

Supplementary Materials

Surrogate Modelling of Potential Progressive Coating Loss Effects in Novel Aero-Engine CMC Seal Segments

Alessio Langiu, Giacomo Canale and Felice Rubino

This document collects five parts of the study that the manuscript refers to but does not need to carry in full. These are the implementation detail of the four surrogates compared, the behaviour of the recommended twin across the whole loss range, the worked prediction and gallery for the two-layer problem, the mesh study behind the crack-tip reading, and the bound on the step a lost cell leaves at its own edge. Section, figure and table numbers are prefixed with S. References to numbered sections without the prefix are to the manuscript.

S1 The Surrogates in Detail

Every surrogate takes the same input and predicts the same output, so that the comparison isolates the model and not the representation. The input is the coverage vector $c \in \{0, 1\}^{30}$ itself, flattened row-major from the 6×5 grid, with no normalisation and no feature engineering. The output is the temperature at all $D = 284\,742$ nodes the solver reports. A surrogate is worth building only if it is cheaper than the solve it replaces, so we report the cost of each model as well as its error.

The *linear superposition baseline* needs no training. We solve the intact segment once and the thirty single-cell losses once each, and set $\Delta_i = T_i - T_{\text{intact}}$. The prediction for a pattern is $T_{\text{intact}} + \sum_{i \text{ lost}} \Delta_i$, one dot product with the loss indicator. It requires 31 finite-element solves in total, all of which are needed anyway, and evaluating it is free. Every residual variant below, marked (a2), (c) and (d2), learns the discrepancy $T - T_{\text{intact}} - \sum_{i \text{ lost}} \Delta_i$ of this baseline and adds the baseline back at prediction time.

Methods (a) and (a2) are a single exact Gaussian process on the whole field. The kernel is squared-exponential with one shared length scale over the thirty input dimensions, a signal amplitude and a white-noise term, so it carries three hyperparameters. We fit these by maximising the log-marginal likelihood with L-BFGS-B from three starting points. Training targets are centred on their mean across the training set. No other scaling is applied. Sharing one kernel across all D outputs is what makes the model tractable. This is the separable, or intrinsic-coregionalization, multi-output Gaussian process [1, 2], fitted in the matrix-normal case with the output covariance taken as the identity. The per-node amplitudes the field actually carries enter only the predictive variance (Section 4.4 of the manuscript), not the mean, which is the same whatever the inter-output covariance. That one factorisation serves every output is a standard consequence of that structure [3, 4]. What is specific here is only its use at the scale of a full 284 742-node field. The log-marginal likelihood summed over the outputs is $-\frac{1}{2} \text{tr}(K^{-1}YY^\top) - \frac{D}{2} \log |K|$ up to an additive constant, and $YY^\top$ is an $n \times n$ Gram matrix which does not depend on the hyperparameters. We precompute it once, at cost $O(n^2D)$, after which each optimiser step is one Cholesky factorisation of an $n \times n$ matrix, $O(n^3)$, instead of a solve with 284 742 right-hand sides. The posterior mean is $K_*(K + \sigma^2 I)^{-1}Y$, so one factorisation serves every node at once. This is what usually rules an exact full-field Gaussian process out, and it is not expensive at all.

Methods (b) and (c) are the standard reduced-basis construction. We take the twenty leading principal components (PCA) of the training fields, which for (c) are the components of the residual field rather than of the temperature. We fit one independent Gaussian process (GP) per mode, each with a constant times a squared-exponential kernel plus white noise, optimised with two restarts and with the mode coefficients standardised internally. The prediction reassembles the field from the twenty predicted coefficients. Method (d) replaces the twenty Gaussian processes with a single multilayer perceptron from the thirty input bits to the twenty coefficients, with two hidden layers of 128 units, ReLU activations, and Adam with early stopping. Method (d2) is the same network on the residual. Since (d) keeps the basis of (b) and changes the regressor alone, and not the representation of the field, the pair isolates the regressor from the basis.

S2 The Twin Across the Loss Range

Figure S1 repeats the worked prediction of Section 4.3.1 of the manuscript, on the three loss regimes of its caption. At every regime the prediction reproduces the finite-element peak to a tenth of a degree, and the

residual is small and distributed across the face, with no preferential concentration at the cell boundaries. Its worst value over the whole face is 0.56, 0.54 and 0.92 °C. The $k = 5$ row is the pattern of Figure 7 of the manuscript, and the two agree exactly: both are the same refit of the same model on the same split, and the bound they report is the 0.54 °C quoted there.




Figure S1: The full-field twin (method a) across the three loss regimes, on held-out patterns: low loss ($k = 1$), maintenance ($k = 5$) and severe ($k = 13$). Each row shows the loss pattern (black cells lost), the finite-element interface temperature, the Gaussian-process prediction, and their difference; within a row the finite-element and twin panels share a colour scale. The worst error over the whole face grows only from 0.56 to 0.92 °C from one lost cell to thirteen.

S3 Two Coating Layers: a Worked Prediction and a Gallery

Figure S2 works one two-layer prediction through, and Figure S3 repeats the exercise on four further held-out patterns, from a lightly worn segment to one that has lost almost all of its coating. The peak-temperature error stays within 1.7 °C across that whole range, and at near-total loss the twin still returns the peak to 0.1 °C. At its worst point on the interface, however, the two-layer twin departs from the solve by 4.5 to 6.1 °C on those patterns, an order of magnitude more than the single-layer case. This is the price of the second layer and not of the loss regime.




Figure S2: A worked two-layer prediction on a pattern the twin never saw, coloured by coating state (grey intact, orange abrasable gone, red both gone). The finite-element field and the twin’s prediction are close enough that their difference is the informative panel: the twin reproduces the solve to within a few degrees at its worst point and to hundredths of a degree on average, at a batched cost of milliseconds per pattern against a solve that takes minutes.




Figure S3: The two-layer twin across the loss range, on patterns it never saw. From top to bottom: (a) 3 abrasable cells lost, 2 of them through the EBC; (b) 11 and 6; (c) 15 and 9; (d) 28 and 15, near-total loss. Colour marks the coating state per cell (grey intact, orange abrasable gone, red both gone); alongside it each row carries the finite-element field, the twin’s prediction and their difference. The peak-temperature error stays within 1.7°C across the whole range, and at near-total loss (d) the twin still returns the peak to 0.1°C .

S4 Mesh Objectivity of the Crack-Tip Reading

Section 4.8 of the manuscript reports that the stress at the crack tip has no mesh-converged value. This section carries the study behind that statement, and the reason the manuscript nevertheless draws a conclusion from the same model.

Two refinement directions were run separately, on the same seeded cracks, the same load case (the published 500/490/400 fluxes) and the same crack paths, so that each isolates one element dimension. The crack is the same in millimetres at every level, though its extent is quantised on each grid and so varies between 6.75 and 7.5 mm. *In-plane* refinement shrinks Δx in the plane the crack runs in; *through-thickness* refinement raises the number of element layers n_t across the 6 mm substrate. Alongside the peak, each run reports the mean stress within a fixed physical radius of 1.5 mm of the tip, with the crack cells excluded — a reading that belongs to the material rather than to the discretisation.

Table S1: The two refinement directions, read at the tip of the second seeded crack. The peak is the worst principal stress at the element centroid; $\bar{\sigma}_{1.5}$ is the mean over a fixed 1.5 mm radius, and n the number of cells that average covers.

Direction	Level	Elements	Peak [MPa]	$\bar{\sigma}_{1.5}$ [MPa]	n
in-plane ($n_t = 6$)	$\Delta x = 1.5$ mm	11 760	191.4	140.0	3
	$\Delta x = 1.0$ mm	26 280	245.7	130.5	7
	$\Delta x = 0.75$ mm	46 560	296.2	143.9	10
	$\Delta x = 0.5$ mm	105 120	360.1	138.3	25
	$\Delta x = 0.375$ mm	187 200	418.8	139.1	44
through-thickness ($\Delta x = 0.75$ mm)	$n_t = 3$	23 280	286.9	138.2	10
	$n_t = 6$	46 560	296.2	143.9	10
	$n_t = 12$	93 120	297.8	146.4	10
	$n_t = 24$	186 240	300.3	147.8	10

Three readings of the table matter.

First, **the peak diverges under in-plane refinement and converges under through-thickness refinement**. Sixteen times the elements in plane raise it by 119%, from 191.4 to 418.8 MPa, with no sign of a limit; eight times the layers through the thickness raise it by 4.7% in all, by increments of 9.3, 1.6 and 2.5 MPa which are an order of magnitude smaller and do not fall in a regular way, so we report the range and do not extrapolate it. The element dimension that governs the divergence is therefore the one measured in the plane the crack runs in, that is the distance from the tip. This is what an elastic crack-tip field predicts, and it is why refining the mesh cannot settle the tip stress: the local exponents of the in-plane series, computed between the five levels of the table as p in $\sigma \propto \Delta x^{-p}$, are 0.62, 0.65, 0.48 and 0.52: they sit around the $r^{-1/2}$ of that field, though four exponents scattered by 30% of their mean are too few to call the value.

Second, **in the direction the peak diverges the fixed-radius reading stays bounded, and without a trend**. Over the same sixteenfold in-plane refinement $\bar{\sigma}_{1.5}$ takes the values 140.0, 130.5, 143.9, 138.3 and 139.1 MPa: it is not monotone, and it stays inside a band of 9.7% of its mean while the peak grows by 119% without turning over. The contrast is therefore between a quantity that is monotone and unbounded and one that is bounded and directionless, which is the distinction that matters here; it is not a claim that the fixed-radius reading has converged. The number of cells it averages grows from three to forty-four over the same series, so neither its level nor its scatter is an artefact of averaging few cells. This is a measured property rather than a fitted exponent, and it is why the reading, and not the peak, is the one that can be quoted.

Third, **what residual movement the fixed-radius reading has comes from the other direction**. Through the thickness it drifts from 138.2 to 147.8 MPa, 7.0%, by increments of 5.7, 2.5 and 1.5 that fall regularly and extrapolate to about 150 MPa. Neither quantity runs away through the thickness; only the peak does, and only in plane.

The elastic card of the slab

The slab carries the orthotropic card of the deck, the same one the full-mesh model of Section 4.9 carries, so the two structural models of the paper share a material. The slab is nevertheless insensitive to that anisotropy: it reads an in-plane stress, and in plane the material is uniform to 1.2% ($E_1 = E_2$, $\alpha_1 = \alpha_2$, and $G_{12} = 151\,800$ against $E/2(1 + \nu) = 150\,000$ MPa). The anisotropy is through-thickness, the direction the slab does not read. It is also read where it is weakest: E_3/E_1 falls from 0.730 at 20 °C to 0.603 at 1300 °C, and the card is taken at 20 °C.

The reading of the manuscript’s Section 4.8, the stress in the cell immediately ahead of the tip rather than the global peak, is tabulated separately below across both refinement directions. The two quantities coincide

once the mesh is fine enough — at $\Delta x = 1.0$ mm and below the global peak *is* the tip cell — and differ by three per cent on the coarsest slab, where the peak sits one cell away from the tip.

Table S2: Worst principal stress ahead of the second crack tip, load case 500/490/400, as the slab discretisation is refined. Through-thickness refinement converges; in-plane refinement does not, and cannot, the tip field being singular.

Δx (mm)	n_t	elements	tip stress (MPa)
1.5	3	5880	177.9
1.5	6	11760	185.3
1.5	12	23520	187.6
1.5	24	47040	188.1
1.0	6	26280	245.7
0.75	6	46560	296.2

The arrest length does not converge either

The same in-plane series carries the length at which the second seeded crack arrests, and it moves more than the manuscript could report from the two levels available when it was written. At $n_t = 6$ the arrest length is 9.0, 11.0, 12.8, 14.0 and 15.4 mm at $\Delta x = 1.5, 1.0, 0.75, 0.5$ and 0.375 mm: monotone in the refinement, with increments of 2.0, 1.8, 1.2 and 1.4 mm which do not fall in a regular way — the last is the larger of the two, by 0.2 mm, which is less than one cell of the finest grid. A series of that shape supports no extrapolation to a limit, and none is attempted here. The first seeded crack does not move at all: 9.0 mm at every level, which is the 8 mm seed of the manuscript quantised on the grid as above.

Two things follow, and they should be separated. The *qualitative* result is unchanged and is what the manuscript rests on: neither crack grows without limit at any discretisation run — the first is never driven at all, the second advances and arrests. The *quantitative* one does not survive: the arrest length has not converged, and the figure of 20 mm quoted in the manuscript as the length at which growth becomes unstable is not a margin against it: the length grows at every refinement and settles at none, so nothing in the series holds it below that value. The arrest length should therefore be read as a number that grows with the discretisation, of the same nature as the tip stress itself, and no clearance from 20 mm should be inferred from it.

Two limits of the study should be read with it. The threshold of 180 MPa is calibrated for element death at a given element size, not for a fixed-radius reading, so the two numbers are not a margin when placed side by side and none is claimed here. And the tip position is located on the element grid, so it shifts by less than one element between levels; the effect is below the mesh spacing and does not account for anything above.

S5 The Step-Induced Augmentation at Cell Scale

This model lets the gas path enter as an imposed flux. A lost cell, however, is not simply a patch of altered resistance: it is a rearward-facing step of the thickness of the coating removed, and a separated shear layer raises the local heat-transfer coefficient where it reattaches [5]. This section bounds how much of that survives at the scale the model resolves. It settles the step term only. The driving-temperature term is a different question and is taken up at the end.

The step term is dimensionless, and can therefore be estimated without the gas-side operating point, which is proprietary to the engine manufacturer and which we do not hold.

The reattachment is internal to the patch. The removed stack is 0.50 mm for the abradable alone and 0.75 mm with the barrier beneath it, against a cell of 12.2 mm circumferentially and 9.894 mm in the trailing-edge column: L/h lies between 13 and 24. Reattachment occurs a few step heights downstream, the distance depending on the Reynolds number and on the state of the boundary layer at separation [6]. Across the range $x_r = 4h$ to $8h$ it falls between 2.0 and 6.0 mm, in every case shorter than the cell. The condition for the shear layer to reattach inside the lost patch rather than to bridge it is just $L/h > x_r/h$, and this component satisfies it with room to spare. The peak in the heat-transfer coefficient therefore lies *inside* the lost cell, not downstream of it, and the cell mean passes over it.

What the substrate sees is the mean. That is not a modelling convenience but the consequence of Section 3.1.1 of the manuscript: the lateral damping length is $\lambda \approx 11$ mm, about one cell, so conduction integrates the surface distribution over precisely the scale on which it varies. The same smoothing removes the only spatial signature the two depths would otherwise leave, the shift of the reattachment point from 3.0 to 4.5 mm at $x_r = 6h$.

The mean is unity to about one point. Take three zones along the cell: a recirculation region from the step to reattachment, exchanging less than the undisturbed surface; a band of width $2h$ about the reattachment point, exchanging more; and a redeveloping boundary layer over the remainder, taken as undisturbed. With the peak at twice the undisturbed value, the recirculation at 0.7 of it and $x_r = 6h$, the length-weighted mean over one cell is that of Table S3.

Table S3: Cell-mean heat-transfer augmentation over a lost cell, relative to the undisturbed surface, for the two loss depths and the two cell lengths. Peak $\times 2$ over a band of $2h$ at $x_r = 6h$, recirculation at 0.7, remainder undisturbed.

cell L (mm)	lost stack	h (mm)	L/h	x_r (mm)	cell mean
12.2	ABR	0.50	24.4	3.00	1.008
12.2	ABR+EBC	0.75	16.3	4.50	1.012
9.894	ABR	0.50	19.8	3.00	1.010
9.894	ABR+EBC	0.75	13.2	4.50	1.015

The gain at reattachment is cancelled by the deficit in the recirculation, and the two depths give the same mean to two decimals. That is not a coincidence of the chosen values. Every zone length scales with the step height, so the mean has the closed form

$$\bar{a} = 1 + \frac{h}{L} \left[2(a_p - 1) + \frac{x_r}{h}(a_r - 1) \right], \quad (\text{S1})$$

Here the depth enters only through h/L , and the bracket carries the cancellation. It vanishes when $a_p = 1 + (x_r/2h)(1 - a_r)$, which at $x_r = 6h$ and $a_r = 0.7$ is $a_p = 1.9$, within a tenth of the value the measurements support. The bracket is 0.2 here, so the whole effect is $0.2h/L$, at most one and a half parts in a hundred. Equation (1) also shows what would break the cancellation, and it is not depth: it is a peak much sharper than twice the undisturbed exchange, or a recirculation that does not run cold.

How far the conclusion survives its inputs. We sweep the three borrowed quantities over their classical ranges: x_r/h from 4 to 8, the peak from twice to three times the undisturbed value, and the recirculation from 0.7 to 1. Over that sweep the cell mean stays between 0.97 and 1.30. The upper corner requires every unfavourable choice at once, on the shortest cell and the deepest step: no deficit at all in the recirculation, a threefold peak, and the longest reattachment. Removing the single most doubtful assumption, the recirculation deficit, and keeping everything else central leaves the mean at 1.12. The estimate is a bound, not a measurement, and it is reported as one. The script that produces both tables is `scripts/step_augmentation.py`.

Scope of the Result. The objection that the loss geometry itself invalidates a prescribed flux is settled: at cell scale the step-induced augmentation is a one-per-cent effect centrally and a tens-of-per-cent effect at the unfavourable extreme, not a factor. The other half of the assumption is not, and nothing here should be read as if it were. Under a convective condition the exposed substrate face runs some 124 to 248 °C colder than the outer surface of the intact stack it replaces (Section 4.8.1 of the manuscript), which raises the driving temperature difference and with it the heat entering the component. That term needs the gas-side operating point to be evaluated. It is not dimensionless, and the caveat stated above stands because of it.

References

- [1] Edwin V. Bonilla, Kian Ming A. Chai, and Christopher K. I. Williams. Multi-task Gaussian process prediction. In *Advances in Neural Information Processing Systems 20 (NIPS 2007)*, pages 153–160, 2007.
- [2] Mauricio A. Álvarez, Lorenzo Rosasco, and Neil D. Lawrence. Kernels for vector-valued functions: A review. *Foundations and Trends in Machine Learning*, 4(3):195–266, 2012. doi: 10.1561/22000000036.

- [3] Jonathan Rougier. Efficient emulators for multivariate deterministic functions. *Journal of Computational and Graphical Statistics*, 17(4):827–843, 2008. doi: 10.1198/106186008X384032.
- [4] Stefano Conti and Anthony O’Hagan. Bayesian emulation of complex multi-output and dynamic computer models. *Journal of Statistical Planning and Inference*, 140(3):640–651, 2010. doi: 10.1016/j.jspi.2009.08.006.
- [5] J. C. Vogel and J. K. Eaton. Combined heat transfer and fluid dynamic measurements downstream of a backward-facing step. *ASME Journal of Heat Transfer*, 107(4):922–929, 1985. doi: 10.1115/1.3247522.
- [6] B. F. Armaly, F. Durst, J. C. F. Pereira, and B. Schönung. Experimental and theoretical investigation of backward-facing step flow. *Journal of Fluid Mechanics*, 127:473–496, 1983. doi: 10.1017/S0022112083002839.