale testing.

The same assumptions that make Michell theory useful also limit its accuracy. The theory represents the hull through a linearised source distribution tied to the geometric hull slopes. For genuinely slender forms, this representation captures the dominant wave-making physics. For fuller ships, or for cases where bow stagnation, viscous separation, transom flow, or trim dominate, the linearised slope boundary condition can no longer stand in for the effective wave-making source. Empirical methods such as Holtrop–Mennen address broader design regimes through statistical correction terms [2,3]; high-fidelity CFD resolves more of the flow directly, at a much higher computational cost. Between these extremes is a useful open question: can a fast thin-ship solver be calibrated without discarding its physical structure?

This paper answers that question for slender hull families, and it does so in the same conceptual slot opened by an earlier study from the present group. Bašić et al. [4] extended Michell’s theory by adding an analytic boundary-layer tangency correction to the effective hull slope, showing that this single change improved the predicted resistance of five hull forms (Wigley, Series 60, Delft 372, a yacht hull, and KCS) relative to both the unmodified theory and Holtrop–Mennen. The present paper substitutes a learned bounded correction in that same slot. Relative to [4], the new elements are a data-driven spatial correction pattern in the same effective-slope closure, LOFO rather than per-hull validation, and explicit out-of-distribution failures that the earlier per-hull protocol could not expose.

We do not replace the resistance solver with a generic neural predictor. Instead, a neural model predicts a bounded correction field to the effective hull slope on the same structured centreplane grid that the Michell solver already uses, and the solver is then run unchanged on the corrected slope field. The learned component is therefore restricted to a hydrodynamically interpretable role: it modifies the slope distribution that drives the wave-making integral, and nothing else.

Four quantitative points summarise the empirical content. On a five-family slender fleet (DTMB, NPL-4A, Wide-Light Canoe, Wigley, Delft 372; 128 held-out samples), leave-one-family-out evaluation gives a total resistance MAPE of 0.0741 when a single configuration is selected on training folds and replayed on each held-out family. The reproduced Bašić et al. [4] corrected-Michell setup gives 0.0976 on the same manifest, while an uncorrected Michell replay gives 0.0804. Thus, the neural correction improves the reproduced predecessor row by about 24% and the uncorrected replay by about 7.9%. Over three independently seeded retrainings of that locked configuration, the aggregate MAPE is 0.0679 ± 0.0041 . A per-family oracle that picks the best campaign checkpoint for each held-out family reaches 0.0568, reported as an upper bound rather than as an honest LOFO headline. Holtrop–Mennen on the same manifest gives 0.1026 and is treated as a secondary empirical reference. Under the primary single-config headline, four of the five slender families improve materially, and Wigley is essentially unchanged.

Direct comparison with the predecessor paper is possible only for the families that they share: Wigley, Series 60, Delft 372, the Wide-Light yacht canoe, and KCS. DTMB and NPL-4A are new to the present study; they are evaluated against the same reproduced predecessor solver configuration on the current dataset, rather than against published 2020 figures.

The method is not presented as a universal ship-resistance model. Dedicated recovery campaigns on KCS and Series 60 leave fuller forms outside the recoverable range of this correction architecture (KCS best 0.230 versus reproduced Bašić 2020 baseline 0.077; S60 best 0.113 versus 0.093, never crossing its baseline). One slender-adjacent out-of-distribution case regresses against the reproduced Bašić 2020 baseline (0.278 versus 0.148). These failures

are retained in this paper because they mark a physical boundary: the learned correction can reshape an effective source distribution, but it cannot supply the stagnation pressure, separated flow, or strong interference physics that fuller or unrelated-geometry hulls require.

2. Related Work

Michell's 1898 wave-resistance integral is the canonical thin-ship result and is still attractive as a fast geometry-driven baseline [1]. The known limitations are well documented in textbook treatments [5–7]: linearisation of the free-surface and hull boundary conditions, neglect of viscous and stagnation effects, and rigid coupling between hull slopes and the wave-making source distribution. Tuck and co-workers extended thin-ship theory with higher-order corrections, transom-stern handling, and efficient evaluation of the wave-pattern integrals [8,9]. Later analytical and semi-empirical work reintroduced specific physical mechanisms piecewise: Doctors and others added transom-flow corrections for displacement-class hulls [10], while Insel-Molland and Couser et al. calibrated form-factor and interference corrections for slender high-speed monohulls and catamarans [11,12]. The direct predecessor of the present work, Bašić et al. [4], introduced a boundary-layer tangency correction inside the thin-ship machinery itself: rather than appending a closure term, they modified the effective hull slope that the Michell integral already consumes, demonstrating improvements over the original theory and Holtrop–Mennen on Wigley, Series 60, Delft 372, a yacht hull, and KCS. Each of these methods preserves the underlying potential-flow structure but introduces a hand-designed closure or correction tailored to a hull class.

The wave-resistance literature has continued to engage with these limitations on both the potential-flow and viscous sides. Ozdemir et al. [13] reported that an iterative boundary-element method tracks URANS predictions on the Wigley hull at intermediate Froude numbers but underestimates resistance for $Fn > 0.4$ once boundary-layer separation matters, sharpening the regime in which a corrected linear solver can be expected to behave. Wang et al. [14] and Kim et al. [15] examined the longitudinal-cut method recommended by ITTC 2021 and found systematic differences between the wave-pattern resistance that it estimates and the wave-making resistance proper, reinforcing that the wave-component budget itself is non-trivial to recover even with high-fidelity tools. Terziev et al. [16] showed that a CFD-derived form factor on the KCS varies appreciably at low Froude numbers, which is relevant to the present paper's documented failure to recover KCS-class fuller forms inside a thin-ship correction. Hakim et al. [17] illustrated that slender-body methods remain in active use for U-shaped vs V-shaped comparisons, the kind of slender design-screening regime our calibration is targeted at.

Higher-fidelity panel and free-surface methods—the Dawson-class Rankine-source approach [18] and modern viscous CFD—relax the linearisation, but at substantially higher cost. The practical intermediate-fidelity tier today is Holtrop–Mennen power prediction [2,3], which is a regression on a broad merchant-hull population and does not target slender thin-ship geometries specifically. We include it as a familiar empirical reference point rather than as a direct competitor, and we discuss this distinction in Section 6.

A separate line of work uses machine learning to regress total resistance or power directly from hull geometry and operating condition. Cepowski applied feed-forward networks to added resistance prediction [19]; Margari et al. trained ANNs on the MARAD systematic-series resistance data [20]; Perez-Arribas and others have explored Gaussian-process surrogates for resistance regression [21]; and broader surveys of machine-learning approaches to ship hydrodynamics now exist [22]. These methods generally treat the solver as opaque and the network as a direct geometry-to-resistance map. The resulting predictors

are fast but lose the interpretability of the underlying potential-flow calculation, and their generalisation to unseen hull families depends entirely on the diversity of the training set.

The contribution here sits between these two lines, and it is most directly a sequel to the work of Bašić et al. [4]. We do not replace the Michell solver with a regressor, and we do not hand-design a new closure term. Instead, we restrict the learned object to a bounded perturbation of the input to the existing solver—specifically, the effective hull-slope field that the linearised boundary condition consumes. The resistance is still produced by the unchanged Michell integration. The 2020 paper showed that an analytic boundary-layer correction in this same slot already moves the predicted resistance in the right direction; the present paper asks the next question—whether a learned correction in the same slot can do better on a held-out fleet, and where it stops working. This places the learned correction in the same conceptual slot as Tuck-style closures, Insel–Molland form-factor corrections, or the analytic tangency correction of [4], but the closure form is data-driven on a structured grid rather than analytically prescribed. The cost is that the correction has no analytical form; the benefit is that its generalisation properties can be measured by leave-one-family-out evaluation, and its scope of validity can be probed empirically.

Against this background and the research gap that it identifies (a fast thin-ship solver that can be calibrated from data without discarding its physical structure or its measurable generalisation behaviour), the contributions of this work are as follows:

1. A solver-preserving neural calibration layer for Michell resistance, in which learning is restricted to a bounded effective-slope correction field on the solver's existing centreplane grid.
2. An honest leave-one-family-out evaluation protocol with explicitly separated headline (training-fold-tuned single configuration) and oracle (per-family on held-out) numbers, plus seed and ablation reports.
3. A documented physical-validity boundary: failure analyses on KCS, Series 60, and a slender-adjacent out-of-distribution case, used to scope the method's domain of applicability.

The remainder of this paper is organised around exactly that probe.

3. Materials and Methods

Michell's thin-ship wave resistance integral [1], for a hull with centreplane offset $y = f(x, z)$ moving at speed U in water of density ρ , is

$$R_w = \frac{4\rho g^2}{\pi U^2} \int_1^\infty K(\lambda) [P(\lambda)^2 + Q(\lambda)^2] d\lambda.$$

The Michell kernel in this notation is

$$K(\lambda) = \frac{\lambda^2}{\sqrt{\lambda^2 - 1}}.$$

With $k_0 = g/U^2$, the source amplitudes are

$$P(\lambda) = \iint_{\Omega} \beta(x, z) \exp(k_0 \lambda^2 z) \cos(k_0 \lambda x) dx dz.$$

$$Q(\lambda) = \iint_{\Omega} \beta(x, z) \exp(k_0 \lambda^2 z) \sin(k_0 \lambda x) dx dz.$$

where $\beta(x, z) = \partial f / \partial x$ is the longitudinal hull slope and Ω is the projected centreplane domain. These two relations are the only hydrodynamic content of the wave-resistance solver; total resistance is reconstructed inside the solver pipeline as $R_t = R_w + R_F + R_{vp} + R_{tr}$, with the friction component computed from ITTC-1957 [23,24] and viscous-pressure and

transom corrections supplied by the existing project pipeline. The neural calibration targets only β .

The neural model emits a spatial correction field $\delta\beta_\theta(x, z; F_n)$ on the same structured centreplane grid that the solver already uses, conditioned on Froude number F_n and parameters θ . The corrected slope used by the source-amplitude integrals is

$$\beta_{\text{eff}}(x, z) = \beta_0(x, z) + \delta\beta_\theta(x, z; F_n), \quad \delta\beta_\theta = m \tanh(\hat{f}_\theta), \quad |\delta\beta_\theta| \leq m.$$

where β_0 is the geometric hull slope and m is the campaign-swept range cap, with values $\pi/8$, $\pi/16$, and $\pi/32$. This bounded correction is the only quantity learned; the integration kernel, the friction line, and the post-processing remain identical to the uncalibrated solver. As an immediate consequence, for any sample i the model-induced change in total resistance equals the wave-component change exactly:

$$R_t^{\text{AI}}(i) - R_t^{\text{base}}(i) = R_w^{\text{AI}}(i) - R_w^{\text{base}}(i).$$

Thus, any error reduction reported in this paper is attributable entirely to the wave-resistance term, not to a friction-line offset.

The notation is deliberately aligned with the tangency correction used in the 2020 analytical method [4]. In that work, a boundary-layer displacement thickness changed the effective surface tangent by an angle $\Delta\beta$; here, the network predicts the analogous slope perturbation $\delta\beta_\theta$ directly on the solver grid. Figure 1 shows the original tangency-correction geometry. The present model should be read as a data-driven estimator of this effective-slope correction, not as a replacement for the Michell integral.




Figure 1. Tangency correction geometry from [4]. Black arrows denote the baseline surface normals and grey arrows the rotated (corrected) normals; the rotation angle $\Delta\beta$ changes the effective boundary tangent followed by the outer flow; the present neural model predicts the analogous $\delta\beta_\theta(x, z; F_n)$ field on the Michell centreplane grid.

Each grid sample is represented as five channels: longitudinal position normalised by waterline length, vertical position normalised by waterline length, beam normalised by waterline length, longitudinal beam derivative, and the baseline slope β_0 . The default grid is 81×21 . The Froude number is treated as a global scalar rather than a per-vertex feature. An earlier approach represented this grid as an unordered point cloud; the architecture introduced here instead uses 2D convolutions over (z, x) , exploiting the lattice structure that the solver already imposes.

The model itself is a convolutional encoder–decoder with same-padding 3×3 kernels and channel widths: (32, 64, 128) in the encoder and (64, 32) in the decoder. The decoder emits a single $\delta\beta$ channel at the input spatial resolution. Each block applies a 2D con-

volutional layer followed by FiLM modulation and GELU activation, where the FiLM modulation [25] maps the scalar F_n to per-channel scale and shift coefficients,

$$h' = \text{Conv}(h), h_{\text{cond}} = h' \odot (1 + \gamma(F_n)) + \delta(F_n),$$

with $\gamma, \delta : \mathbb{R} \rightarrow \mathbb{R}^{C_{\text{out}}}$ a small linear map. We chose FiLM rather than a per-vertex F_n channel because F_n is a global operating condition; whether this choice is empirically justified is treated as an ablation question in Section 5 rather than as a methodological assumption. Two scalar phase parameters are produced from the globally pooled encoder latent and feed the same phase model already used by the solver bridge. The output $\delta\beta$ is bounded by the perturbation cap above; no waterplane, smoothness, or symmetry projections are applied at evaluation time. Figure 2 shows the data flow.




Figure 2. Convolutional encoder–decoder with per-block FiLM modulation. The encoder ingests a five-channel centreplane grid, the decoder emits a single bounded $\delta\beta$ channel, and the Froude number is broadcast through FiLM scale-and-shift coefficients at every block.

Training is leave-one-family-out (LOFO). For each held-out family, all samples from that family are removed from training and from any hyperparameter-selection step; the resulting trained model is evaluated once through the Michell solver on the held-out samples to produce that fold’s predictions. Training proceeds in two phases: Phase A trains a Conv2d residual-resistance surrogate on training-fold samples with the same encoder backbone but a scalar regression head; its purpose is to warm-start the encoder, and its loss is mean squared error on residual resistance evaluated against measurements. Phase B transfers the matched encoder into the deflection model, freezes nothing, and optimises the LOFO solver-replay MAPE on the training folds against the unmodified Michell solver.

Because the solver is treated as a black-box function of the deflection field, Phase B uses Simultaneous Perturbation Stochastic Approximation [26]: each forward pass calls the unmodified solver, and we do not implement an adjoint or a differentiable Michell prototype. SPSA estimates the gradient with two solver evaluations per gradient sample regardless of parameter dimension, which is the dominant practical advantage given the

approximately 10^5 trainable parameters. AdamW [27] then applies the update. Differentiable surrogate variants of Phase B were also tested; they consistently optimised the surrogate rather than the true solver MAPE, so the final recipe uses solver-in-the-loop SPSA. A small per-iteration anchor regulariser penalises the L2 norm of the predicted field, $\lambda_{\text{anchor}} \|\delta\beta_\theta\|_2^2$, discouraging collapse onto the field cap m . The campaign sweeps m over $\{\pi/8, \pi/16, \pi/32\}$ and λ_{anchor} over $\{0.01, 0.1, 1.0\}$ with two seeds per cell, for 18 configurations in total. Figure 3 summarises the solver-in-the-loop flow.




Figure 3. Solver-in-the-loop training pipeline. Solid arrows show the data flow; the dashed link marks the transfer of encoder weights from Phase A to Phase B. The neural model emits $\delta\beta$ on the centreplane grid; the Michell solver consumes the corrected slope and returns total resistance; SPSA estimates the gradient of evaluation MAPE with respect to network parameters; AdamW applies the update.

Hyperparameter selection is the most leakage-sensitive step in any LOFO report, so the protocol is made explicit. Two distinct headline numbers are reported in this paper. The single-config LOFO number, which we treat as primary, picks one campaign configuration using only training-fold MAPE aggregated over the five slender folds, and that configuration's checkpoints are replayed on each held-out family. Selection sees no held-out data, and this is the protocol that we recommend readers compare to. The per-family oracle number, reported as an upper bound, instead picks the campaign checkpoint with the lowest MAPE on the held-out family for each fold. This is not an honest LOFO claim; it is the upper bound on what per-family campaign selection could deliver if held-out data were available, and we report it explicitly as such. Section 5.2 additionally reports seed robustness from independent training runs and a three-seed no-FiLM ablation.

The constant-field ablation uses the same LOFO splits and solver bridge but replaces the Conv2d field with one spatially uniform scalar $\delta\beta_{\text{const}}$ per fold. For each held-out family, the scalar is selected on the training folds by an exhaustive 33-point grid search over the locked single-config bound $[-\pi/32, \pi/32]$ with fixed phase parameters $(\alpha, \gamma) = (3.8, 0.0)$,

and then replayed once on the held-out family. This is a deterministic one-parameter comparator rather than a neural architecture.

Table 1 lists the locked values used throughout.

Table 1. Locked hyperparameters for the deflection model and Phase B training.

Hyperparameter	Value
Centreplane grid	81×21
Encoder channels	32, 64, 128
Decoder channels	64, 32
Conv kernel	3×3 , same padding
Activation	GELU [28]
FiLM map	Linear, $1 \rightarrow 2C_{\text{out}}$ per block
$\delta\beta$ bound m	$\pi/8, \pi/16, \pi/32$ (swept)
Anchor weight λ_{anchor}	0.01, 0.1, 1.0 (swept)
Optimiser	AdamW
Learning rate	5×10^{-5}
SPSA samples per step	4
SPSA perturbation radius c	10^{-3}
Phase B epochs	Up to 15 (early stop on training-fold MAPE)
Seeds per cell	2 campaign seeds, plus 3 disjoint robustness seeds

4. Dataset and Experimental Protocol

4.1. Hull Families

The primary evaluation set covers five slender monohull families used for the leave-one-family-out (LOFO) evaluation. Two fuller monohull families are retained as documented boundary cases, and two further hulls are reported as out-of-distribution (OOD) probes. The principal particulars of every monohull form, computed at the tested waterline from the same hull geometry that the solver consumes, are summarised in Table 2.

Table 2. Principal particulars of the monohull forms, evaluated at the tested waterline: n is the held-out sample count; F_n is the maximum tested length Froude number. Coefficients for Delft 372 are not recomputed here because its surface mesh is not part of the geometry export (the form is the demihull of the Delft 372 catamaran, evaluated in isolation); its resistance samples are still included in the LOFO aggregate. Entries marked “n/a” are not computed for this form.

Family	Role	n	L_{wl}/B_{wl}	B_{wl}/T	C_B	C_P	C_{WP}	$\max F_n$
DTMB [29]	LOFO slender	45	7.41	3.10	0.499	0.609	0.775	0.47
NPL-4A [30]	LOFO slender	17	10.39	1.79	0.460	0.683	0.761	1.01
Wide-Light Canoe [31]	LOFO slender	15	5.23	9.84	0.507	0.639	0.718	0.80
Wigley [32]	LOFO slender	38	10.01	1.60	0.443	0.665	0.666	0.48
Delft 372 [33]	LOFO slender	13	n/a	n/a	n/a	n/a	n/a	0.65
Series 60 ($C_B = 0.60$) [34]	Boundary (fuller)	—	7.63	2.45	0.589	0.608	0.699	0.46
KCS [35]	Boundary (full-form container)	—	7.22	2.98	0.625	0.702	0.822	0.29
Begovic W2 [36]	OOD (slender-adjacent)	13	4.39	4.54	0.448	0.895	0.872	1.87
Full-Form Tanker [37]	OOD (full-form)	11	5.41	3.20	0.758	0.762	0.847	0.26

The experimental resistance data for these families derive from established towing-tank campaigns and systematic series: the DTMB model series [29], the NPL high-speed round-bilge series [30], the Wigley mathematical hull [32], the Delft 372 catamaran demihull in isolation [33], the Wide-Light Canoe form [31], Series 60 [34], the KRISO Container Ship [35], and the Begovic warped-chine planing series [36]. The out-of-distribution tanker EFD data were sourced from NMRI model tests [37].

The two out-of-distribution probes can now be defined concretely against this envelope. The full-form tanker has a block coefficient $C_B = 0.76$, far above the slender training envelope ($C_B \leq 0.51$), and represents the viscous-dominated, high-displacement regime that thin-ship theory is not expected to capture. The slender-adjacent probe is the Begovic W2 hull: it is labelled slender-adjacent because its block coefficient ($C_B = 0.45$) falls inside the slender training range, yet it lies outside that distribution on every other axis: a much shorter length-to-beam ratio ($L_{wl}/B_{wl} = 4.4$ against 5.2–10.4 for the slender fleet), a far higher prismatic coefficient ($C_P = 0.90$ against ≤ 0.68), and a high-trim planing regime reaching $F_n \approx 1.9$. This is precisely why a block-coefficient-only view would misleadingly suggest that the method should transfer, while the fuller prismatic form and planing trim place it out of distribution. The KCS ($C_B = 0.63$) is a full-form container hull, distinct from the tanker, and is treated as a fuller-form boundary case rather than an OOD probe.

For clarity, the “slender-adjacent OOD” case referred to in the results and in the out-of-distribution tables is the same geometry as the “Begovic W2” hull shown in the resistance and Froude-sweep figures; the two labels denote a single hull.

Catamarans are excluded from the present paper. The current architecture provides the Froude number as the only global conditioning variable and, therefore, cannot represent demihull-spacing interference physics in a controlled way; catamarans are deferred to a follow-up study with an additional global condition s/L .

The primary metric is mean absolute percentage error on total resistance,

$$\text{MAPE} = \frac{1}{N} \sum_{i=1}^N \frac{|R_t^{\text{pred}}(i) - R_t^{\text{meas}}(i)|}{|R_t^{\text{meas}}(i)|}.$$

Aggregate MAPE values over a fleet are sample-weighted across the contributing families.

4.2. Protocol

For each held-out family, all samples from that family are excluded both from training and from any hyperparameter selection. A checkpoint is then evaluated once through the Michell solver on the held-out samples to produce that fold’s predictions. The headline aggregate is built across folds without family-level weighting (i.e., it equals the sample-weighted aggregate over all 128 slender held-out samples).

Two distinct headline numbers are reported:

1. Single-Config LOFO (primary): One campaign configuration is selected using only training-fold MAPE pooled across slender folds, and that single configuration’s per-fold checkpoints are replayed on the held-out family. Hyperparameter selection sees no held-out data.
2. Per-Family Oracle (upper bound): For each held-out family, the campaign checkpoint with the lowest MAPE on the held-out family is chosen. This uses held-out data and is therefore not an honest LOFO headline; it is reported only as the upper bound on what per-family campaign selection could deliver.

The single-config number is the primary protocol claim; the oracle number is informative because it bounds the headroom that better configuration adaptation could in principle close.

4.3. Baselines

We compare against the following:

- Uncorrected Michell evaluation on the same measured-attitude manifest, with the 2020 boundary-layer and phase corrections disabled.
- A reproduced Bašić 2020 baseline from the existing measured-attitude predecessor pipeline.

- The Holtrop–Mennen baseline [2] generated on the same slender-family manifest and the same Reynolds and Froude conditions, so its sample count matches the headline tables.

The reproduced Bašić 2020 baseline is not an uncorrected Michell evaluation. It uses the Bašić-style corrected Michell configuration preserved from the predecessor work, evaluated on the current manifest under measured attitude for dynamic cases. A parametric NLLS calibration variant explored during development is discussed in Section 6.5 as a development-phase reference point; it is not included as a formal comparator in the headline tables because it was not evaluated under the present LOFO protocol.

We do not compare against a tuned Tuck second-order or Doctors-style transom correction because no production implementation of either is in the current solver pipeline; that comparison is recommended for future work (Section 8.2).

Four secondary checks are used to assess robustness and scope:

- Grid Resolution: The locked trained models are evaluated on 41×11 , 81×21 (default), and 121×31 centreplane-grid exports. The default grid is the one used for training and for the headline tables.
- Bootstrap Confidence Intervals: Per-family and aggregate 95% intervals are computed post hoc from held-out per-sample errors using resampling, with the single-config LOFO predictions treated as the primary interval target.
- Robustness and Ablation: Three independently seeded retrainings of the locked headline cell are reported, plus three independently seeded runs with FiLM conditioning removed (FiLM blocks reduce to plain Conv2d followed by GELU activation), and a one-scalar-per-fold constant-field LOFO ablation.
- Out-of-Distribution Probes: The tanker and the slender-adjacent out-of-distribution case are reported as out-of-scope checks. They do not contribute to any headline aggregate.

The tables and figures in this manuscript are drawn from locked internal result bundles covering the slender LOFO replay, mono-campaign summary, robustness reruns, no-FiLM ablation, KCS and Series 60 recovery sweeps, and the Holtrop–Mennen baseline. The Data Availability Statement states the public artefact commitment for the headline checkpoints, replay scripts, splits, prediction tables, and verification scripts; exploratory internal sweeps may remain outside that archive.

4.4. Scope of the Analysis

Three conditions bound how the results that follow should be read, and we state them here so that the findings are taken with appropriate caution. First, the method addresses steady, calm-water total resistance at the tested attitude; added resistance in waves, manoeuvring, and seakeeping are out of scope. Second, the calibration is trained and validated only on slender displacement monohulls, and the boundary and out-of-distribution cases in Table 2 are reported to show where this scope ends, not to claim general applicability. Third, the learned object is a bounded perturbation of the effective-slope field, which is then consumed by an unchanged Michell solver. Every conclusion therefore inherits the validity limits of linearised thin-ship theory: the correction can reshape a linear source distribution, but it cannot add viscous separation, stagnation pressure, or strongly nonlinear free-surface physics that the solver itself does not represent. We return to the forward-looking implications of these bounds in Section 7.

A final, practical caveat concerns sample size. Several families contribute small held-out sets (NPL-4A, $n = 17$; Delft 372, $n = 13$; and Wide-Light Canoe, $n = 15$) that are too small to support reliable single-family inference on their own. The LOFO protocol is built around this limitation rather than ignoring it: hyperparameters are selected only on the pooled training folds, and each configuration is tested on a fully held-out family

whose geometries were never seen during training, so the aggregate LOFO MAPE over the combined 128 held-out samples reflects cross-family generalisation rather than per-family fitting. To keep the small-sample uncertainty explicit, we report per-family and aggregate 95% bootstrap confidence intervals (Section 5.2); as one would expect, the widest interval belongs to the smallest family, Delft 372. Per-family point values should therefore be read alongside their intervals, while the aggregate remains the quantity that the protocol is designed to support.

5. Results

5.1. Slender-Fleet Headline

We report two distinct numbers, with their selection criteria stated up front. The single-config LOFO number is the primary headline. The campaign cell with the lowest mean training-fold MAPE across both campaign seeds is $(\beta_{\text{cap}}, \lambda_{\text{anchor}}) = (\pi/32, 0.01)$; selection uses only training-fold data. Replaying this configuration's per-fold checkpoints on the held-out family, seed-averaged across the two campaign seeds, gives an aggregate held-out MAPE of 0.0741 over 128 samples. The reproduced Bašić 2020 corrected-Michell setup gives 0.0976 on the same samples, an approximately 24% relative reduction. An uncorrected Michell evaluation with the 2020 boundary-layer and phase corrections disabled gives 0.0804, so the locked neural correction is a 7.9% relative reduction against the uncorrected solver.

The per-family single-configuration LOFO values are listed in Table 3. Throughout the results tables, “AI MAPE” denotes the MAPE of the neural-calibrated Michell prediction, as distinct from the uncorrected-Michell and reproduced Bašić 2020 baseline MAPEs.

Table 3. Single-configuration LOFO held-out MAPE per family, seed-averaged across the two campaign seeds. Improvement values are computed from the unrounded MAPEs.

Family	<i>n</i>	AI MAPE (Single-Config, Seed-Mean)	Reproduced Bašić 2020 Baseline MAPE	Improvement (Exact)
Delft 372	13	0.1271	0.1363	+0.0091
DTMB	45	0.0669	0.0892	+0.0223
NPL-4A	17	0.1092	0.1322	+0.0231
Wide-Light Canoe	15	0.0540	0.1575	+0.1035
Wigley	38	0.0568	0.0553	−0.0014
Aggregate	128	0.0741	0.0976	+0.0235

The uncorrected Michell replay is listed separately in Table 4 because it is a different comparator from the reproduced Bašić 2020 corrected-Michell row.

Table 4. Uncorrected Michell evaluation on the same measured-attitude slender manifest, with the 2020 boundary-layer and phase corrections disabled.

Family	<i>n</i>	Uncorrected Michell MAPE
Delft 372	13	0.1201
DTMB	45	0.0712
NPL-4A	17	0.1272
Wide-Light Canoe	15	0.0911
Wigley	38	0.0526
Aggregate	128	0.0804

Four of the five families improve materially against the reproduced Bašić 2020 baseline. Wide-Light Canoe and NPL-4A see the largest gains, simply because their baseline errors are the largest. Wigley is essentially unchanged: the seed-averaged post-correction MAPE

of 0.0568 is statistically indistinguishable from the baseline 0.0553. We do not claim Wigley as a successful family in the single-config setting.

If the best campaign checkpoint is allowed to be chosen separately for each held-out family, using MAPE on that held-out family as the criterion, the aggregate falls to 0.0568 and all five families improve. This is the per-family oracle, reported only as a leakage-aware upper bound; it is not an honest LOFO headline. The gap between the single-config result (0.0741) and the oracle (0.0568) approximates the headroom available for better configuration adaptation. The constant-field ablation reported below is also included in the main baseline ladder. Figure 4 summarises the locked Conv2d headline against the one-scalar constant-field ablation, uncorrected Michell, reproduced Bašić 2020, and Holtrop–Mennen.




Figure 4. Per-family held-out MAPE for the primary slender-fleet baseline ladder: locked Conv2d single-config neural calibration, one-scalar constant-field ablation, uncorrected Michell replay, reproduced Bašić 2020 baseline, and Holtrop–Mennen empirical baseline.

For convenience, the aggregate baseline, headline, and oracle numbers are collected in Table 5.

Table 5. Aggregate MAPE over the five-family slender fleet ($n = 128$). The single-configuration row is the headline; the per-family oracle row is an upper bound and not an honest LOFO claim.

Selection Rule	AI MAPE	<i>n</i>	Notes
Uncorrected Michell evaluation	0.0804	128	No 2020 boundary-layer or phase correction
Reproduced Bašić 2020 baseline	0.0976	128	Reproduced predecessor-method setup
Holtrop–Mennen baseline	0.1026	128	Empirical; see Holtrop comparison below
Constant-field LOFO ablation	0.0729	128	One spatially uniform $\delta\beta$ per fold
Single-config (train-fold-tuned)	0.0741	128	Primary headline
Per-family oracle (held-out MAPE)	0.0568	128	Upper bound only

The single selected cell is also stable under a leave-one-fold-out check of the selection step. Omitting Delft 372, NPL-4A, Wide-Light Canoe, or Wigley from the selection aggregate still chooses ($\pi/32, 0.01$). Omitting DTMB chooses the neighbouring ($\pi/32, 0.1$) cell; its held-out aggregate differs from the locked cell only in the sixth decimal place of the campaign CSV, so this is not a meaningful selection difference. Across the nine hyperparameter cells (three caps times three anchor weights, each with two campaign seeds), the seed-mean held-out aggregate MAPE ranges from 0.0716 to 0.0743, with a median of 0.0741; the oracle advantage comes from per-family checkpoint selection rather than from a single hidden dominant cell.

5.2. Robustness and Ablation

Three independently seeded retrains of the locked headline cell on the five slender folds give best-per-fold aggregate held-out MAPEs of 0.0718, 0.0696, and 0.0622, for an overall mean of 0.0679 ± 0.0041 (population standard deviation across seeds).

The per-family seed breakdown is given in Table 6.

Table 6. Per-family seed robustness across three independent retrains of the locked headline cell.

Family	<i>n</i>	Mean MAPE (3 Seeds)	s.d.
Delft 372	13	0.1324	0.0078
DTMB	45	0.0624	0.0021
NPL-4A	17	0.0773	0.0119
Wide-Light Canoe	15	0.0472	0.0072
Wigley	38	0.0563	0.0127

Two observations follow: First, the three-seed mean of 0.0679 is better than the two-seed campaign mean of 0.0741 attached to the locked headline cell. This does not change the paper's headline protocol: 0.0741 remains the predeclared single-config result, while 0.0679 ± 0.0041 is a post-lock robustness readout from three fresh reruns of the same cell. Second, Wigley and Delft 372 carry the largest seed variance, consistent with their small held-out sample counts and narrow improvement margins. The aggregate range across seeds (0.0622 to 0.0718) does not overlap the Bašić 2020 baseline of 0.0976, so the gain itself is not seed-fragile, but specific per-family numbers should be read with the seed standard deviations attached.

Replacing each FiLM block with a plain 2D convolution followed by GELU activation (removing per-block global Froude conditioning entirely) gives an aggregate held-out MAPE of 0.0726 ± 0.0007 across three no-FiLM seeds on the same five folds.

The per-family no-FiLM comparison is listed in Table 7.

Table 7. Three-seed no-FiLM ablation against the three-seed FiLM mean.

Family	<i>n</i>	No-FiLM Mean	No-FiLM s.d.	FiLM Mean	FiLM s.d.
Delft 372	13	0.1254	0.0005	0.1324	0.0078
DTMB	45	0.0674	0.0002	0.0624	0.0021
NPL-4A	17	0.1156	0.0033	0.0773	0.0119
Wide-Light Canoe	15	0.0558	0.0011	0.0472	0.0072
Wigley	38	0.0480	0.0003	0.0563	0.0127
Aggregate	128	0.0726	0.0007	0.0679	0.0041

We therefore keep FiLM, but we treat the ablation as a modest design-support result rather than as a central architectural claim. The aggregate mean improves by 0.0047, comparable to the observed aggregate seed standard deviation of the FiLM runs (0.0041), and the no-FiLM variant also beats the baseline. The retained FiLM model is preferred mainly because it does not hurt the aggregate result, improves NPL-4A strongly (0.1156 to 0.0773), and keeps the architecture compatible with future global conditioning variables such as catamaran spacing.

A constant-field ablation fits only one spatially uniform $\delta\beta$ scalar per LOFO fold, selected on the training folds by a 33-point grid search over the locked $\pi/32$ bound and then replayed once on the held-out family. This scalar comparator gives an aggregate held-out MAPE of 0.0729, slightly better than the two-seed locked Conv2d campaign value of 0.0741 and worse than the three-seed Conv2d rerun mean of 0.0679. The result changes the interpretation: most of the aggregate gain over the reproduced Bašić 2020 and uncorrected

Michell rows is achievable with a bounded global slope offset, while the spatial Conv2d field is mainly justified by the per-family oracle headroom, the NPL-4A/Wide-Light Canoe spatial patterns, and future configuration adaptation, rather than by this aggregate scalar gate alone. The per-family constant-field values are listed in Table 8.

Table 8. Constant-field LOFO ablation using one spatially uniform bounded slope offset per fold. The scalar is selected on training folds only and replayed once on the held-out family.

Family	<i>n</i>	Constant-Field MAPE	Best $\delta\beta_{\text{const}}$
Delft 372	13	0.1302	0.0245
DTMB	45	0.0622	0.0245
NPL-4A	17	0.1041	0.0184
Wide-Light Canoe	15	0.0525	0.0245
Wigley	38	0.0600	0.0307
Aggregate	128	0.0729	–

Per-family 95% bootstrap intervals on held-out per-sample errors for the single-config LOFO predictions are reported in Table 9. The aggregate single-config interval is 0.0648 to 0.0838. The corresponding reproduced Bašić 2020 baseline interval is 0.0870 to 0.1089, so the headline gain over that reproduced predecessor row is separated under this resampling check. The interval overlaps the uncorrected Michell aggregate (0.0804), which is why the uncorrected replay is reported as a stronger secondary comparator rather than treated as a floor.

Table 9. Bootstrap 95% confidence intervals on held-out per-sample errors for the single-config LOFO predictions.

Family	<i>n</i>	MAPE	95% Low	95% High
Delft 372	13	0.1271	0.0719	0.1850
DTMB	45	0.0669	0.0548	0.0803
NPL-4A	17	0.1092	0.0874	0.1307
Wide-Light Canoe	15	0.0540	0.0404	0.0672
Wigley	38	0.0568	0.0472	0.0673
Aggregate	128	0.0741	0.0648	0.0838

The DTMB and Wigley intervals are tight; NPL-4A and Wide-Light Canoe are tight relative to their large baseline errors; Delft 372 has the widest interval because it contributes only 13 held-out samples.

Holtrop–Mennen on the same slender manifest gives an aggregate MAPE of 0.1026, worse than both the neural model and the reproduced Bašić 2020 baseline. DTMB is the one family on which Holtrop–Mennen is competitive (0.0457); on the other four, its errors are substantially higher.

The per-family Holtrop–Mennen errors, alongside the reproduced Bašić 2020 baseline, are listed in Table 10.

This comparison should be read carefully: Holtrop–Mennen is an empirical power-prediction regression on a broad merchant-hull population, not a slender-wave calibration. We include it as a non-Michell empirical reference, and because it shows that the neural method is not simply recovering a generic correction.

Because the friction line, transom correction, and viscous-pressure terms in the solver pipeline are unchanged between the reproduced Bašić 2020 baseline and the neural model, the entire change in R_t between the two predictions is a wave-resistance change by construction. This identity is not claimed as a new hydrodynamic result; it is a useful accounting

check showing that the neural model is not borrowing accuracy from a retuned friction line or a modified transom closure.

Table 10. Holtrop–Mennen empirical baseline on the slender manifest, alongside the reproduced Bašić 2020 baseline.

Family	<i>n</i>	Holtrop–Mennen MAPE	Reproduced Bašić 2020 Baseline
Delft 372	13	0.1651	0.1363
DTMB	45	0.0457	0.0892
NPL-4A	17	0.1163	0.1322
Wide-Light Canoe	15	0.1923	0.1575
Wigley	38	0.1070	0.0553
Aggregate	128	0.1026	0.0976

Replays on 41×11 , 81×21 (default), and 121×31 give the following results:
The aggregate replay sensitivity across the three grids is reported in Table 11.

Table 11. Grid sensitivity of the per-family-oracle replay MAPE. Delta values are computed from the unrounded replay MAPEs.

Grid	AI MAPE (Per-Family-Oracle Replay)	Delta vs. Default (Exact)
41×11	0.0593	+4.39%
81×21	0.0568	0.00%
121×31	0.0568	−0.10%

The 121×31 agreement with the default to within 0.1% indicates that the default grid is not under-resolving the result. The coarse 41×11 replay disagrees by 4.4%, which fails the originally planned 1% invariance criterion. We attach the locked artefact to the 81×21 grid and report 121×31 as the high-resolution check; we do not claim that the learned field is a continuous object that can be sampled arbitrarily coarsely. The failed coarse grid check means that the Conv2d model has learned the default grid-sampling pattern as part of the artefact, not only an underlying coordinate-continuous field. The same grid sensitivity holds for the single-config variant up to the same tolerance. Figure 5 plots the three grid replays.




Figure 5. Aggregate held-out MAPE on the per-family-oracle evaluation across three centreplane-grid resolutions. The 81×21 default and 121×31 high-resolution evaluations agree to within 0.1%; the coarse 41×11 evaluation disagrees by 4.4%.

5.3. Boundary Cases and Out-of-Distribution Probes

KCS recovery remained a clear failure across the entire recovery sweep: the best held-out MAPE of 0.230 is more than triple the reproduced Bašić 2020 baseline (0.077). S60's recovery is closer to the baseline but never reaches it: the best configuration delivers 0.113 against 0.093. Across both campaigns, no configuration improved on the reproduced Bašić 2020 baseline.

The OBJ-derived body-plan comparison in Figure 6 helps explain why these failures are not simply unlucky optimiser outcomes. The slender training families present comparatively fine sections, while Series 60, KCS, and the tanker carry much fuller sections and larger adverse-pressure/stagnation regions. A slope-only correction can reshape the effective Michell source distribution on the former group; it cannot add the nonlinear pressure physics missing from the thin-ship model for the latter group.




Figure 6. Transverse body-plan sections extracted from the local OBJ hull geometry for all monohull forms reported in this paper. Grey curves are x-normal frame intersections, orange lines mark the bottom-aligned baseline and centreline, and the blue line marks the waterline.

The best recovered configuration found for each boundary family is listed in Table 12.

Table 12. Boundary-case recovery: best configuration per family across the dedicated KCS and S60 sweeps. Neither family crosses its reproduced Bašić 2020 baseline.

Family	Best Configuration	AI MAPE	Reproduced Bašić 2020 Baseline	Improves?
KCS	ramp 2.0, cap $\pi/16$, anchor 0.1	0.230	0.077	No
Series 60 ($C_B = 0.60$)	ramp 1.0, cap $\pi/8$, anchor 0.1	0.113	0.093	No

Figure 7 plots the boundary-case recovery sweep.




Figure 7. KCS and Series 60 recovery sweeps over field cap, anchor weight, and longitudinal ramp. Each scatter shows held-out MAPE vs the reproduced Bašić 2020 baseline; no configuration crosses the baseline for either family. Boundary-case failure is interpreted physically in the Discussion.

Two further geometries lie outside the declared scope: the slender-adjacent case is not present in any training family; the tanker is a full-form hull far outside scope.

The two out-of-distribution probes are summarised in Table 13.

Table 13. Out-of-distribution probes outside the declared scope of this paper.

Case	<i>n</i>	AI MAPE	Reproduced Bašić 2020 Baseline	Comment
Slender-adjacent OOD (Begovic W2)	13	0.278	0.148	Regresses against baseline
Full-form tanker	11	0.228	0.311	Both errors high; AI lower than baseline

The slender-adjacent case is the more important warning: its geometry resembles that of the training fleet, yet the calibration worsens the prediction. We report it explicitly as a degradation case rather than a transparency observation: out-of-distribution slender geometry can be actively harmed by the learned correction, and any deployment of this method should screen incoming hulls for membership in the training distribution. The tanker improvement is reported for completeness; both predictions remain too inaccurate to be useful, and the case is far outside scope. Figure 8 plots both probes.

As a minimal screening check, we computed nearest-neighbour distance in a standardised hull-coefficient space ($L/B, B/T, C_B, C_P, C_{WP}$) using the five slender training families as the reference set. The five training families have distance 0 to themselves. The slender-adjacent Begovic W2 probe has nearest-neighbour distances of 9.18 standardised units to NPL-4A, 12.06 to the tanker, 6.40 to KCS, and 4.10 to Series 60. This simple coefficient-space screen would therefore flag the slender-adjacent failure before applying the learned correction. It is reported as a diagnostic, not as a validated deployment rule.

Out-of-distribution cases — outside slender-fleet training scope




Figure 8. Per- F_n predicted resistance for the two out-of-distribution probes against measurement and the reproduced Bašić 2020 baseline. The slender-adjacent case regresses against the baseline across most of the curve, while the tanker improves on a worse baseline, but both predictions remain too inaccurate to be useful.

6. Discussion

6.1. Interpretation of the Headline Result

The single-config LOFO result, 0.0741, should be read against two classical rows: 0.0976 for the reproduced Bašić 2020 corrected-Michell setup, and 0.0804 for the uncorrected Michell replay. It is not an interpolation of a resistance curve and not a learned mapping from geometry to resistance. It is a perturbation of the slope field that the unmodified Michell solver consumes. As the wave-component identity in Section 3 makes explicit, every unit of change in total resistance relative to a fixed classical row comes from the wave term; the friction line, transom closure, and viscous-pressure components in the solver pipeline are numerically identical between the baseline and AI rows. The interpretation is therefore narrow but clean: a 2D convolutional model conditioned on Froude numbers can reshape the effective wave-making source distribution of slender hulls in a way that improves replayed wave resistance, while remaining inside a bounded perturbation envelope. This narrowness is also why the method works at all on a 128-sample held-out set—we are not asking the network to learn ship hydrodynamics; we are asking it to predict a small, bounded correction to one input field of a solver that already knows the rest.

The comparison to uncorrected Michell is deliberately conservative. The single-config aggregate improves from 0.0804 to 0.0741, but the bootstrap interval on the neural aggregate (0.0648 to 0.0838) overlaps the uncorrected point, so the paper does not claim a statistically separated aggregate win over the strongest classical replay. The contribution is instead that the learned correction improves four of the five families against the reproduced Bašić 2020 row, documents where it fails, and provides a per-family adaptable field structure whose upper-bound oracle reaches 0.0568. Against uncorrected Michell, the result should be read as a modest held-out gain with transparent uncertainty, rather than as a decisive aggregate separation.

The gap between the single-config LOFO headline (0.0741) and the per-family oracle (0.0568) measures how much different families benefit from different campaign settings. At the configuration level, the oracle picks $\pi/8$ caps with strong anchor for NPL-4A, a tight $\pi/32$ cap with weak anchor for DTMB, and so on. This is a real signal: families with different bow-shoulder shape and beam-to-length ratio benefit from different field amplitudes. A meta-conditioning mechanism that selects the field cap from training-family geometry alone could close part of the 0.017 aggregate gap; we regard that as a prospect for future work, rather than claiming that the oracle number is reachable today.

6.2. Domain of Validity and Boundary Cases

The slender families share a property that makes a bounded slope perturbation plausible: their wave-making is dominated by the linearised source distribution that Michell theory already captures, and the residual error has the character of a small effective-shape correction (boundary-layer displacement near the stern, mild overprediction at sharp bow shoulders, and similar). The learned $\delta\beta$ field has enough freedom to encode where and how strongly the hull contributes to wave-making, while the range cap prevents the optimiser from inventing pathological solutions that satisfy training MAPE but distort the source distribution beyond physical plausibility.

KCS and S60 are not unlucky training cases. The recovery campaigns swept field caps, anchor weights, and longitudinal ramps and found no configuration that improves on the reproduced Bašić 2020 baseline. The failure is consistent with the physics: fuller ships and finer container hulls with sharp bow stagnation introduce nonlinear pressure mechanisms that a linear thin-ship source distribution cannot encode, and adding a bounded perturbation to that distribution does not supply the missing physics. The neural correction is a calibration of the existing thin-ship source model; it cannot add the missing pressure physics. We note that the 2020 analytic-correction paper of the same group [4] did report improvements on KCS and Series 60, but under a per-hull validation protocol rather than a held-out family protocol; the present LOFO setting deliberately denies the model any samples from the held-out family during training, and the difference between those two protocols is part of what the present results measure. The slender-adjacent out-of-distribution case is a sharper and more important warning. Its geometry is slender-adjacent rather than full-bodied, which would naively suggest that the method should help, but it is not represented in any training family, and the AI prediction is markedly worse than the predecessor baseline (0.278 versus 0.148). We read this as the clearest evidence in this paper that the calibration is a within-distribution method. Its safe deployment region is not “any slender hull”; it is “geometries similar to those in the training fleet”. The coefficient-space nearest-neighbour screen reported in the Results section flags Begovic W2 before replay, which converts this case from an unobservable deployment failure into a testable fallback condition: outside the training-family geometry envelope, use the reproduced Bašić 2020 or uncorrected Michell replay rather than the neural correction.

6.3. Resistance Curves and Wave-Making Interference

Figure 9 shows the total resistance versus Froude number for all hulls tested in this paper: the five slender LOFO families, the two boundary-case recovery families, and the two out-of-distribution probes. The slender panels show the single-config LOFO replay, KCS and Series 60 show their best dedicated recovery runs, and the out-of-distribution panels show the slender-trained probe predictions. This view makes the operational envelope visible in the same units that a designer would inspect: the calibrated curve tracks several slender families more closely than the reproduced Bašić 2020 baseline, while the KCS, Series 60, and out-of-distribution panels show why those cases are treated as documented boundary or out-of-distribution failures rather than hidden inside an aggregate metric.

The comparatively large variations in resistance with Froude number visible in Figure 9, particularly the marked rise and the local extrema at higher F_n , are not numerical artefacts but the expected signature of transverse and divergent wave-making interference. As speed increases, the bow and stern wave systems move through successive in- and out-of-phase conditions, producing the classical “humps and hollows” of the wave-resistance curve; these features are intrinsic to thin-ship theory and are well documented in both analytical and experimental wave-resistance studies of slender hulls [9,10]. Because the underlying Michell solver resolves the interference lengths natively through the oscillatory

kernel, the neural calibration does not have to learn this structure from scratch: the learned $\delta\beta$ correction modulates the amplitude of the existing peaks rather than reconstructing the oscillatory curve, which is one reason a bounded slope perturbation suffices to track the measured curves across the tested speed range.




Figure 9. Total resistance versus Froude number for all hulls tested in this paper. Slender-family panels use the single-config LOFO replay; KCS and Series 60 use the best dedicated recovery runs; the out-of-distribution cases use the slender-trained probe predictions.

6.4. Physical Interpretation of the Learned Field

The learned field should also stand up as something physical, not merely a regularised fit. Figure 10 shows the $\delta\beta$ field across six plan-view waterline sections. At each sampled point on the hull curve, a grey tick marks the baseline normal direction (β_0) and a coloured tick marks the AI-corrected normal ($\beta_0 + \delta\beta$), so the angular gap between the two ticks gives the local correction and the colour gives its magnitude on the shared coolwarm scale. Two things stand out: The strongest corrections sit on the after-body, most clearly for NPL-4A and more faintly for Series 60, with the dominant negative patch growing behind $x/L \approx 0.5$. This is what one would expect from boundary-layer displacement-thickness growth, an effect that classical thin-ship theory leaves out and that Insel–Molland-style form-factor corrections only capture piecewise. The analytic tangency correction of [4] reshapes the same effective-slope field, in the same after-body region, for the same physical reason; what the network adds on top is a per-family pattern that the fixed analytic form could not express. The bow, by contrast, stays quiet across all panels, and lower-gain families such as Delft 372 and Wigley show only small angular gaps between their grey and coloured ticks. We make no claim of a quantitative match to any particular displacement-thickness model, but the spatial pattern is clearly not arbitrary.

The constant-field ablation sharpens this interpretation. A LOFO comparator that learns only one spatially uniform $\delta\beta$ per fold reaches 0.0729 aggregate MAPE, slightly better than the locked two-seed Conv2d campaign headline (0.0741) and still below the three-seed Conv2d rerun mean (0.0679). This means that the aggregate improvement cannot be used by itself to prove that spatial Conv2d structure is necessary: a bounded global effective-slope offset captures much of the same correction. The spatial model remains useful for a different claim, namely, that families such as NPL-4A and Wide-Light Canoe admit structured after-body corrections, and that per-family configuration adaptation has substantial headroom. The result therefore turns the ablation from a threat into a clearer mechanism statement: the current aggregate gain is mostly a low-dimensional effective-slope calibration, while the spatial field is the route to the oracle headroom and to physically interpretable local corrections.




Figure 10. Plan-view waterline sections for six monohull families (per-family-oracle checkpoints). Grey ticks show the baseline normals (β_0); coloured ticks show the AI-corrected normals ($\beta_0 + \delta\beta$), with colour encoding the correction magnitude (red: positive, blue: negative). The strongest angular shifts concentrate in the NPL-4A after-body behind $x/L \approx 0.5$; lower-gain families such as Delft 372 and Wigley remain close to their baselines.

6.5. Baselines, Robustness, and Grid Dependence

The Holtrop–Mennen result (0.1026) is best read as a sanity check, not as the central scientific opponent. Holtrop–Mennen is calibrated on a much broader merchant population than the slender families used here, so a slender-only manifest is not its home turf, and its weakness on Wide-Light Canoe and NPL-4A is partially a consequence of that. We include it as a familiar empirical reference point, not because we are claiming that the calibration generalises to Holtrop’s regimes.

A parametric NLLS calibration of the same Michell solver was explored during development as a non-neural reference point. In that scheme, a low-dimensional scalar deflection profile was fitted per hull via nonlinear least squares against measured resistance curves. Over the slender fleet, it achieved an aggregate MAPE of approximately 0.08149. That value is useful context, but it is not a publication gate or a formal comparator: the fit sees each hull’s measurements and was not evaluated under the present LOFO protocol. It

therefore cannot answer the same generalisation question as the single-config neural replay. A LOFO-equivalent low-dimensional deflection baseline remains an area for future work.

The seed-mean aggregate is 0.0679 ± 0.0041 across three independent reruns of the headline configuration. The standard deviation is non-trivial, especially for Delft 372 and Wigley, which have small held-out counts and tight per-sample dynamic ranges. We therefore distinguish two claims that are sometimes conflated. The method is a useful slender-hull calibration: the seed range (0.0622 to 0.0718) does not overlap the reproduced Bašić 2020 baseline 0.0976. The training procedure is not a deterministic recipe that finds the same optimum from any seed: optimiser noise produces meaningful spread, particularly for small-sample folds. Both are true, and this paper reports both rather than leaning on a single locked checkpoint. The matching three-seed no-FiLM ablation is weaker in aggregate (0.0726 ± 0.0007) than the FiLM runs (0.0679 ± 0.0041), so Froude conditioning is retained. The margin is modest, and the no-FiLM variant also passes the baseline gates; the ablation supports the design choice without becoming this paper's main claim.

The 81×21 and 121×31 results agree to within 0.1%, which is good evidence that the default grid is not under-resolving the model. The coarse 41×11 replay disagrees by 4.4% and does not meet the planned 1% criterion; therefore, the result is grid-conditional. This is not merely a plotting detail: in a Michell integral, correction structure near the grid scale can interact with the oscillatory kernel. The present artefact should be read as an 81×21 centreplane-grid calibration with a successful 121×31 high-resolution check, not as a sampling-rate-free continuous field. Coordinate-based field parameterisations, such as implicit neural representations, and direct spectral checks of $\delta\beta$ are attractive next steps if grid-independence becomes a primary deployment requirement.

6.6. Relation to Prior Work

These results sit between two neighbouring strands of the literature, and it is helpful to place them there. Purely data-driven surrogates regress total resistance or power directly from hull geometry and operating condition [19–22], and on favourable in-distribution test sets they can reach percentage errors broadly comparable to ours. The difference is structural, not just numerical. Those models treat the solver as a black box and learn a geometry-to-resistance map, so their predictions are not bound to remain solutions of the underlying flow equations, and their behaviour on unseen hull families rests entirely on how well the training set covers them. Our solver-in-the-loop calibration instead limits the learned object to a bounded perturbation of the effective-slope field that an unchanged Michell solver consumes, so the final resistance always comes out of the linearised potential-flow integral. This is what guarantees physically structured wave fields and keeps the limits of validity explicit and testable, as the documented KCS, Series 60, and slender-adjacent failures show. Set against the analytic predecessor closures (Tuck-style higher-order corrections, Insel–Molland form factors, and the boundary-layer tangency correction of [4]), the present approach keeps the same philosophy of correcting the input field rather than the output number, but it swaps a hand-designed closure for a per-family pattern whose generalisation is measured under LOFO rather than assumed. The contribution, then, is a method that trades the closed analytical form of those corrections for measurable, bounded, physics-preserving data-driven calibration at comparable accuracy on slender hulls.

This paper shows that a bounded neural correction to the effective hull slope, predicted by a 2D convolutional encoder–decoder over a structured centreplane grid and conditioned on Froude numbers, materially improves Michell thin-ship resistance prediction on a five-family slender fleet, under a training-fold-only hyperparameter selection protocol, with documented seed variance and a reported FiLM ablation. It does not claim univer-

sal ship-resistance prediction, KCS-class recovery, or out-of-distribution generalisation. The boundary cases are evidence for the calibration's scope.

7. Conclusions

We have presented a physics-preserving neural calibration of Michell's thin-ship wave-resistance solver, in which a bounded correction to the effective hull-slope field is learned while the solver itself is left unchanged. Evaluated under a strict leave-one-family-out (LOFO) protocol, with hyperparameters selected on training folds only, over a five-family slender fleet of 128 held-out samples, the calibrated solver attains a total-resistance mean absolute percentage error (MAPE) of 0.0741. Two reference points frame this result on the same manifest: the reproduced "Bašić 2020" baseline, the corrected-Michell configuration of Bašić et al. [4], gives an MAPE of 0.0976, and an uncorrected Michell replay gives an MAPE of 0.0804. The calibration is therefore a 24% relative MAPE reduction against the reproduced Bašić 2020 baseline and a 7.9% relative MAPE reduction against the uncorrected solver. Supporting readouts are consistent with this headline: three independently seeded retrainings of the locked configuration give $\text{MAPE } 0.0679 \pm 0.0041$ (population standard deviation), a leakage-aware per-family oracle reaches MAPE 0.0568 (an upper bound on configuration-adaptation headroom), and the Holtrop–Mennen empirical method gives MAPE 0.1026 as a secondary reference. Because the solver and friction line are identical between the baseline and calibrated runs, the entire change in total resistance is carried by the wave-resistance term. A key finding of the ablation study is that most of this gain is captured by a single, spatially uniform slope offset per family, underscoring the dominant role of a global displacement correction.

This method does not extend to fuller-form ships, and we report this boundary explicitly rather than omitting it. Dedicated recovery campaigns on the KRISO Container Ship (KCS) and on Series 60 ($C_B = 0.60$, "S60") leave both families worse than their reproduced Bašić 2020 baselines (KCS best MAPE 0.230 versus baseline MAPE 0.077; S60 best MAPE 0.113 versus baseline MAPE 0.093), and the slender-adjacent out-of-distribution case (the Begovic W2 hull) also regresses (MAPE 0.278 versus baseline MAPE 0.148). A simple coefficient-space nearest-neighbour screen flags that slender-adjacent case as lying outside the training-family geometry envelope, converting a silent deployment risk into a testable pre-screening condition. Together, these cases define the physical boundary of the method: the learned correction can reshape an effective source distribution within the validity regime of linear thin-ship theory, but it cannot supply the stagnation pressure, separated flow, or strong interference mechanisms that fuller or unrelated geometries require.

The practical outcome is a fast, solver-preserving resistance calibration for slender-hull early-stage design, with a clearly defined within-distribution scope, an honest LOFO headline, and documented boundary cases that delimit where the method should and should not be applied.

8. Limitations and Future Work

8.1. Limitations

The scope of the present method is slender-only by construction. Dedicated KCS and Series 60 recovery campaigns confirm that fuller forms lie outside the recoverable range of this correction architecture, and the slender-adjacent out-of-distribution case shows that even slender-shaped hulls can regress when they lie outside the training distribution. Within the slender fleet itself, Wigley is essentially a tie under the headline protocol: its single-config seed-averaged MAPE of 0.0568 is statistically indistinguishable from the reproduced Bašić 2020 baseline of 0.0553, so it is not counted as a successful family at the headline-config level.

Several caveats apply to the headline number itself. The reported approximately 24% aggregate reduction is the conservative single-config reading; the per-family oracle reaches 0.0568 but uses held-out MAPE in its selection and is therefore reported as an upper bound rather than a deliverable. The uncorrected Michell replay is a stronger comparator than the reproduced Bašić 2020 row on this manifest—the neural headline of 0.0741 sits against 0.0804 for uncorrected Michell—so the gain over the uncorrected solver is modest and its bootstrap interval overlaps that comparator. Two ablations further temper the architectural claim. The three-seed no-FiLM model gives 0.0726 ± 0.0007 , weaker than the three-seed FiLM aggregate of 0.0679 ± 0.0041 but still clearing the baseline gates, so FiLM is retained as design support rather than as a primary claim. A one-scalar-per-fold constant-field ablation reaches 0.0729, slightly better than the locked two-seed Conv2d campaign value of 0.0741, so the aggregate gain by itself does not establish that spatial Conv2d structure is necessary.

The remaining limitations are technical. Coarse-grid invariance fails the 1% criterion: the 41×11 replay sits +4.4% off the default grid, so the model is tied to the 81×21 centre-plane grid and is not a continuous, sampling-rate-free field representation. The learned field is bounded but not zero at the waterplane, with symmetry and waterplane projections reported as diagnostics rather than enforced by post-processing. Catamarans are deferred: the current architecture has no global conditioning variable for demihull spacing s/L , so catamaran interference physics is not represented in this work.

8.2. Future Work

The most direct follow-ups address the gaps identified above. A warm-start-matched no-FiLM ablation, in which a matching no-FiLM Phase A surrogate is trained from scratch, would isolate Froude conditioning from encoder warm-start effects more cleanly than the current ablation can. The single-config-to-oracle gap (0.0741 versus 0.0568) measures what a learned cap or anchor selector could recover; a meta-conditioner over $(\beta_{\text{cap}}, \lambda_{\text{anchor}})$ trained on training-family geometry alone is the natural next step. The slender-adjacent regression case suggests an explicit in-distribution test before applying the calibration—for example, a geometry-similarity score against the training fleet—so that out-of-scope hulls fall back to the reproduced Bašić 2020 baseline rather than silently degrading.

Three architectural extensions follow: Adding the demihull spacing-to-length ratio as a global FiLM condition and evaluating interference-physics LOFO on a catamaran fleet is the most immediate. A coordinate-based field representation, such as an implicit neural representation of $\delta\beta(x, z; F_n)$, would address the coarse-grid 4.4% gap directly. SPSA is a black-box workaround for a non-differentiable solver; an adjoint or differentiable Michell prototype would reduce optimisation cost and likely seed variance.

Finally, the empirical scope can be widened. The reproduced Bašić 2020 baseline overlaps only on Wigley, Delft 372, Wide-Light Canoe, KCS, and Series 60; DTMB and NPL-4A are new relative to that earlier paper, so any predecessor-method comparison remains partial by construction. Additional slender families and repeated tank measurements, where available, would tighten per-family confidence intervals, particularly for Delft 372.

Author Contributions: Conceptualization, J.B., I.B., M.B., D.M., and B.B.; methodology, J.B., I.B., and D.M.; software, J.B.; validation, J.B., I.B., M.B., and D.M.; formal analysis, J.B., M.B., and B.B.; investigation, J.B., I.B., and M.B.; data curation, J.B. and M.B.; writing—original draft preparation, J.B.; writing—review and editing, J.B., I.B., M.B., D.M., and B.B.; visualization, J.B.; supervision, D.M. and B.B.; project administration, J.B., D.M., and B.B. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The public artefact will contain the training-fold split definitions, locked single-config checkpoints, replay scripts, per-sample prediction tables for the 128 held-out slender samples, bootstrap script, constant-field ablation artefacts, and the numerical outputs underlying the reported tables and figures. Internal exploratory sweeps and recovery-campaign development logs may be omitted, but the headline replay path and trained weights needed to verify this paper's primary claims are part of the artefact commitment.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Michell, J.H. The wave-resistance of a ship. *Philos. Mag.* **1898**, *45*, 106–123. [\[CrossRef\]](#)
2. Holtrop, J.; Mennen, G.G.J. An approximate power prediction method. *Int. Shipbuild. Prog.* **1982**, *29*, 166–170. [\[CrossRef\]](#)
3. Holtrop, J. A statistical re-analysis of resistance and propulsion data. *Int. Shipbuild. Prog.* **1984**, *31*, 272–276.
4. Bašić, J.; Blagojević, B.; Andrun, M. Improved estimation of ship wave-making resistance. *Ocean Eng.* **2020**, *200*, 107079. [\[CrossRef\]](#)
5. Larsson, L.; Raven, H.C. *Ship Resistance and Flow*; SNAME: Alexandria, VA, USA, 2010.
6. Bertram, V. *Practical Ship Hydrodynamics*, 2nd ed.; Butterworth-Heinemann: Oxford, UK, 2012. [\[CrossRef\]](#)
7. Bertram, V.; Schneekluth, H. *Ship Design for Efficiency and Economy*, 2nd ed.; Butterworth-Heinemann: Oxford, UK, 1998. [\[CrossRef\]](#)
8. Tuck, E.O. Wave resistance of thin ships and catamarans. *Appl. Ocean. Res.* **1987**, *9*, 68–77. [\[CrossRef\]](#)
9. Tuck, E.O.; Lazauskas, L.; Scullen, D.C. Sea wave pattern evaluation: An evaluation of methods for computation of wave resistance using thin-ship theory. *Ship Technol. Res.* **1999**, *46*, 3–14.
10. Doctors, L.J. A numerical study of the resistance of transom-stern monohulls. In Proceedings of the 5th International Conference on High-Performance Marine Vehicles (HIPER), Launceston, Australia, 8–10 November 2006.
11. Insel, M.; Molland, A.F. An investigation into the resistance components of high speed displacement catamarans. *Trans. R. Inst. Nav. Archit.* **1992**, *134*, 1–20.
12. Couser, P.R.; Molland, A.F.; Armstrong, N.A.; Utama, I.K.A.P. Calm water powering predictions for high-speed catamarans. In Proceedings of the FAST'97 (4th International Conference on Fast Sea Transportation), Sydney, Australia, 21–23 July 1997.
13. Ozdemir, Y.H.; Cosgun, T.; Barlas, B. Comparative study on ship wave making resistance via viscous and IBEM solvers. *Ocean Eng.* **2021**, *241*, 110074. [\[CrossRef\]](#)
14. Wang, X.; Zhu, R.; Xu, D.; Wang, H.; Gu, Y. Mechanism analysis and prediction of longitudinal cut wave pattern resistance based on CFD simulation. *J. Ocean Eng. Sci.* **2025**, *10*, 258–273. [\[CrossRef\]](#)
15. Kim, M.; Seo, J.; Kim, M.A.; Park, J.; Seok, W. Examination of wave pattern resistance estimation method by experiments and numerical analyses. *Int. J. Nav. Archit. Ocean Eng.* **2025**, *17*, 100679. [\[CrossRef\]](#)
16. Terziev, M.; Tezdogan, T.; Demirel, Y.K.; Mizzi, S.; Incecik, A. Exploring the effects of speed and scale on a ship's form factor using CFD. *Int. J. Nav. Archit. Ocean Eng.* **2021**, *13*, 147–162. [\[CrossRef\]](#)
17. Hakim, M.L.; Tuswan; Firdaus, A.; Mursid, O. Investigating the comparison of ship resistance components between U and V-shaped hulls. *J. Teknol.* **2023**, *85*, 153–164. [\[CrossRef\]](#)
18. Dawson, C.W. A practical computer method for solving ship-wave problems. In *Proceedings of the 2nd International Conference on Numerical Ship Hydrodynamics, Berkeley*; University Extension Publication: Berkeley, CA, USA, 1977; pp. 30–38.
19. Cepowski, T. The prediction of ship added resistance at the preliminary design stage by the use of an artificial neural network. *Ocean Eng.* **2020**, *195*, 106657. [\[CrossRef\]](#)
20. Margari, V.; Kanellopoulou, A.; Zaraphonitis, G. On the use of artificial neural networks for the calm water resistance prediction of MARAD systematic series' hullforms. *Ocean Eng.* **2018**, *165*, 528–537. [\[CrossRef\]](#)
21. Pérez-Arribas, F.; Calderón-Sánchez, J. A Gaussian-process-based machine learning approach for hull resistance prediction. In Proceedings of the 16th International Conference on Computer Applications and Information Technology in the Maritime Industries (COMPIT), Cardiff, UK, 15–17 May 2017.
22. Panda, J.P. Machine learning for naval architecture, ocean and marine engineering. *J. Mar. Sci. Technol.* **2023**, *28*, 1–26. [\[CrossRef\]](#)
23. International Towing Tank Conference. Report of the Resistance Committee. In Proceedings of the 8th International Towing Tank Conference, Madrid, Spain, 15–23 September 1957.

24. International Towing Tank Conference. ITTC—Recommended Procedures and Guidelines: Resistance Test, 7.5-02-02-01; International Towing Tank Conference, 2014. Available online: <https://www.ittc.info/media/8001/75-02-02-01.pdf> (accessed on 19 May 2026).
25. Perez, E.; Strub, F.; de Vries, H.; Dumoulin, V.; Courville, A. FiLM: Visual reasoning with a general conditioning layer. *Proc. AAAI Conf. Artif. Intell.* **2018**, *32*, 3942–3951. [[CrossRef](#)]
26. Spall, J.C. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Trans. Autom. Control* **1992**, *37*, 332–341. [[CrossRef](#)]
27. Loshchilov, I.; Hutter, F. Decoupled weight decay regularization. In Proceedings of the International Conference on Learning Representations (ICLR), New Orleans, LA, USA, 6–9 May 2019.
28. Hendrycks, D.; Gimpel, K. Gaussian Error Linear Units (GELUs). *arXiv* **2016**, arXiv:1606.08415. [[CrossRef](#)]
29. Yeh, H.Y.H. Series 64 Resistance Experiments on High-Speed Displacement Forms. *Mar. Technol. SNAME News* **1965**, *2*, 248–272. [[CrossRef](#)]
30. Bailey, D. *The NPL High Speed Round Bilge Displacement Hull Series*; Number 4 in Marine Technology Monograph; Royal Institution of Naval Architects: London, UK, 1976.
31. Claughton, A.R. *SYRF Wide-Light Project*; Technical Report; Sailing Yacht Research Foundation: Jamestown, RI, USA, 2015.
32. Kajitani, H.; Miyata, H.; Ikehata, M.; Tanaka, H.; Adachi, H.; Namimatsu, M.; Ogiwara, S. *The Summary of the Cooperative Experiment on Wigley Parabolic Model in Japan*; Technical Report; Report presented to the Resistance Committee of the 17th ITTC; The University of Tokyo: Tokyo, Japan, 1983.
33. van't Veer, R. *Experimental Results of Motions, Hydrodynamic Coefficients and Wave Loads on the 372 Catamaran Model*; Technical Report 1129; Delft University of Technology: Delft, The Netherlands, 1998.
34. Toda, Y.; Stern, F.; Longo, J. Mean-Flow Measurements in the Boundary Layer and Wake and Wave Field of a Series 60, $C_B = 0.6$ Ship Model—Part 1: Froude Numbers 0.16 and 0.316. *J. Ship Res.* **1992**, *36*, 360–377. [[CrossRef](#)]
35. Kim, W.J.; Van, S.H.; Kim, D.H. Measurement of flows around modern commercial ship models. *Exp. Fluids* **2001**, *31*, 567–578. [[CrossRef](#)]
36. Begovic, E.; Bertorello, C. Resistance assessment of warped hullform. *Ocean Eng.* **2012**, *56*, 28–42. [[CrossRef](#)]
37. Kume, K.; Hasegawa, J.; Iwasaki, Y.; Namimatsu, M.; Uto, S.; Ohashi, K. *Tank Tests of KVLCC2 Model Ship for SIMMAN 2008 Technical Report*; NMRI Report No. 2005-04; National Maritime Research Institute: Tokyo, Japan, 2005.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.