



Multi-fidelity Kriging-assisted structural optimization of whole engine models employing medial meshes

Hau Kit Yong¹ · Leran Wang¹ · David J. J. Toal¹ · Andy J. Keane¹ · Felix Stanley²

Received: 12 November 2018 / Revised: 18 February 2019 / Accepted: 19 February 2019
© The Author(s) 2019

Abstract

Finite element models of whole gas turbine engines, also known as whole engine models (WEMs), which consist of three-dimensional solid elements are not commonly used in design optimization studies due to the high computational cost of solving them for many designs. WEMs consisting of two-dimensional shell elements can be a suitable replacement for high-fidelity solid WEMs as they approximate the responses well while being significantly quicker to solve. However, in a surrogate-assisted optimization study, the accumulation of errors in the shell WEM evaluations can result in the construction of a surrogate model that can be somewhat misleading compared to the solid WEM response surface. Such a surrogate model could return promising designs that, when validated using solid WEMs, turn out to be suboptimal or infeasible. A novel approach which combines medial meshing and multi-fidelity surrogate modelling techniques is proposed to increase the feasibility of conducting whole engine optimization studies. We demonstrate the workflow for generating medial meshes on an engine inter casing geometry. The accuracy of medial mesh simulations with respect to solid mesh simulations is evaluated and discussed in the context of their suitability as a source of low-fidelity structural information for multi-fidelity surrogate models. The impact of this combination of techniques is subsequently illustrated using two case studies. The first case study is the optimization of an intermediate compressor casing for minimum mass with constraints on the casing stiffness. The results show that the multi-fidelity approach is able to find optimum designs that are equivalent to the expensive single-fidelity approach of using only solid mesh evaluations but at a significantly lower computational cost. The second case study is the optimization of a whole engine geometry. This case study serves to demonstrate the effectiveness of the multi-fidelity approach for solving realistic design problems.

Keywords Surrogate modeling · Multi-fidelity · Kriging · Whole engine models · Medial mesh

Responsible Editor: Erdem Acar

✉ Hau Kit Yong
hky1g12@soton.ac.uk

Leran Wang
leran.wang@soton.ac.uk

David J. J. Toal
d.j.j.toal@soton.ac.uk

Andy J. Keane
andy.keane@soton.ac.uk

Felix Stanley
felix.stanley@rolls-royce.com

¹ Faculty of Engineering and Physical Sciences, University of Southampton, Southampton, UK

² Rolls-Royce plc., Derby, UK

1 Introduction

Due to advances in computational power and simulation efficiency, whole engine models (WEMs) have become a staple tool in the design process of gas turbine engines in industry (Voutchkov et al. 2006; Arkhipov et al. 2009; Toal et al. 2014). Analyzing WEMs allows engineers to capture both the physics of inter-component interactions and also emergent behaviour that would otherwise be lost if the analysis was done on each component in isolation. For example, WEMs facilitate studies on the vibrational response of the engine, the magnitude of which needs to be minimized to increase fatigue life and reduce cabin noise (Garcia 2008). WEMs are also useful for understanding the engine's structural behaviour during extreme scenarios, such as in the event of a Fan-Blade-Off (Bettebghor et al. 2013). The assessment of whether the engine casing is able to contain the detached fan blade requires a WEM that is

accurate yet cheap to run, due to the high computational cost of the involved non-linear transient analysis.

However, the high cost of solving a WEM that is meshed using solid elements restricts the number of simulations that can be performed in the overall design process. This means that these simulations are often only used as gatekeepers for design decisions after they have been proposed, instead of as key drivers behind these decisions. It is thus necessary to simplify the WEM such that they can take on a more active role in design. The simplification needs to reduce the cost of solving the WEM without significantly affecting its structural behaviour when compared to the original model. The simplification method used in this work is medial object transformation, which is an approach for dimensionally reducing geometry.

By reducing solid CAD geometries to their medial objects, the resulting surfaces can be meshed using shell elements. We define shell meshes of this type as a medial mesh. In general, medial meshes can be much quicker to solve than solid meshes of the same geometries as larger shell elements can be used while maintaining solution accuracy. When medial meshes are used in design studies that involve only symmetric thickness changes, the time savings are even more notable as remeshing is not required. The thickness modifications can be realized by simply overwriting the physical properties of each shell element. Voutchkov et al. (2006) demonstrated this geometry modification approach in a multi-objective optimization study on a whole jet engine which employed manually constructed shell meshes. The simulation time was reduced but at a cost of increased manual effort and time during mesh construction. The MANTLE software (Modelling and Analysis in the NeuTral LinE) has been developed in Rolls-Royce to automate the medial object transformation process (Stanley 2010), and it is used to produce the medial meshes in the current work.

Evaluations of whole engine designs using medial meshes can be used to construct a relatively inexpensive surrogate model of the overall structural behaviour. The predictive accuracy of this surrogate model, with respect to whole engine evaluations using expensive solid meshes, can be further improved by augmenting the data set of medial mesh evaluations with a number of the expensive solid mesh evaluations within a multi-fidelity surrogate modelling framework. In the current work, Kriging and its auto-regressive multi-fidelity variant, Co-Kriging (Kennedy and O'Hagan 2000), are used, as the formulation provides estimates of the prediction error which can be subsequently used to search for infill points. Toal et al. (2014) demonstrated the use of Co-Kriging models for reducing the frequency of performing expensive whole engine transient thermo-mechanical simulations in the optimization of a high-pressure compressor for minimum specific fuel

consumption. The thermal results were supplemented with a larger number of relatively cheap steady-state mechanical analyses. Brooks et al. (2011) optimized a transonic compressor rotor for maximum isentropic efficiency using Co-Kriging models where the different levels of fidelities were defined by the mesh resolution used for CFD. Zhang et al. (2015) utilized single- and double-sector fluid volumes of a combustor feed arm geometry to construct Co-Kriging models which were subsequently used to minimize pressure loss. To the best of the authors' knowledge, structural optimization of whole engine models using mixed-dimensional finite element models and multi-fidelity surrogate models has not been demonstrated in the literature.

The following paper is organized as follows. Section 2 describes the medial object transformation process and how it is implemented in the MANTLE software, as well as a study to validate the simulation results of medial meshes against expensive solid meshes. Section 3 introduces Kriging and Co-Kriging. Section 4 presents two case studies to demonstrate the effectiveness of the multi-fidelity surrogate modelling framework for reducing optimization costs. Section 5 is a summary of our conclusions.

2 Medial mesh generation

2.1 The medial object

Medial object transformation is a process whereby a geometry is converted to a compact, lower-dimensional, skeletal form which preserves useful information about the properties of the original model, such as its topology and symmetry. Blum (1967) defined the medial object as the locus of a maximally-inscribed disc (for 2D shapes) or sphere (for 3D boundary representations of geometry) as it rolls within the boundaries of an input geometry (Fig. 1). This locus can take the form of $(\mathbf{x}, r(\mathbf{x}))$, where \mathbf{x} is the disc/sphere centre and r is the radius of the disc/sphere. The medial object can also be described in terms of minimum distances (Gürsoy and Patrikalakis 1991). For closed geometries, minimum distance-to-boundary values can be computed for all internal points. Interpolating the set of internal points that have non-unique minimum distance values also gives the medial object.

These definitions of the medial object are known to be challenging to apply for arbitrary geometries in a robust manner, as the medial object is highly sensitive to perturbations on the boundary. The formulation has since undergone significant changes to increase the stability of the computed medial object, as seen in the works of Tek and Kimia (2001), Katz and Pizer (2003), and Macrini et al. (2011), amongst many others. Here, we emphasize

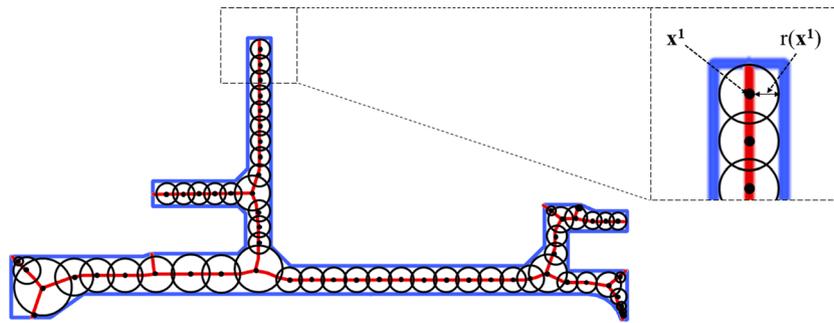


Fig. 1 A 2D shape boundary (blue curves) and its medial object (red curves) with no artificial branches at the end-faces (adapted from Wang et al. (2017)). The black dots and the solid black circles denote the medial points and the corresponding maximal discs respectively.

that the proprietary transformation algorithm in MANTLE does not implement any of these stability modifications, as it is currently only applied to engineering geometries which have surfaces that are smooth and well-defined. The reader is referred to Tagliasacchi et al. (2016) for a survey of the state of the art in skeletonization algorithms, including those for medial object transformation.

The key advantage of the medial object transformation approach for dimensionally reducing solid 3D geometries is its ability to accurately represent complex junctions. This is in contrast to conventional mid-surface generation tools in commercial CAD software that are commonly based on face-pairing mechanisms (Woo 2014; Zhu et al. 2016; Kulkarni et al. 2017). Figure 2a shows a cutout of an engine casing geometry which contains multiple structural members of different shapes, sizes, and orientations that intersect at several junctions. The corresponding surfaces that were generated using the mid-surface by face-pairing algorithm in Siemens NX9 Modelling are shown in Fig. 2b. The junctions consist of disconnected mid-surface patches, and this highlights the inability of these face-pairing methods for handling geometries with no obvious face pairs. Existing techniques to automatically extend and connect these patches are also unreliable. In comparison, the surfaces that were generated by the MANTLE software for the same geometry are fully connected (Fig. 2c).

2.2 MANTLE and the medial mesh workflow

The MANTLE software, first proposed in Stanley (2010), is currently implemented via a suite of tools consisting of MATLAB functions and Siemens NX plugins that were written using the Siemens NX9 Open API. The tools are grouped into three main packages: M-2D, M-3D, and M-Core. M-2D and M-3D are the primary tools for the medial object transformation of axisymmetric and non-axisymmetric geometries, respectively. They contain functions for medial point cloud generation, duplicate

Each medial point is defined by its coordinate x^i and the radius of its maximal disc $r(x^i)$, where i is the medial point identifier. Note that not all medial points and maximal discs are shown

node removal, and medial surface topology preservation. M-Core contains auxiliary operations that help complete the medial mesh generation process. This includes the definition of material properties, the transfer of thickness information between medial meshes, the application of boundary conditions, and post-processing of simulation results.

The MOT algorithm in MANTLE is unique in that it computes the medial point cloud from a surface mesh of the solid geometry, instead of from its boundary representation (Wang et al. 2014, 2017). This approach has several advantages:

1. The problem of artificial branch generation at the end-faces is eliminated since these faces can be left unmeshed and are thus ignored in the transformation process (see Fig. 1). Parallel computation of the medial object is enabled as the subdivision of geometries into segments no longer produces these branches at the location of the split-faces.
2. Boundary conditions and inter-component interface definitions that were applied to geometric entities can be transferred automatically to the medial mesh by tagging the surface nodes with information about their originating geometry.

The medial mesh generation process for the engine casing geometry from Fig. 2 is illustrated in Fig. 3.

1. The repeating pattern of the stator vane arrangement allows us to represent the casing using only a small segment of the original geometry. Tags are applied to the relevant geometric entities. The model decomposition and tagging tasks are currently performed manually in the modelling environment.
2. A surface mesh of the segment is generated. An element size which allows the resulting mesh to adequately represent existing curvatures (of fillets, for example) in the solid geometry is normally used.

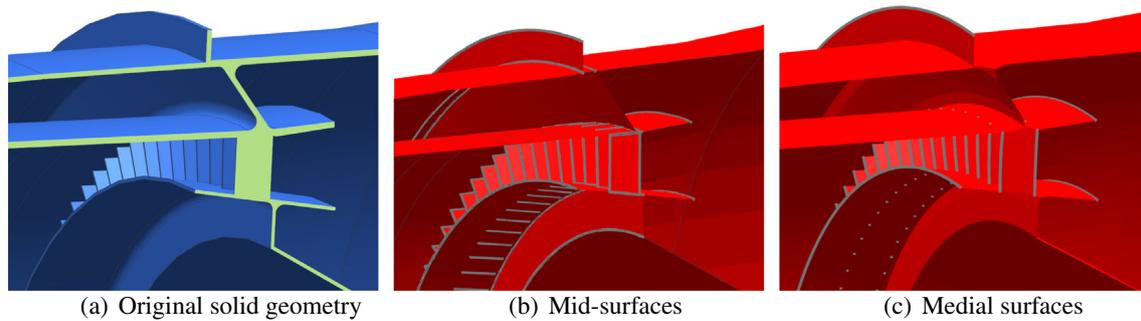


Fig. 2 Comparison of the **b** mid-surface by face-pairing algorithm in Siemens NX9 Modelling with the **c** medial surfaces from the medial object transformation algorithm in MANTLE for the complex junction

shown in **a**. Free edges are highlighted in grey. The mid-surfaces in **b** have been passed through the automatic edge stitching modules in Siemens NX9 Modelling

- M-3D takes the surface mesh as input, computes the medial points, and constructs the medial mesh of the segment. Tags that were originally applied to faces and edges are transferred automatically to element faces and nodes.
- M-Core processes the existing split-face tags and performs the necessary segment-duplication and node-merging operations to build the complete medial mesh of the original geometry. This mesh can be solved directly but is often converted to a shell polygon for further mesh sensitivity studies. In this case, M-Core is also able to transfer thickness information between medial meshes of different resolutions.

2.3 Validation of medial mesh simulations against solid mesh simulations

In this section, the simulation accuracy of medial meshes with respect to solid meshes is demonstrated through two studies involving different types of simulations. The first study is a static structural simulation using a simple whole engine model, where the deformations at several common points in both models are compared. The second study is a free-free modal simulation on the casing of a low-pressure turbine from the same engine, where comparisons are made

in terms of natural frequencies and mode shapes. Here we consider the evaluations from mesh-insensitive solid finite element models to be the most accurate source of structural information available to engine designers.

Figure 4 shows a section view of the whole engine. The geometry was built such that its longitudinal axis is parallel to the global x -axis, and its transverse plane is parallel to the global y - z plane. All studies were performed on a workstation which contained a 4-core Intel Xeon E5-4640 CPU clocked at 2.40 GHz and 256 GB of installed memory.

2.3.1 Static structural analysis

For the solid mesh, four-sided CTETRA tetrahedral elements are used. For the medial mesh, a combination of CTRIA3 triangular shell elements and CQUAD4 quadrilateral shell elements is used. For both cases, each of the four engine subsystems are meshed independently and then assembled using RBE3 interpolation constraint elements (Siemens 2014) that are generated by M-Core. The set of boundary conditions used are obtained from a sample simulation that had been previously performed in Rolls-Royce and they correspond to operating conditions at cruise (Toal et al. 2014). Thrust loads in the form of uni-directional forces are applied to every compressor and turbine stage,

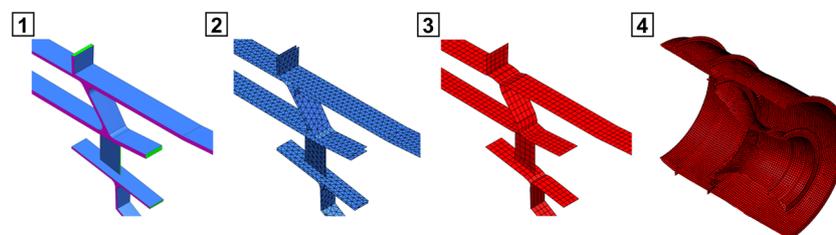


Fig. 3 The medial mesh generation workflow of an engine casing geometry within the MANTLE software. (1) The casing is reduced to a single-stator-vane segment. The end-faces (shown in green) and the split-faces (shown in pink) have been tagged. (2) A surface mesh of the segment is generated. The tagged faces are observed to have been excluded. (3) M-3D takes the surface mesh as input and generates its

medial mesh while preserving the tagging information. (4) M-Core uses the tags to identify split-faces and subsequently constructs the complete medial mesh of the original geometry. Here, a section view of the complete medial mesh is provided to show the duplication of the medial mesh segments from step 3

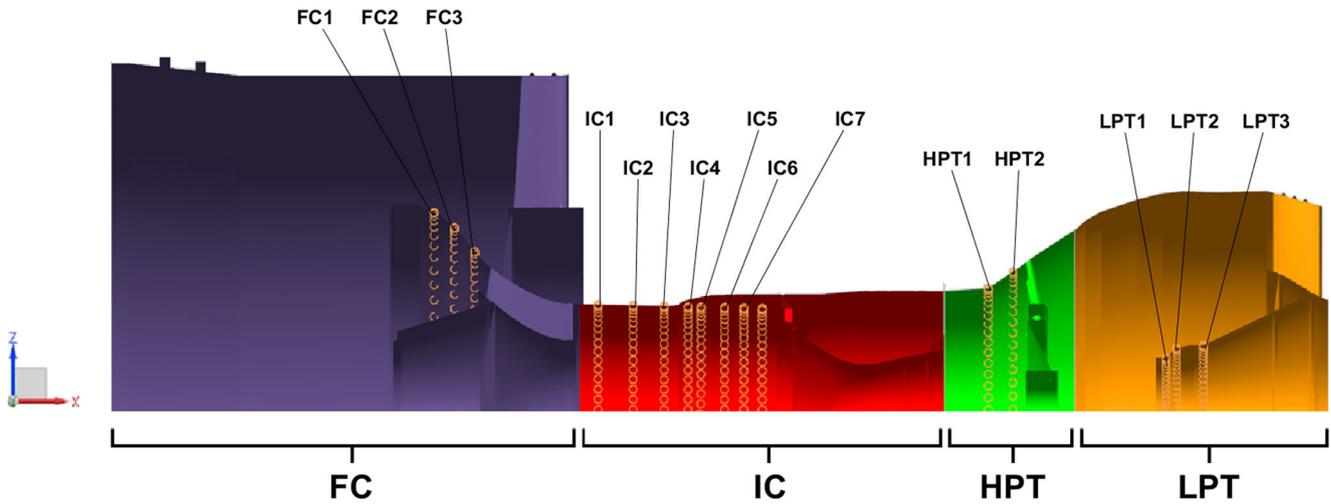


Fig. 4 Section view of the whole engine model, its subsystems, and the locations of the output sets that were used for post-processing in the static structural analysis. There are 4 subsystems in this engine

model: the fan casing (FC), the intermediate compressor casing (IC), the high-pressure turbine casing (HPT), and the low-pressure turbine casing (LPT)

while static pressure loads act normally on every surface. The engine is constrained in the *y*- and *z*-axes at the fan casing and low-pressure turbine casing mounts, and in the *x*-axis at the intermediate compressor casing thrust lugs. The simulations are solved using the SOL101 static structural solver in Siemens NX9 Nastran.

The displacements at a total of 960 points are extracted. The points are grouped into 15 output sets of 64 points each, and the points in each set are evenly distributed on the circumference of the WEM at a constant *x*-position (see Fig. 4). The output sets are post-processed to give maximum local radial displacement values which are then

used as design constraints in Section 4. This maximum local radial displacement value is an approximation of the maximum pinch point metric for tip clearance studies of gas turbine engine casings (Arkhipov et al. 2009). The calculation first involves the fitting of a least-squares circle to the deformed points in each output set on the *y-z* plane. The radial displacement of each point is then computed as its distance from the *y-z* plane origin minus the radius of the least-squares circle. The least-squares circles approach accounts for rigid body translation.

The results of the mesh sensitivity study are shown in Fig. 5. A percentage difference of < 1% over all maximum

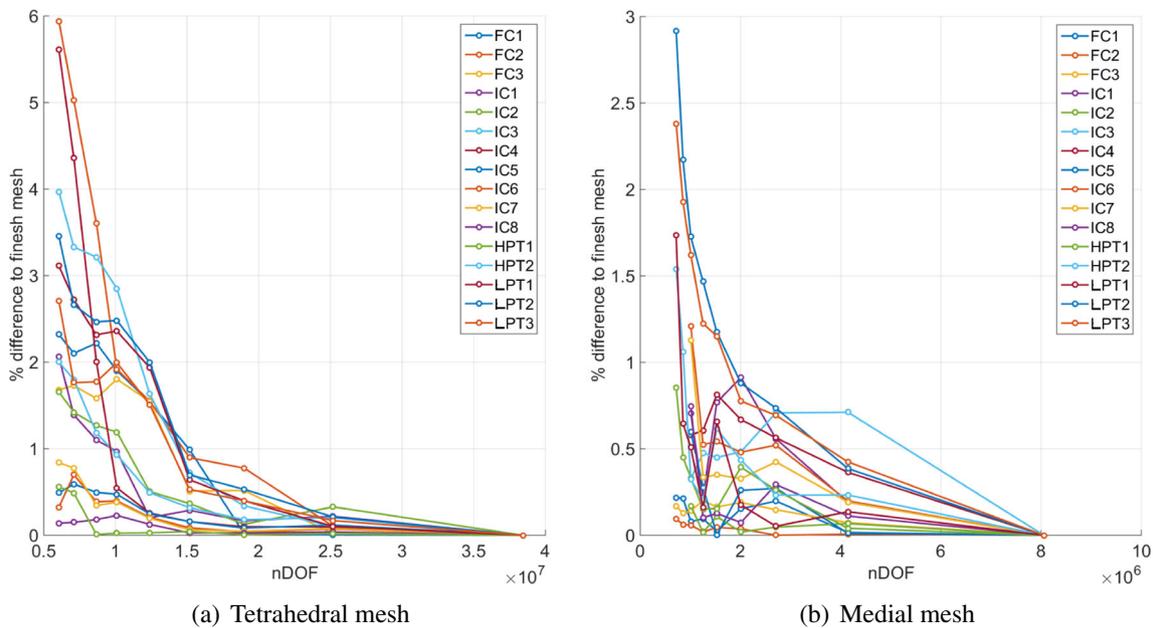


Fig. 5 Mesh sensitivity study of the solid and medial whole engine models. See Fig. 4 for a description of the response locations

Table 1 Comparison of the maximum local radial displacements of all outputs between the converged tetrahedral and medial WEMs

Output	TM (mm)	MM (mm)	% Difference
FC1	2.3334	2.2962	1.60
FC2	2.0973	2.0703	1.29
FC3	1.9489	1.9071	2.14
IC1	1.4540	1.4094	3.06
IC2	1.3291	1.2855	3.29
IC3	1.2579	1.1993	4.66
IC4	1.2112	1.1633	3.95
IC5	1.0742	1.0217	4.88
IC6	1.0168	0.9595	5.63
IC7	0.9278	0.8728	5.93
HPT1	0.2757	0.2701	2.04
HPT2	0.2013	0.1772	11.98
LPT1	0.1519	0.1387	8.71
LPT2	0.2292	0.2062	10.01
LPT3	0.3312	0.3022	8.76

local radial displacement values that are computed by the finest mesh of the corresponding type is taken to be the criteria for attaining mesh independence. Thus, the sixth coarsest meshes tested for both the solid and medial cases are considered to have converged. Table 1 compares the maximum local radial displacement values between the converged solid and medial meshes, and Table 2 is a summary of their computational requirements. For an average loss in accuracy of 5.2%, the medial-meshed engine model can be solved in 8.6% the wall time of the solid-meshed engine model.

2.3.2 Free-free modal analysis

To validate the medial mesh for dynamic analyses, the Modal Assurance Criterion (MAC) is used as it is regarded as a more suitable measure of accuracy than displacement results from static structural simulations (Chen 2001; Pastor

Table 2 Comparison of the computational requirements for solving the converged tetrahedral and medial whole engine models

Solution details	TM	MM
Element count	2793845	351635
Node count	5586772	337455
Solution time (HH:MM:SS)	7:34:58	00:39:03
Max memory usage (GB)	19.969	0.751
Max disk usage (GB)	90.284	39.824
Total I/O (GB)	555.955	189.693

et al. 2012; Towner and Band 2012). The MAC is a statistical indicator that is used to represent the similarity in mode shapes between two modal data sets. Besides being used for assessing mesh quality, the MAC is also commonly used to validate simulation results against physical experiments.

The MAC value between two mode shape vectors $\{\phi_r\}$ and $\{\phi_t\}$ is a normalized scalar product of the two vectors and is calculated as follows:

$$MAC(r_i, t_j) = \frac{|\{\phi_r\}_i^T \{\phi_t\}_j|^2}{(\{\phi_r\}_i^T \{\phi_r\}_i)(\{\phi_t\}_j^H \{\phi_t\}_j)} \quad (1)$$

The subscripts r and t denote the reference and test data sets respectively, while i and j are mode numbers. $\{\phi_t\}_j$ thus refers to the modal vector of the j th mode shape from the test data set. The superscript T is a conjugate transpose operator. The results from comparing n modal vectors from the two models can be arranged into an $n \times n$ MAC matrix. The MAC values can range from 0 (no correlation in mode shapes) to 1 (perfect modal correlation). An accurate model would thus produce a MAC matrix that has values near 1 in the main diagonal and near-zero values everywhere else.

The meshes used in this study are the low-pressure turbine portions of the converged whole engine meshes from the previous study. The meshes were solved using the

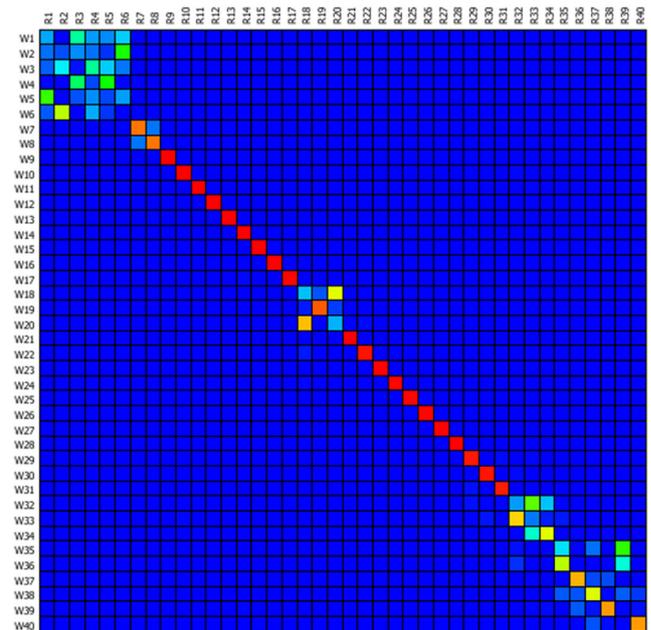


Fig. 6 The matrix of Modal Assurance Criterion values for the first 40 modes of the LPT casing. The horizontal axis, R (reference model), and the vertical axis, W (working model), represent the modes for the solid and medial meshes respectively. Modes 7 to 17 are critical for the current modal evaluation (modes 1 to 6 are rigid body modes and hence have near-zero frequencies, while modes 18 and above occur at frequencies past the maximum operating rotational frequency of the low-pressure turbine)

Table 3 Comparisons of Modal Assurance Criterion values and natural frequencies in the main diagonal of the matrix in Fig. 6

Mode	MAC value	TM (Hz)	MM (Hz)	% Difference
7	0.89	176.60	178.55	1.10
8	0.89	176.68	178.68	1.13
9	1.00	292.08	291.70	0.13
10	1.00	446.05	452.98	1.55
11	1.00	446.56	453.50	1.56
12	1.00	794.33	813.85	2.46
13	1.00	813.87	834.56	2.54
14	1.00	876.62	891.13	1.66
15	1.00	876.98	891.68	1.68
16	1.00	885.95	901.37	1.74
17	1.00	926.73	931.42	0.51

SOL103 real eigenvalue solver from Siemens NX9 Nastran under free-free conditions. Figure 6 shows the MAC matrix plotted in the form of a heat map, while Table 3 lists the MAC values in the main diagonal of the matrix. Within the range of interest, all main diagonal MAC values are near 1, while all off-diagonal MAC values are approximately zero. These results show that the medial mesh is also a good replacement for the expensive solid mesh for dynamic analyses.

3 Kriging and Co-Kriging

3.1 Kriging

Kriging (Sacks et al. (1989)) is a popular surrogate modeling approach for facilitating engineering design studies that are based on deterministic computational responses. It is particularly useful for predicting responses from numerical simulations as it only makes assumptions about the smoothness of the response surface but not the degrees of freedom. The Kriging formulation also lends itself to the derivation of a prediction error metric which is useful for gauging the quality of the surrogate when high simulation costs make it difficult to do so otherwise (Keane and Nair 2005; Forrester et al. 2008; Sóbester et al. 2014).

The smoothness assumption says that the observed values for two points that are close together in the design space should be similar. The similarity in values, or correlation, between two points $\mathbf{x}^{(i)}$ and $\mathbf{x}^{(j)}$ in a k -dimensional design space can be expressed as,

$$R_{i,j} = \text{Corr}[f(\mathbf{x}^{(i)}), f(\mathbf{x}^{(j)})] = \exp\left(-\sum_{l=1}^k \theta_l |x_l^{(i)} - x_l^{(j)}|^{p_l}\right). \tag{2}$$

A correlation matrix, \mathbf{R} , can then be formed from all possible pairs of points in a given training data set containing n points,

$$\mathbf{R} = \begin{pmatrix} R_{1,1} & \cdots & R_{1,n} \\ \vdots & \ddots & \vdots \\ R_{n,1} & \cdots & R_{n,n} \end{pmatrix}. \tag{3}$$

θ_l and p_l are the hyperparameters of the l th design variable which control the rate of correlation decrease and the degree of smoothness, respectively. A maximum likelihood approach is used to find the values for these hyperparameters which best return the observed function values, \mathbf{f} , for the training data. The natural logarithm of the likelihood function can be written as,

$$\ln(L) = -\frac{n}{2} \ln(2\pi) - \frac{n}{2} \ln(\sigma^2) - \frac{1}{2} \ln |\mathbf{R}| - \frac{(\mathbf{f} - \mathbf{1}\mu)^T \mathbf{R}^{-1} (\mathbf{f} - \mathbf{1}\mu)}{2\sigma^2}. \tag{4}$$

Setting the partial derivatives of $\ln(L)$ to zero gives us the maximum likelihood estimates (MLE) of μ and σ^2 ,

$$\hat{\mu} = \frac{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{f}}{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{1}}, \tag{5}$$

$$\hat{\sigma}^2 = \frac{(\mathbf{f} - \mathbf{1}\hat{\mu})^T \mathbf{R}^{-1} (\mathbf{f} - \mathbf{1}\hat{\mu})}{n}. \tag{6}$$

Substituting the MLEs of μ and σ^2 into (4) gives us what is called the concentrated likelihood function,

$$\ln(L) = -\frac{n}{2} \ln(\hat{\sigma}^2) - \frac{1}{2} \ln |\mathbf{R}|, \tag{7}$$

which is then maximized to obtain the MLEs of the hyperparameters. For all Kriging and Co-Kriging models in the current work, a hybridized particle swarm algorithm was employed by the proprietary optimization software OPTIMATv2 (Toal et al. 2008, 2009, 2011) to optimize the hyperparameters.

With the hyperparameters tuned, the Kriging model can now be used to predict the response at unsampled points. The vector of correlations between an unsampled point, $\mathbf{x}^{(*)}$, and the points in the current training set can be expressed as,

$$\mathbf{r} = \begin{pmatrix} \text{Corr}[f(\mathbf{x}^{(1)}), f(\mathbf{x}^{(*)})] \\ \vdots \\ \text{Corr}[f(\mathbf{x}^{(n)}), f(\mathbf{x}^{(*)})] \end{pmatrix} = \begin{pmatrix} R_{1,*} \\ \vdots \\ R_{n,*} \end{pmatrix}. \tag{8}$$

The augmented correlation matrix for this new data set thus becomes,

$$\tilde{\mathbf{R}} = \begin{pmatrix} \mathbf{R} & \mathbf{r} \\ \mathbf{r}^T & 1 \end{pmatrix}. \tag{9}$$

Substituting $\tilde{\mathbf{R}}$ into the last term of (4) produces the augmented ln-likelihood function which can be subsequently

maximized to give the Kriging prediction of the response value at $\mathbf{x}^{(*)}$,

$$\hat{f}(\mathbf{x}^{(*)}) = \hat{\mu} + \mathbf{r}^T \mathbf{R}^{-1} (\mathbf{f} - \mathbf{1}\hat{\mu}). \tag{10}$$

This predictor can be used by a global optimizer to quickly search and exploit promising regions in the design space. For a more exploratory approach, the mean square error (MSE) of the Kriging predictor can be used to drive the optimizer towards sparsely sampled regions. The MSE estimate is,

$$s^2(\mathbf{x}^{(*)}) = \hat{\sigma}^2 [1 - \mathbf{r}^T \mathbf{R}^{-1} \mathbf{r}]. \tag{11}$$

The MSE allows us to model the uncertainty in the predictions as a Gaussian distribution, $F(\mathbf{x})$, with mean \hat{f} and variance s^2 . A prediction's improvement over the current best objective, f_{\min} , can be defined as,

$$I(\mathbf{x}^{(*)}) = f_{\min} - F(\mathbf{x}^{(*)}). \tag{12}$$

The expected value of this improvement is thus,

$$\begin{aligned} EI = \frac{f_{\min} - \hat{f}(\mathbf{x}^{(*)})}{2} \left[1 + \operatorname{erf} \left(\frac{f_{\min} - \hat{f}(\mathbf{x}^{(*)})}{s\sqrt{2}} \right) \right] \\ + \frac{s}{\sqrt{2\pi}} \exp \left[-\frac{(f_{\min} - \hat{f}(\mathbf{x}^{(*)}))^2}{2s^2} \right], \end{aligned} \tag{13}$$

where $\operatorname{erf}()$ is the error function.

By progressively adding maximum EI designs to the data set, a balance between exploration and exploitation of the design space is achieved. The three metrics that have been presented (predicted optimum, maximum error, and maximum expected improvement) are used in the case studies in the present work.

3.2 Co-Kriging

When there are multiple analysis models of varying accuracies and evaluation costs, a multi-fidelity surrogate modelling approach can be used to make full use of the available data. In this work, the auto-regressive, multi-fidelity extension of Kriging, called Co-Kriging (Kennedy and O'Hagan 2000; Forrester et al. 2007), is used. The

auto-regressive model says that observations of the high-fidelity response are correct and that any inaccuracies in the Co-Kriging model are caused entirely by the low-fidelity response. This is equivalent to approximating the high-fidelity response by adding a scaled Gaussian process of the low-fidelity response to a Gaussian process of the difference between the two data sets. The reader is referred to Park et al. (2017) for a survey of multi-fidelity surrogate modelling methods, in which the Co-Kriging approach used here is classified as a Bayesian discrepancy framework with a non-informative prior.

The covariance between the two data sets can be expressed using a covariance matrix, \mathbf{C} , as,

$$\begin{aligned} \mathbf{C} &= \begin{pmatrix} \mathbf{C}_{1,1} & \mathbf{C}_{1,2} \\ \mathbf{C}_{2,1} & \mathbf{C}_{2,2} \end{pmatrix}, \\ \mathbf{C}_{1,1} &= \sigma_c^2 \mathbf{R}_c(\mathbf{X}_c, \mathbf{X}_c), \\ \mathbf{C}_{1,2} &= \rho \sigma_c^2 \mathbf{R}_c(\mathbf{X}_c, \mathbf{X}_e), \\ \mathbf{C}_{2,1} &= \rho \sigma_c^2 \mathbf{R}_c(\mathbf{X}_e, \mathbf{X}_c), \\ \mathbf{C}_{2,2} &= \rho^2 \sigma_c^2 \mathbf{R}_c(\mathbf{X}_e, \mathbf{X}_e) + \sigma_d^2 \mathbf{R}_d(\mathbf{X}_e, \mathbf{X}_e). \end{aligned} \tag{14}$$

The subscripts e , c , and d denote the expensive, cheap, and difference models respectively. ρ is the scaling factor. \mathbf{X} is the data set of each model fidelity. The correlation matrices \mathbf{R}_c and \mathbf{R}_d follow the form of (2) and (3).

There are now five hyperparameters to tune (two for each correlation matrix, as well as ρ). The MLE approach from (4) to (6) can be used but with \mathbf{f} replaced by a set of difference values between the high- and low-fidelity data sets:

$$\mathbf{d} = \mathbf{f}_e - \rho \mathbf{f}_c(\mathbf{X}_e). \tag{15}$$

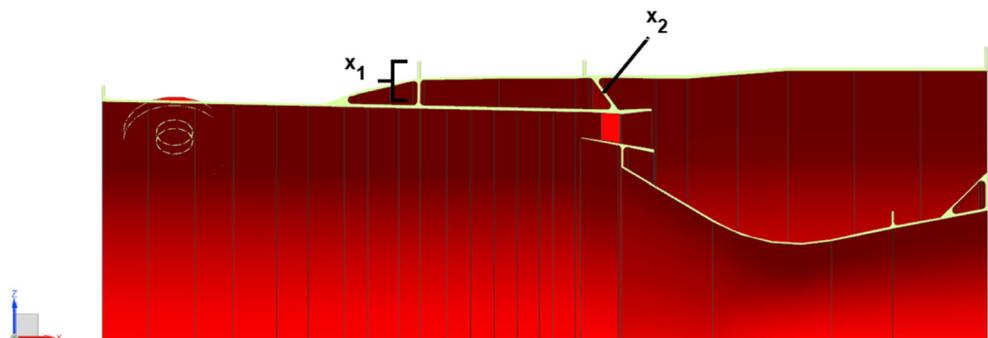
with the hyperparameters estimated, the Co-Kriging predictor can be derived as,

$$\hat{f}_e(\mathbf{x}^{(*)}) = \hat{\mu} + \mathbf{c}^T \mathbf{C}^{-1} (\mathbf{f} - \mathbf{1}\hat{\mu}), \tag{16}$$

where \mathbf{c} is a column vector of covariances between the existing set of observations and a new point. The MSE estimate is given as,

$$s_e^2(\mathbf{x}^{(*)}) = \rho^2 \hat{\sigma}_c^2 + \hat{\sigma}_d^2 - \mathbf{c}^T \mathbf{C}^{-1} \mathbf{c} + \frac{1 - \mathbf{1}^T \mathbf{C}^{-1} \mathbf{c}}{\mathbf{1}^T \mathbf{C}^{-1} \mathbf{1}}. \tag{17}$$

Fig. 7 Section view of the engine intercase and the design variables x_1 and x_2 . x_1 and x_2 are chosen such that their thickness modification would not affect the gas path. x_1 consists of the second flange and the vertical rib which separates the two void boxes, while x_2 consists of the diagonal rib which separates the second void box and the combustor region



4 Application of the multi-fidelity optimization with a medial mesh

In this section, two case studies are presented to demonstrate the benefits of using Co-Kriging surrogate models of solid and medial mesh evaluations (abbreviated as the CK strategy) for the structural optimization of engine geometries. The optimum designs found using this multi-fidelity approach, as well as the optimization times, are assessed with respect to a reference single-fidelity approach which is solely constructed from expensive tetrahedral meshes (the K_{TM} strategy). The properties of a single-fidelity approach employing only medial meshes (the K_{MM} strategy) are also examined.

4.1 Two-variable optimization of an engine intercasings

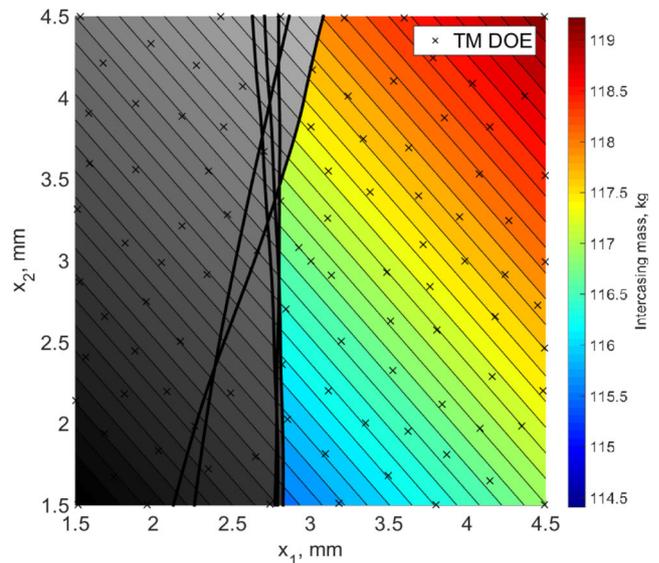
The intercasings from the whole engine geometry shown in Fig. 4 is chosen to be the subject of the first case study. The objective is to minimize the mass of the intercasings, f , while ensuring that its stiffness in terms of the maximum local radial displacement values as described in Section 2.3.1 is not reduced. Figure 7 shows the intermediate compressor casing geometry and the location of the design variables. The variable bounds are set to $\pm 50\%$ of the nominal dimensions. A total of seven design constraints, IC1 to IC7 (see Fig. 4), are present in the current study. The upper limits of the constraints are the maximum local radial displacement values of the nominal design. Preliminary results showed that the noise in the FEA simulations prevented consistently accurate evaluations of the constraints, so the constraint limits were increased by a factor of 0.1% (ranging from $1.519E-04$ to $2.333E-03$ mm) of the nominal values.

Thickness modifications for the solid geometry are realized using the Synchronous Modelling functions in Siemens NX9 Modelling. The geometry is meshed according to the results from the mesh sensitivity studies that were

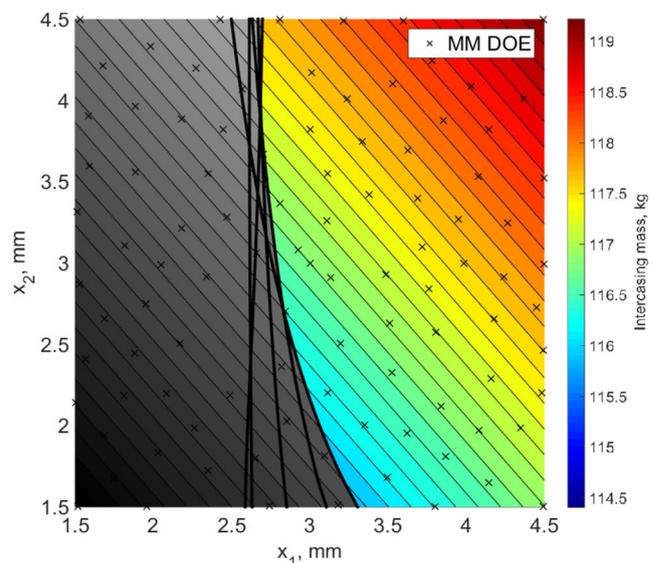
Table 4 Comparison of simulation results using the converged medial and tetrahedral meshes for 100 intermediate compressor casing designs

Output	r^2	% RMSE	% MAE
IC1	0.29	4.16E+3	4.21E+3
IC2	0.72	3.03E+3	3.06E+3
IC3	0.95	7.81E+2	7.99E+2
IC4	0.93	1.83E+2	1.95E+2
IC5	0.99	5.44E+2	5.57E+2
IC6	0.99	6.98E+2	7.07E+2
IC7	0.96	9.04E+2	9.15E+2

performed in Section 2.3.1. The meshing parameters for generating the tetrahedral meshes are taken to be valid across the design space considered and thus remain constant throughout this study. For the medial object, the converged medial mesh from the mesh sensitivity studies is used directly. The process of realizing design changes in the medial mesh is also straightforward as it only involves overwriting the thicknesses of the groups of elements that are associated with each design variable.



(a) TM response surface.



(b) MM response surface.

Fig. 8 Contour plots of Kriging predictions for the 100-point space-filling DOE. The contour plots are generated using Kriging predictions from a 501×501 grid. The coloured area represents the feasible design space. The bold lines represent the constraint limit curves. The colour gradients for both plots have been scaled to the same range. TM, tetrahedral mesh; MM, medial mesh

Table 5 Optimization results for the three surrogate modelling strategies in the intercase case study. The outputs of the optimum designs have been normalized against the nominal design

Optimum parameters	K _{MM}			K _{TM}			CK		
	CI _{2.5}	μ	CI _{97.5}	CI _{2.5}	μ	CI _{97.5}	CI _{2.5}	μ	CI _{97.5}
x_1	3.3059	3.3059	3.3059	2.8310	2.8351	2.8407	2.8332	2.8430	2.8680
x_2	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000
f	0.9924	0.9924	0.9924	0.9890	0.9890	0.9891	0.9890	0.9891	0.9892
IC1	0.9989	0.9989	0.9989	0.9989	0.9989	0.9989	0.9989	0.9989	0.9989
IC2	0.9989	0.9989	0.9989	0.9988	0.9988	0.9988	0.9988	0.9988	0.9988
IC3	0.9983	0.9983	0.9983	0.9988	0.9988	0.9989	0.9987	0.9988	0.9988
IC4	0.9980	0.9980	0.9980	0.9980	0.9981	0.9981	0.9979	0.9980	0.9981
IC5	0.9984	0.9984	0.9984	0.9996	0.9997	0.9997	0.9994	0.9996	0.9997
IC6	0.9991	0.9991	0.9991	0.9999	0.9999	1.0000	0.9997	0.9999	0.9999
IC7	1.0000	1.0000	1.0000	0.9997	0.9998	0.9998	0.9996	0.9998	0.9998

μ is the average output over the 10 repeated runs. The confidence intervals (CI_{2.5} and CI_{97.5}) are calculated through resampling with 50,000 bootstrap samples. The constraint satisfaction values for the K_{MM} strategy are the results from evaluations of the optimum designs using tetrahedral meshes

A preliminary assessment of the suitability of constructing Co-Kriging models from the tetrahedral and medial meshes could be performed. This involved a comparison of simulation results for a set of 100 designs as specified by a space-filling DOE. Table 4 shows the comparison results. Here, the root mean squared error (RMSE) and maximum absolute error (MAE) are presented as a percentage of the range of response values in the observed data. Note that the mass is not compared as in both cases it is obtained directly from the solid CAD model and is thus the same.

The r^2 correlation coefficient values for the outputs IC3 to IC7 satisfy the $r^2 > 0.9$ recommendation in Toal (2015). The low r^2 values for IC1 and IC2 are attributed to the low sensitivities of the two responses toward changes in x_1 and x_2 , where the range of observed values are an

order of magnitude smaller than the remaining outputs. The low sensitivities may be explained by how IC1 and IC2 are relatively far away from x_1 and x_2 , and that their responses are already largely dictated by the presence of the thrust lugs. The small ranges also contributed to the large percentage error values. For this case study, a surrogate model that is constructed from only medial mesh evaluations would thus be able to predict the locations of optima but not their response values.

The contour plots in Fig. 8 show the feasible design space of each mesh type as predicted by the Kriging models that are constructed from the space-filling data sets. Differences in the constraint limits are noticeable. The output IC6 is predicted to be the active constraint at the K_{TM} optimum, while the output IC7 is predicted to be the active constraint

Table 6 r^2 correlation between predictions from the surrogate modelling strategies at the end of the optimizations and the observed values for the 100-point space-filling data set. The best strategy for each output is italicized

Output	K _{MM}			K _{TM}			CK		
	CI _{2.5}	μ	CI _{97.5}	CI _{2.5}	μ	CI _{97.5}	CI _{2.5}	μ	CI _{97.5}
IC1	0.29	0.29	0.29	<i>0.87</i>	<i>0.89</i>	<i>0.90</i>	0.71	0.80	0.86
IC2	0.72	0.72	0.72	<i>1.00</i>	<i>1.00</i>	<i>1.00</i>	0.94	0.96	0.97
IC3	0.95	0.95	0.95	<i>0.97</i>	<i>0.97</i>	<i>0.97</i>	0.93	0.96	0.97
IC4	0.93	0.93	0.93	<i>1.00</i>	<i>1.00</i>	<i>1.00</i>	0.99	0.99	0.99
IC5	0.99	0.99	0.99	<i>1.00</i>	<i>1.00</i>	<i>1.00</i>	1.00	1.00	1.00
IC6	0.99	0.99	0.99	<i>1.00</i>	<i>1.00</i>	<i>1.00</i>	0.96	0.98	0.99
IC7	0.96	0.96	0.96	<i>1.00</i>	<i>1.00</i>	<i>1.00</i>	0.97	0.99	0.99

Table 7 Percentage RMSE of predictions from the surrogate modelling strategies at the end of the optimizations as compared to the observed values for the 100-point space-filling data set. The best strategy for each output is italicized

Output	K_{MM}			K_{TM}			CK		
	CI _{2.5}	μ	CI _{97.5}	CI _{2.5}	μ	CI _{97.5}	CI _{2.5}	μ	CI _{97.5}
IC1	4.16E+2	4.16E+2	4.16E+2	<i>7.15</i>	<i>7.62</i>	<i>8.34</i>	9.61	11.42	15.06
IC2	3.03E+3	3.03E+3	3.03E+3	<i>0.77</i>	<i>0.81</i>	<i>0.83</i>	0.94	0.96	0.97
IC3	7.81E+2	7.81E+2	7.81E+2	4.72	4.83	4.93	<i>0.93</i>	<i>0.96</i>	<i>0.97</i>
IC4	1.83E+2	1.83E+2	1.83E+2	1.66	1.73	1.81	<i>0.99</i>	<i>0.99</i>	<i>0.99</i>
IC5	5.44E+2	5.44E+2	5.44E+2	1.19	1.25	1.32	<i>1.00</i>	<i>1.00</i>	<i>1.00</i>
IC6	6.98E+2	6.98E+2	6.98E+2	1.03	1.07	1.11	<i>0.96</i>	<i>0.98</i>	<i>0.99</i>
IC7	9.04E+2	9.04E+2	9.04E+2	<i>0.81</i>	<i>0.86</i>	<i>0.92</i>	0.97	0.99	0.99

at the K_{MM} optimum. x_2 is expected to go to its lower bound for both of the single-fidelity strategies. However, x_1 is expected to be at ≈ 2.8 and ≈ 3.3 for the K_{TM} and K_{MM} optimum designs, respectively. The optimum design from the K_{TM} strategy is also expected to offer more mass reduction.

The K_{TM} and K_{MM} strategies are initiated with the same 20-point Latin Hypercube DOE, following the 10d recommendation from Jones et al. (1998). The same DOE is also taken to be the initial low-fidelity data set for the CK strategy, while the high-fidelity data set is a 5-point subset of the low-fidelity data that is obtained through the Morris-Mitchell space-filling procedure (Morris and Mitchell 1995).

Each strategy is run for 4 update iterations, where the surrogate models are retrained at the end of each iteration. The K_{TM} and K_{MM} surrogates are updated with 2 maximum EI points, 2 maximum RMSE points, and 1 point at the predicted optimum per iteration, giving a total of 40 evaluations at the end of the optimization. The CK strategy is updated with 1 maximum EI point, 1 maximum RMSE point, and 1 point at the predicted optimum for each mesh fidelity per iteration, ending with a data set of 17 and 32 high- and low-fidelity evaluations respectively.

Each strategy is repeated with 10 differently seeded Latin Hypercube DOEs of the same size.

Table 5 summarizes the optimization results. The optimum design parameters for the K_{TM} and CK strategies are approximately the same across all 10 runs. The mass reduction over the nominal design is $\approx 1.10\%$, with an active constraint IC6 (a constraint with a satisfaction value greater than 0.9995 is taken to be active as the pre-set minimum search distance parameter of $1E-3$ in the optimizer prevented any further evaluations in the feasible design space that is close to the constraint boundaries). The K_{MM} strategy is not able to find an equally good optimum design. The average mass reduction is 69% of the mass reduction achieved by the other two strategies. However, when evaluated using tetrahedral meshes, the K_{MM} optimum designs are able to satisfy the constraint limits in the K_{TM} strategy, as predicted by Fig. 8b.

To assess the global accuracy of the surrogate models, the final surrogate models from each optimization run are used to generate predictions for the 100-point space-filling DOE that is used in the preliminary test (see Table 4). The predictions are compared against the observed values. Tables 6, 7, and 8 show the comparison results. The CK strategy is able to significantly improve the r^2 correlation

Table 8 Percentage MAE of predictions from the surrogate modelling strategies at the end of the optimizations as compared to the observed values for the 100-point space-filling data set. The best strategy for each output is italicized

Output	K_{MM}			K_{TM}			CK		
	CI _{2.5}	μ	CI _{97.5}	CI _{2.5}	μ	CI _{97.5}	CI _{2.5}	μ	CI _{97.5}
IC1	4.21E+3	4.21E+3	4.21E+3	<i>20.56</i>	<i>23.12</i>	<i>26.46</i>	23.32	27.18	33.27
IC2	3.06E+3	3.06E+3	3.06E+3	<i>2.18</i>	<i>2.44</i>	<i>2.71</i>	10.42	11.74	13.78
IC3	7.98E+2	7.98E+2	7.98E+2	<i>11.98</i>	<i>13.20</i>	<i>15.04</i>	12.69	15.61	24.94
IC4	1.95E+2	1.95E+2	1.95E+2	4.32	4.79	5.26	5.11	5.97	6.91
IC5	5.57E+2	5.57E+2	5.57E+2	2.92	3.22	3.60	4.25	4.76	5.07
IC6	7.07E+2	7.07E+2	7.07E+2	2.77	3.07	3.36	5.51	9.36	16.39
IC7	9.15E+2	9.15E+2	9.15E+2	2.30	2.64	3.07	5.74	8.21	11.30

for IC1 and IC2 and also greatly reduce the percentage errors for all response predictions.

The improvement in the prediction accuracy of the Co-Kriging model can also be observed in Fig. 9b where its predictions of the constraint limits in the region of the optimum design are accurate with respect to the K_{TM}

strategy (Fig. 9a). The Co-Kriging predictions of the constraint limits in the region of $\{1.5 < x_1 < 3.0, 3.0 < x_2 < 4.5\}$ are less accurate as the region is sparsely sampled even with the influence of the maximum RMSE criterion. The predictions of the constraint limits in the optimum region from the K_{MM} surrogate model (Fig. 9c)

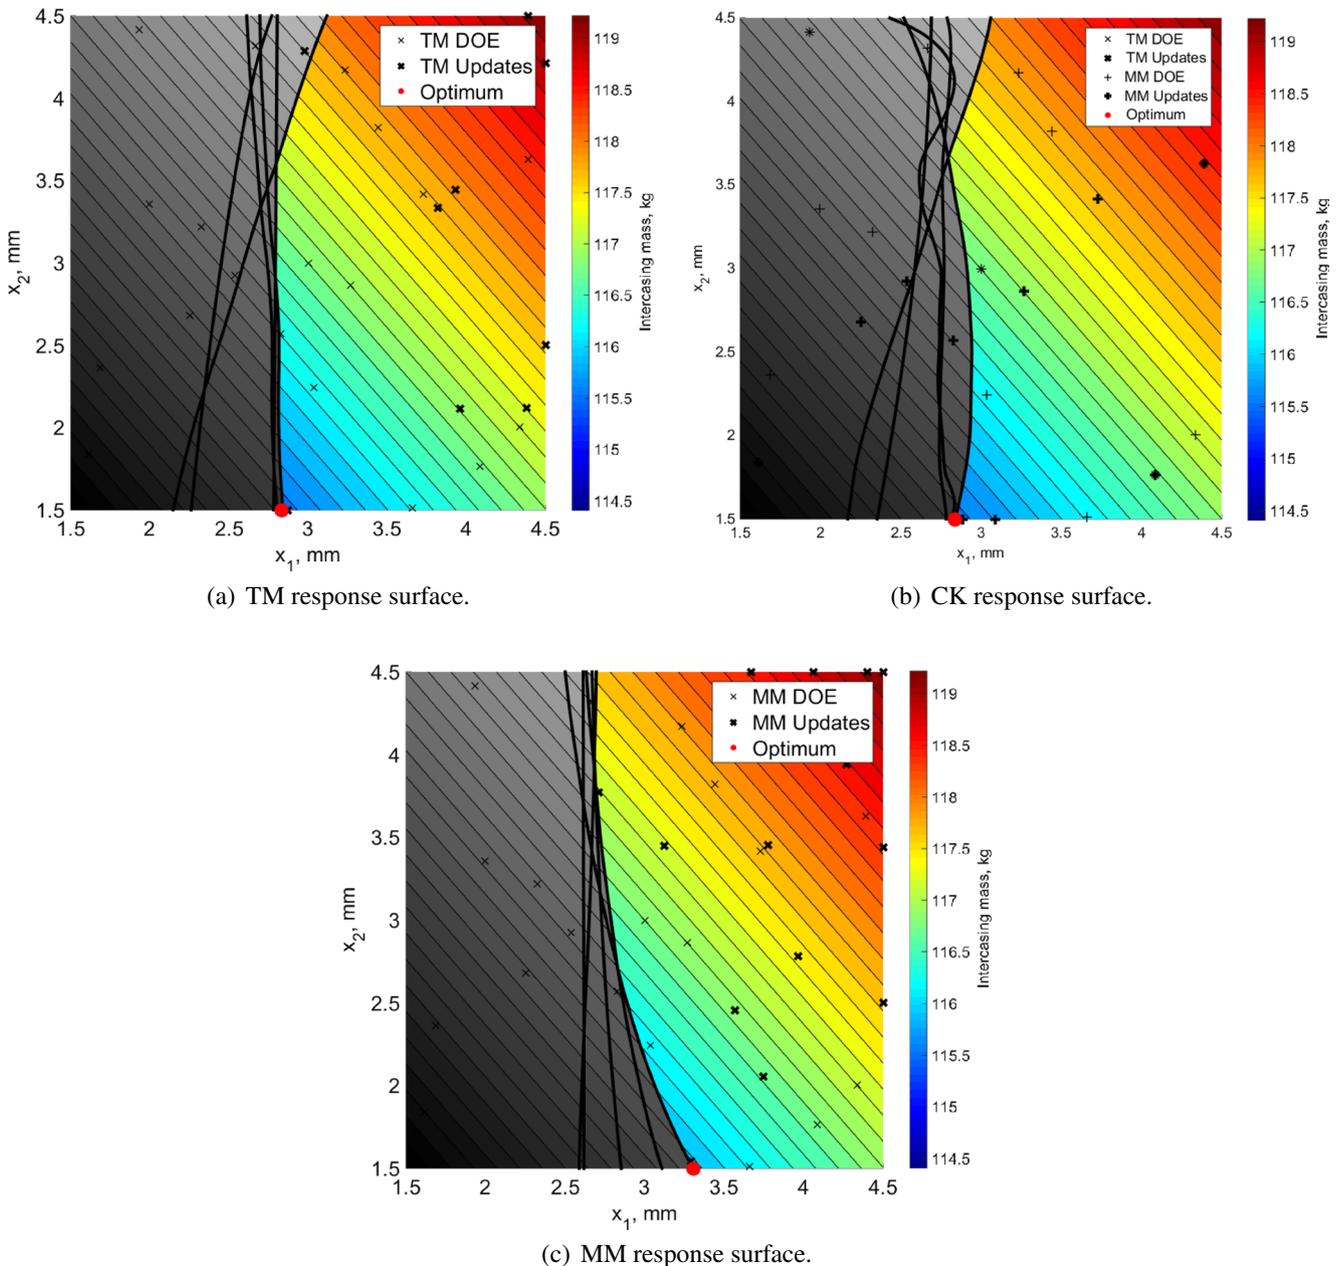


Fig. 9 Contour plots of predictions from the surrogate models in each strategy at the end of the optimization runs for the first of the 10 Latin Hypercube DOEs that were used. The contour plots were generated using Kriging predictions from a 501×501 grid. The bold lines

represent the constraint limit curves. The colour gradients for all three plots have been scaled to the same range. TM, tetrahedral mesh; MM, medial mesh

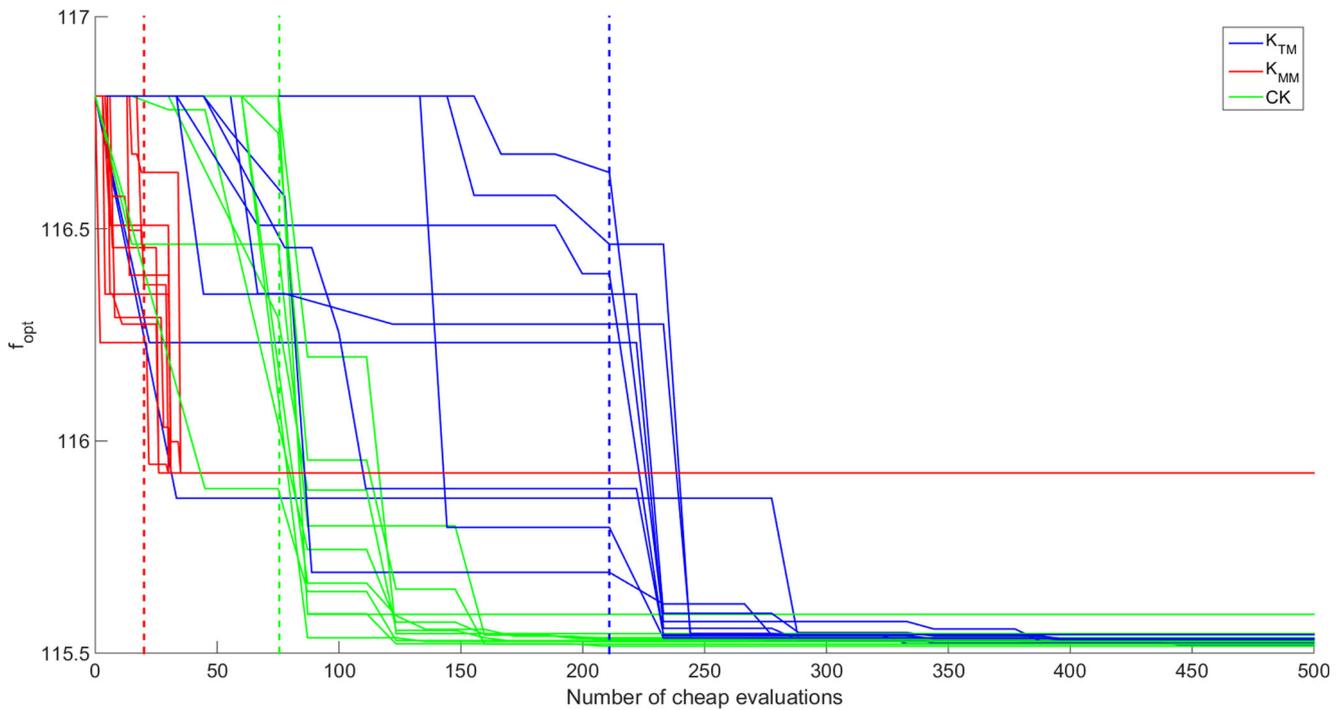


Fig. 10 Optimization histories for the three strategies in the intermediate compressor casing case study. The dashed vertical lines mark the end of the DOE evaluations as well as the start of the surrogate model updating procedure

are not as accurate, even with the high sampling frequency in the vicinity as driven by the maximum EI and predicted optimum criteria.

Figure 10 illustrates the optimization histories for all three strategies where the costs are plotted in units of equivalent number of cheap evaluations. The cost ratio of tetrahedral mesh to medial mesh evaluation times, averaged over all optimization runs, is 11.10. The cost savings

provided by the CK strategy can be clearly observed, with an average convergence cost of 52% when compared to the K_{TM} strategy.

4.2 Eight-variable optimization of a whole engine

Having clearly illustrated the benefits of the CK strategy for optimizing the engine inter casing, the multi-fidelity

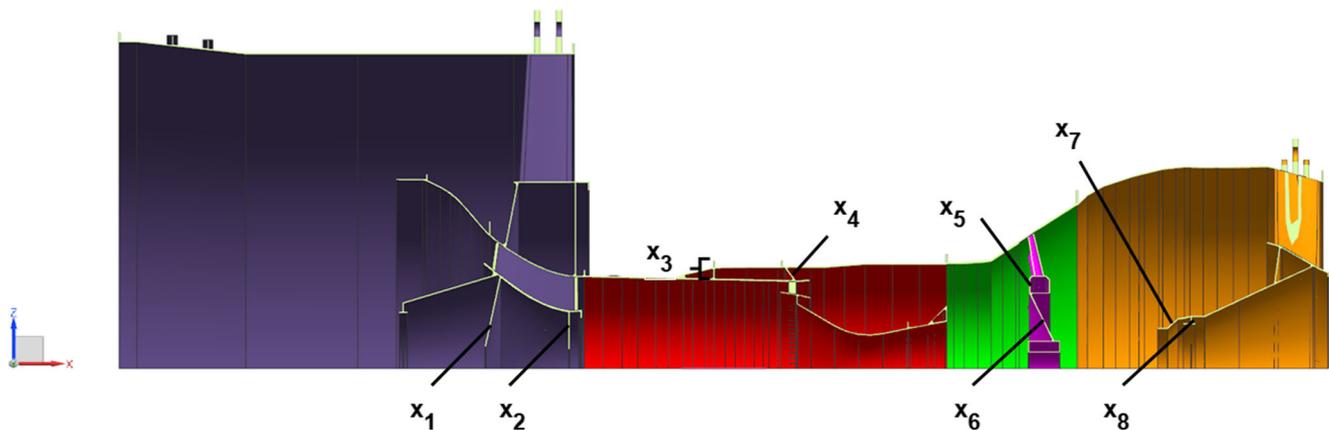


Fig. 11 Section view of the whole engine model and the design variables x_1 to x_8

optimization with medial meshes approach is now applied to a whole engine optimization study to demonstrate its effectiveness in solving realistic design problems involving large-scale engine assemblies.

The objective now is to minimize the whole engine mass, and the number of constraints have been increased to include all 15 maximum local radial displacement outputs across the entire engine (see Fig. 4). Figure 11 shows the eight design variables for the current optimization, where the thicknesses of two structural members for each of the three remaining subsystems have been added. Other than the increased problem size, the parameters of the optimization, as well as the manner in which the FEA simulations are performed, are the same as in Section 4.1.

The K_{TM} and K_{MM} strategies are initiated with an 80-point Latin Hypercube DOE. This DOE is also used as the initial low-fidelity data set for the CK strategy, while a 20-point Morris-Mitchell subset is extracted to be used as the high-fidelity data set. The longer evaluation times of the whole engine models made it infeasibly expensive to generate a space-filling data set of a suitable size for calculating solid and medial mesh correlations, as was done in the first case study. As such, the medial mesh accuracy is assessed based on the initial 80-point DOEs for the two single-fidelity strategies, where although the comparison results may not be as precise, they should still provide

some indication as to whether a multi-fidelity approach is suitable. Table 9 shows the comparison results.

The r^2 correlation coefficient values for all outputs are sufficiently high enough to warrant the use of medial meshes for building the Co-Kriging models. The percentage error values here are also significantly lower when compared to the first case study as the larger number of variables and the presence of inter-subsystem interactions resulted in larger ranges for each of the outputs.

Each strategy is run for 16 update iterations. The surrogate models in the K_{TM} and K_{MM} strategies are updated with 2 maximum EI points, 2 maximum RMSE points, and 1 point at the predicted optimum per iteration, giving a total of 160 evaluations at the end of the optimization. The surrogate models in the CK strategy are updated with 1 maximum EI point, 1 maximum RMSE point, and 1 point at the predicted optimum for each mesh fidelity per iteration, ending with a data set of 68 and 128 high- and low-fidelity evaluations respectively. Each strategy is repeated for a total of times times with the Latin Hypercube DOEs that have been evaluated.

Table 10 summarizes the optimization results. The mass reductions offered by both the K_{TM} and CK strategies are $\approx 1.40\%$. The output HPT1 is also the active constraint for both strategies. There is, however, some variation in the values for variables x_7 and x_8 in the CK strategy. The

Table 9 Comparison between the 80-point Latin Hypercube data sets of medial and solid WEM evaluations. μ is the average output over the 3 repeated runs. The confidence intervals ($CI_{2.5}$ and $CI_{97.5}$) were calculated through resampling with 50,000 bootstrap samples

Output	r^2			% RMSE			% MAE		
	$CI_{2.5}$	μ	$CI_{97.5}$	$CI_{2.5}$	μ	$CI_{97.5}$	$CI_{2.5}$	μ	$CI_{97.5}$
FC1	0.89	0.94	0.97	23.88	28.34	30.93	48.96	56.64	71.97
FC2	0.90	0.95	0.98	40.24	52.97	60.07	74.43	79.96	89.71
FC3	0.93	0.97	0.99	32.07	44.23	50.90	59.79	66.64	75.85
IC1	0.90	0.93	0.94	113.81	145.90	180.72	161.59	176.51	205.30
IC2	0.92	0.94	0.95	79.01	89.69	106.23	104.45	114.35	122.97
IC3	0.91	0.95	0.97	92.74	101.95	111.98	117.53	132.00	146.75
IC4	0.97	0.98	0.98	80.81	81.92	83.20	95.78	106.23	126.11
IC5	0.99	0.99	1.00	35.60	36.56	38.36	45.67	53.33	68.59
IC6	0.81	0.89	0.93	228.30	230.92	235.40	255.17	272.43	303.45
IC7	0.97	0.99	0.99	65.34	68.82	70.80	82.71	92.38	109.80
HPT1	0.87	0.89	0.90	117.98	126.93	135.19	153.58	159.83	170.74
HPT2	0.89	0.90	0.93	84.39	104.50	115.43	115.61	134.85	146.95
LPT1	0.97	0.97	0.97	58.22	61.59	63.47	70.00	74.26	76.51
LPT2	0.97	0.98	0.99	103.94	111.57	115.73	117.05	126.15	131.19
LPT3	0.95	0.96	0.97	121.45	121.68	122.11	143.37	144.56	145.26

Table 10 Optimization results for the three surrogate modelling strategies in the whole engine case study. The outputs of the optimum designs have been normalized against the nominal design. μ is the average output over the 3 repeated runs. The confidence intervals (CI_{2.5} and CI_{97.5}) were calculated through resampling with 50,000 bootstrap samples. The constraint satisfaction values for the K_{MM} strategy are the results from evaluations of the optimum MM designs using TMs. Constraint violations are italicized

Optimum parameters	K _{MM}			K _{TM}			CK		
	CI _{2.5}	μ	CI _{97.5}	CI _{2.5}	μ	CI _{97.5}	CI _{2.5}	μ	CI _{97.5}
x_1	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000
x_2	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000
x_3	1.6016	1.6021	1.6033	1.5142	1.5255	1.5372	1.5222	1.5360	1.5589
x_4	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000
x_5	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000
x_6	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000	1.5000
x_7	2.6806	2.8012	2.9303	3.8725	4.0041	4.0751	2.2121	2.9571	4.3224
x_8	2.6645	2.7268	2.7678	1.5000	1.5000	1.5000	1.5000	2.3408	2.8033
f	0.9864	0.9864	0.9864	0.9860	0.9860	0.9860	0.9861	0.9862	0.9862
FC1	0.9900	0.9900	0.9900	0.9908	0.9908	0.9909	0.9908	0.9908	0.9909
FC2	0.9894	0.9894	0.9894	0.9904	0.9904	0.9904	0.9904	0.9904	0.9904
FC3	0.9896	0.9896	0.9896	0.9905	0.9905	0.9905	0.9905	0.9905	0.9905
IC1	0.9879	0.9879	0.9879	0.9909	0.9909	0.9909	0.9909	0.9909	0.9909
IC2	0.9888	0.9888	0.9888	0.9909	0.9910	0.9910	0.9909	0.9910	0.9910
IC3	0.9754	0.9754	0.9754	0.9878	0.9879	0.9881	0.9876	0.9879	0.9882
IC4	0.9795	0.9795	0.9795	0.9926	0.9927	0.9928	0.9925	0.9926	0.9927
IC5	0.9893	0.9893	0.9893	0.9996	0.9997	0.9998	0.9993	0.9995	0.9997
IC6	0.9841	0.9841	0.9841	0.9897	0.9898	0.9899	0.9898	0.9898	0.9898
IC7	0.9834	0.9834	0.9834	0.9941	0.9941	0.9943	0.9940	0.9941	0.9941
HPT1	<i>1.0114</i>	<i>1.0114</i>	<i>1.0114</i>	0.9995	0.9998	0.9999	0.9995	0.9997	0.9999
HPT2	0.9810	0.9810	0.9810	0.9902	0.9906	0.9911	0.9901	0.9905	0.9912
LPT1	0.9341	0.9341	0.9341	0.9987	0.9989	0.9992	0.9968	0.9977	0.9984
LPT2	0.9331	0.9331	0.9331	0.9876	0.9876	0.9877	0.9865	0.9907	0.9933
LPT3	0.9593	0.9593	0.9593	0.9805	0.9806	0.9806	0.9801	0.9869	0.9907

3 sets of optimum values for the variables x_7 and x_8 are {2.0873, 2.8875}, {2.4615, 2.6348}, and {4.3224, 1.5000}. Two of the three optimum designs are found to have a more equal distribution of material between the two structural members. The remaining optimum design is observed to push x_8 to its lower bound, similar to the optimum designs found by the K_{TM} strategy.

The K_{MM} strategy here is able to find optimum designs that gave similar mass reductions to the two higher-fidelity strategies as well. However, when these optimum designs are evaluated using tetrahedral meshes, the constraint HPT1 is found to be violated by 1.14%, or approximately 3.13E−3 mm. The optimum designs are found to have a

good balance of material between x_7 and x_8 , similar to a majority of the CK optimum designs. Even so, the K_{MM} surrogate models are not able to predict the constraint responses to the levels of accuracy needed to prevent small violations of the constraint limits from occurring.

Figure 12 shows the optimization histories for all three strategies where the costs are plotted in units of equivalent number of cheap evaluations. The cost ratio of tetrahedral mesh to medial mesh evaluation times, averaged over all optimization runs, is 11.65. The cost savings provided by the CK strategy can again be clearly observed, with an average convergence cost of 57% when compared to the K_{TM} strategy.

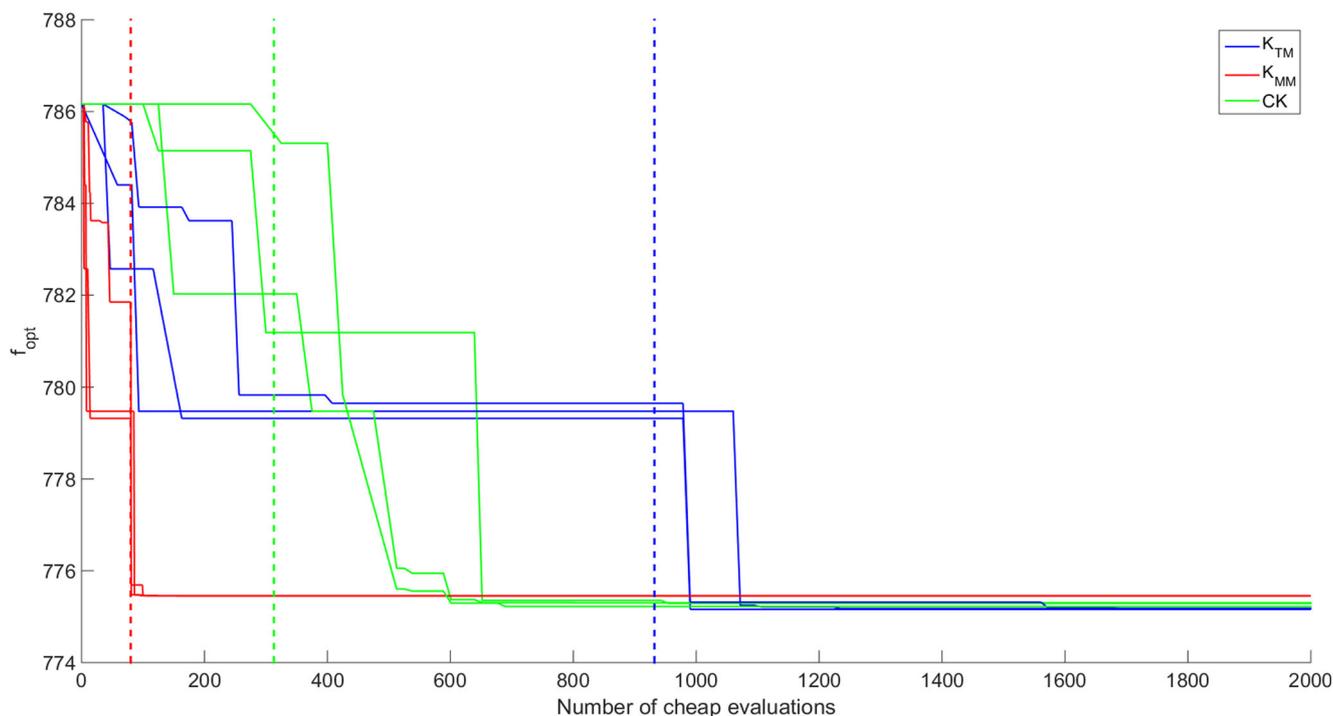


Fig. 12 Optimization histories for the three strategies in the whole engine case study. The dashed vertical lines mark the end of the DOE evaluations as well as the start of the surrogate model updating

procedure. Note that the K_{MM} runs converged to optimum designs that are slightly infeasible with respect to the K_{TM} constraints when they were evaluated using tetrahedral meshes

5 Conclusions

In this paper, a novel approach which combines medial meshing and multi-fidelity surrogate modelling techniques is proposed and demonstrated on two engine design problems. For the intermediate compressor casing problem, the Co-Kriging approach is able to find feasible optimum designs that are on par with the single-fidelity Kriging approach which solely uses solid mesh evaluations. However, the Co-Kriging approach was able to converge to its optimum at 52% the cost of the expensive single-fidelity Kriging approach. A single-fidelity Kriging approach which solely uses medial mesh evaluations was also tested and it is able to find optimum designs that are feasible with respect to the constraint limits of the high-fidelity strategy, even in the presence of active constraints. For the whole engine problem, the Co-Kriging approach was also able to find feasible optimum designs but at only 57% the cost of the expensive single-fidelity Kriging approach. The single-fidelity Kriging with medial meshes approach, however, found optimum designs that were infeasible with respect to the constraint limits of the high-fidelity strategy.

These results have demonstrated that the feasibility of conducting whole engine optimization studies using finite element models can be greatly improved through a combination of medial meshing and multi-fidelity surrogate

modelling techniques. The low manual effort required to generate these dimensionally-reduced whole engine geometries, as enabled by the MANTLE software, also means that the geometries are cheap and easy to update as individual subsystem designs progress throughout their design cycles. Future work aims to further reduce the cost of the multi-fidelity approach. Given the correlation and error values between two data sets of different fidelities, research for establishing the most efficient distribution of computational resources between the high-fidelity and low-fidelity evaluations is necessary. The robustness of the multi-fidelity approach also needs to be demonstrated using more complex optimization problems with multi-modal response surfaces.

6 Replication of results

The MANTLE software, as well as all finite element models generated using the software, are protected intellectual properties of Rolls-Royce. OPTIMATv2 is a Rolls-Royce proprietary optimization toolbox. The engine geometry shown in Figs. 2, 3, 4, 7, and 11 are created using a Rolls-Royce model provided through the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 234344 (www.crescendo-fp7.eu).

Acknowledgements The authors acknowledge Rolls-Royce for granting permission to publish this paper.

Funding information The research leading to these results has received funding from Rolls-Royce under the ASCEND (Aero Structures Collaborative Engine iNtegration and Design) project funded by CleanSky 2.

Compliance with Ethical Standards

Conflict of interests The author declares no conflict of interest.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Arkhipov AN, Karaban VV, Putchkov IV, Filkorn G, Kieninger A (2009) The whole-engine model for clearance evaluation. In: Structures and dynamics, Parts A and B, vol 6. ASME, Orlando, pp 9–17. <https://doi.org/10/dgmsg4>
- Bettebghor D, Blondeau C, Toal D, Eres H (2013) Bi-objective optimization of pylon-engine-nacelle assembly: Weight vs. tip clearance criterion. *Struct Multidiscip Optim* 48(3):637–652. <https://doi.org/10/f5cp9c>
- Blum H (1967) A transformation for extracting new descriptors of shape. In: Models for the perception of speech and visual form. MIT Press, Cambridge, pp 362–380
- Brooks CJ, Forrester AIJ, Keane AJ (2011) Multi-fidelity design optimisation of a transonic compressor rotor. In: 9th European conference of turbomachinery fluid dynamics and thermodynamics, Turkey
- Chen G (2001) FE model validation for structural dynamics PhD thesis. University of London, London
- Forrester AIJ, Sobester A, Keane AJ (2007) Multi-fidelity optimization via surrogate modelling. *Proceedings of the Royal Society A: Mathematical. Phys Eng Sci* 463(2088):3251–3269. <https://doi.org/10/crb9j8>
- Forrester AIJ, Sobester A, Keane AJ (2008) Engineering design via surrogate modelling: A practical guide. Wiley, New York
- Garcia JV (2008) Development of valid models for structural dynamic analysis. PhD thesis. Imperial College London, London
- Gürsoy HN, Patrikalakis NM (1991) Automated interrogation and adaptive subdivision of shape using medial axis transform. *Adv Eng Softw Work* 13(5):287–302. <https://doi.org/10/crqs2w>
- Jones DR, Schonlau M, Welch WJ (1998) Efficient global optimization of expensive black-box functions. *J Glob Optim* 13(4):455–492. <https://doi.org/10/fg68nc>
- Katz RA, Pizer SM (2003) Untangling the Blum medial axis transform. *Int J Comput Vis* 55(2):139–153. <https://doi.org/dm345d>
- Keane AJ, Nair P (2005) Computational approaches for aerospace design: the pursuit of excellence. Wiley
- Kennedy MC, O'Hagan A (2000) Predicting the output from a complex computer code when fast approximations are available. *Biometrika* 87(1):1–13. <https://doi.org/10/bps6bs>
- Kulkarni YH, Sahasrabudhe A, Kale M (2017) Leveraging feature generalization and decomposition to compute a well-connected midsurface. *Engineering with Computers* 33(1):159–170. <https://doi.org/10/f9mqf8>
- Macrini D, Dickinson S, Fleet D, Siddiqi K (2011) Bone graphs: Medial shape parsing and abstraction. *Comput Vis Image Underst* 115(7):1044–1061. <https://doi.org/10/bm6cdg>
- Morris MD, Mitchell TJ (1995) Exploratory designs for computational experiments. *Journal of Statistical Planning and Inference* 43(3):381–402. <https://doi.org/10/bbz8bx>
- Park C, Haftka RT, Kim NH (2017) Remarks on multi-fidelity surrogates. *Struct Multidiscip Optim* 55(3):1029–1050. <https://doi.org/10/f92h5f>
- Pastor M, Binda M, Harcarik T (2012) Modal assurance criterion. *Proc Eng* 48:543–548. <https://doi.org/10/gfr975>
- Sacks J, Welch WJ, Mitchell TJ, Wynn HP (1989) Design and analysis of computer experiments. *Stat Sci* 4(4):409–423. <https://doi.org/10/fn276g>
- Siemens (2014) NX nastran quick reference guide. Siemens
- Sóbester A, Forrester AIJ, Toal DJJ, Tresidder E, Tucker S (2014) Engineering design applications of surrogate-assisted optimization techniques. *Optim Eng* 15(1):243–265. <https://doi.org/gftpgg>
- Stanley F (2010) Dimensional reduction and design optimization of gas turbine engine casings for tip clearance studies. PhD thesis. University of Southampton, Southampton
- Tagliasacchi A, Delame T, Spagnuolo M, Amenta N, Telea A (2016) 3D skeletons: A state-of-the-art report. *Comput Graphics Forum* 35(2):573–597. <https://doi.org/10/f8q7s5>
- Tek H, Kimia BB (2001) Boundary smoothing via symmetry transforms. *J Math Imaging Vision* 14(3):211–223. <https://doi.org/10/cpczbh>
- Toal DJJ (2015) Some considerations regarding the use of multi-fidelity Kriging in the construction of surrogate models. *Struct Multidiscip Optim* 51(6):1223–1245. <https://doi.org/f7fg2v>
- Toal DJJ, Bressloff NW, Keane AJ (2008) Kriging hyperparameter tuning strategies. *AIAA J* 46(5):1240–1252. <https://doi.org/10/cm3jm4>
- Toal DJJ, Forrester AIJ, Bressloff NW, Keane AJ, Holden C (2009) An adjoint for likelihood maximization. *Proceedings of the Royal Society A: Mathematical. Phys Eng Sci* 465(2111):3267–3287. <https://doi.org/10/dgzkhx>
- Toal DJJ, Bressloff NW, Keane AJ, Holden CME (2011) The development of a hybridized particle swarm for kriging hyperparameter tuning. *Eng Optim* 43(6):675–699. <https://doi.org/10/fbsnfc>
- Toal DJJ, Keane AJ, Benito D, Dixon JA, Yang J, Price M, Robinson T, Remouchamps A, Kill N (2014) Multifidelity multidisciplinary whole-engine thermomechanical design optimization. *J Propuls Power* 30(6):1654–1666. <https://doi.org/10/gftpgj>
- Towner R, Band J (2012) An analysis technique/automated tool for comparing and tracking analysis modes of different finite element models. In: 53rd AIAA/ASME/ASCE/AHS/ASC structures, Structural Dynamics and Materials Conference American Institute of Aeronautics and Astronautics, Honolulu, Hawaii. <https://doi.org/10/gfr97s>
- Voutchkov I, Keane A, Fox R (2006) Robust structural design of a simplified jet engine model, using multiobjective optimization. In: 11th AIAA/ISSMO multidisciplinary analysis and optimization conference american institute of aeronautics and astronautics, Portsmouth, Virginia. <https://doi.org/10/gfr97t>
- Wang L, Toal DJJ, Keane AJ, Stanley F (2014) An accelerated medial object transformation for whole engine optimisation. In: Volume 2B: Turbomachinery, ASME, Düsseldorf, Germany, p V02BT45A011. <https://doi.org/10/gfr967>

- Wang L, Toal DJJ, Keane AJ, Stanley F (2017) A whole engine optimization based on medial object transformations. In: 27th ISABE Conference, Manchester
- Woo Y (2014) Abstraction of mid-surfaces from solid models of thin-walled parts: A divide-and-conquer approach. *Comput Aided Des* 47:1–11. <https://doi.org/10/gfr97f>
- Zhang X, Toal DJJ, Keane AJ, Witham F, Gregory J, Ravikanti M, Aurifeille E, Stow S, Rogers M, Zedda M (2015) Isothermal combustor prediffuser and fuel injector feed arm design optimization using the prometheus design system. *J Eng Gas Turbines Power* 138(6):061504–061504–17. <https://doi.org/10/gftpgn>
- Zhu H, Shao Y, Liu Y, Li C (2016) Mid-surface abstraction for complex thin-wall models based on virtual decomposition. *Int J Comput Integr Manuf* 29(8):821–838. <https://doi.org/10/gftpgp>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.