

RESEARCH ARTICLE

# Bayesian assessments of aeroengine performance with transfer learning

Pranay Seshadri<sup>1,2,\*</sup> , Andrew B. Duncan<sup>1,2</sup> , George Thorne<sup>3</sup>, Geoffrey Parks<sup>4</sup>,  
Raul Vazquez Diaz<sup>3</sup> and Mark Girolami<sup>2,4</sup> 

<sup>1</sup>Department of Mathematics, Imperial College London, 180 Queen's Gate, London SW7 2AZ, United Kingdom

<sup>2</sup>Data-Centric Engineering, The Alan Turing Institute, 96 Euston Road, London NW1 2DB, United Kingdom

<sup>3</sup>Civil Aerospace, Rolls-Royce plc, 100 Victory Road, Derby DE24 8EN, United Kingdom

<sup>4</sup>Department of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, United Kingdom

\*Corresponding author. E-mail: [p.seshadri@imperial.ac.uk](mailto:p.seshadri@imperial.ac.uk)

**Received:** 21 December 2021; **Revised:** 29 July 2022; **Accepted:** 17 August 2022

**Keywords:** Gaussian processes; jet-engines; Bayesian inference; measurements; thermofluids

## Abstract

Aeroengine performance is determined by temperature and pressure profiles along various axial stations within an engine. Given limited sensor measurements, we require a statistically principled approach for inferring these profiles. In this paper we detail a Bayesian methodology for interpolating the spatial temperature or pressure profile at axial stations within an aeroengine. The profile at any given axial station is represented as a spatial Gaussian random field on an annulus, with circumferential variations modelled using a Fourier basis and radial variations modelled with a squared exponential kernel. This Gaussian random field is extended to ingest data from multiple axial measurement planes, with the aim of transferring information across the planes. To facilitate this type of *transfer learning*, a novel planar covariance kernel is proposed. In the scenario where frequencies comprising the temperature field are unknown, we utilise a sparsity-promoting prior on the frequencies to encourage sparse representations. This easily extends to cases with multiple engine planes whilst accommodating frequency variations between the planes. The main quantity of interest, the spatial area average is readily obtained in closed form. We term this the Bayesian area average and demonstrate how this metric offers far more representative averages than a sector area average—a widely used area averaging approach. Furthermore, the Bayesian area average naturally decomposes the posterior uncertainty into terms characterising insufficient sampling and sensor measurement error respectively. This too provides a significant improvement over prior standard deviation based uncertainty breakdowns.

## Impact Statement

This article offers systematic improvements for the way sensor data is synthesized in aeroengines, with applications to gas turbines, compressors, steam turbines, and other forms of turbomachinery. Specific contributions of this article include:

- a more rigorous approach for calculating area averages in turbomachinery;
- metrics for spatial sampling and measurement imprecision uncertainty using the law of total covariance—that is, are more sensors required or simply better sensor quality;

- utilization of sparsity promoting priors for estimating circumferential distributions of aerothermal quantities, and
- a framework to transfer information across numerous measurement planes.

## 1. Introduction

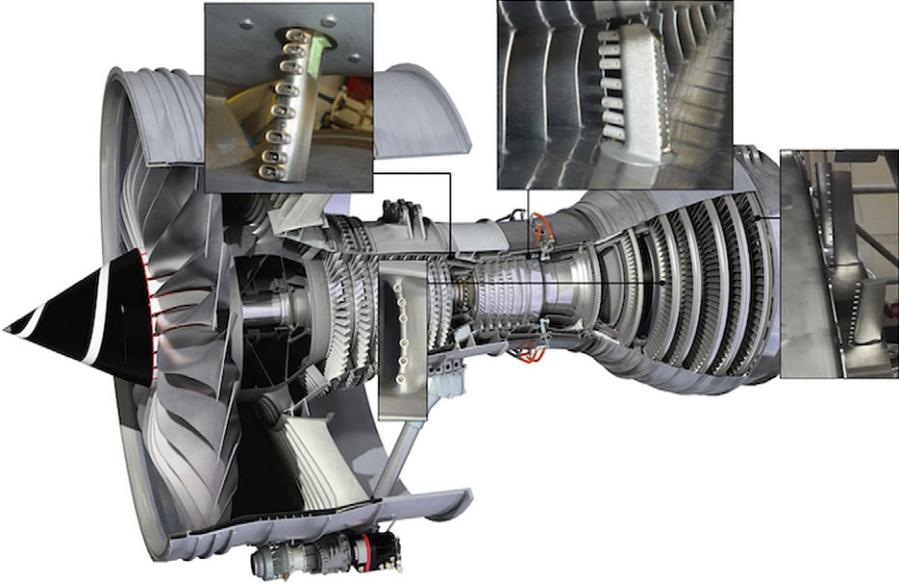
Temperature and pressure measurements are vital in both the prognostics of existing in-flight engines and the understanding of new engine architectures and component designs. There are two reasons for this. First, over many running cycles, an engine will undergo a certain level of degradation. This typically manifests as an increase in blade tip and seal clearances (Seshadri et al., 2014); an accumulation of dirt and other contaminants within the gas path, and blade surface damage owing to oxidation, sulfidation, and the impact of foreign objects (Aust and Pons, 2019). These factors increase the amount of work the compressor has to do to achieve a certain pressure rise and the amount of work the turbine has to do to deliver the power required. The consequence of this increased workload is higher temperatures in both the compressor and turbine sections, measured via temperature probes; the pressure rise is measured via pressure probes. One of these measurements, the engine gas temperature (EGT)<sup>1</sup> forms an important metric for forecasting the remaining useful life of an engine (Marinai, 2004; Bonnet, 2007). It can be found on the engine performance panel in aircraft cockpit displays, as shown in Figure 1. Note that this 1D value, among others, is a reported average across a 2D nonuniform spatial field.

The second reason why temperature and pressure measurements are so critical is because they are used to compute subsystem (e.g., low-, intermediate-, and high-pressure compressor, and turbine components) efficiencies. This is done by defining a control volume around the subsystem of interest and ascertaining what the average stagnation flow properties are at the inlet and exit, whilst accounting for work being done both into and out of the system. At each measurement plane, circumferentially positioned rakes—with radially varied probes on each rake—are used to measure pressure and temperature values (see Figures 2 and 3). These measurements are aggregated through 1D area- or mass-averages of the circumferentially

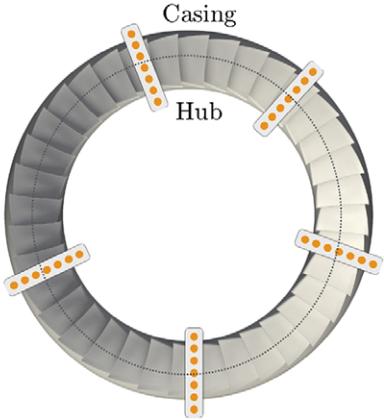


**Figure 1.** Cockpit display of a twin-engine aircraft with a close-up (inset) of the engine performance parameters. The engine gas temperature (EGT) for both engines is shown within the blue boxes. Source: Flightradar24 (2021). Image reproduced with permission from FlightRadAr24 under a Creative Commons Attribution 4.0 license.

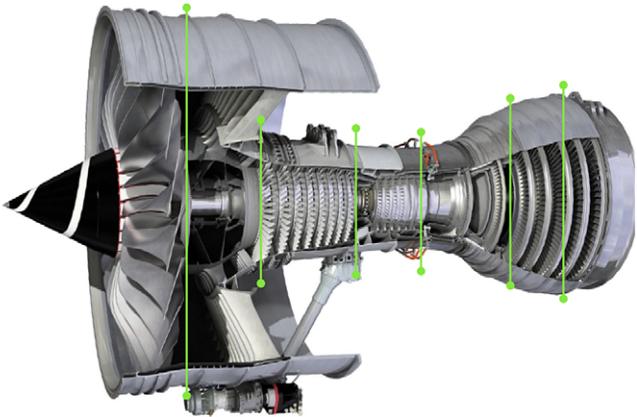
<sup>1</sup> Also known as the turbine gas temperature (TGT).



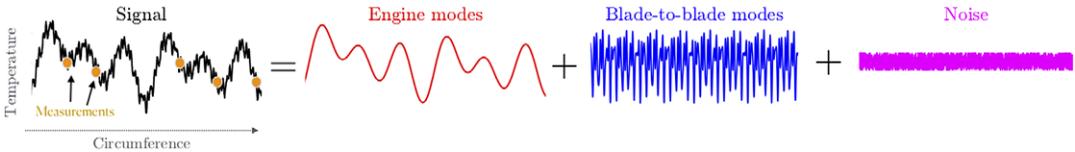
**Figure 2.** Characteristic temperature and pressure rakes at a few locations in an aeroengine. Source: Rolls-Royce plc.



Close-up of an isolated measurement plane with rakes



Engine axial view with locations of measurement planes



**Figure 3.** Close-up of an axial measurement plane in an engine. Each plane is fitted with circumferentially scattered rakes with radially placed probes. The circumferential variation in temperature (or pressure) can be broken down into various modes, as shown. Engine cutaway image source: Rolls-Royce plc.

and radially scattered measurements at a given axial plane. Identifying which subsystem needs to be improved based on its efficiency rating, feeds into research and development programmes for current and new engines. Furthermore, if the uncertainty in a given subsystem's calculated efficiency is deemed too large, then it is likely that a decision on adding more instrumentation or improving the precision of the existing sensors will follow. As both (a) research and development programmes for improving the performance of a given subsystem, and (b) the enhancement of the engine measurement apparatus, are extremely expensive, it is imperative that the decisions made be based upon accurate and precise temperature and pressure values.

### **1.1. 1D performance values**

As in many other engineering disciplines, 1D metrics are often used for performance assessments in turbomachinery. When provided with radial and circumferentially placed temperature or pressure measurements, area-based averages are often the norm for arriving at 1D values. These are typically estimated by assigning each sensor a weight based on the sector area it covers. This weight will depend on the total number of sensors and the radial and circumferential spacing between them (Stoll et al., 1979; Francis and Morse, 1989). This *sector area-average* is computed by taking the weighted sum of each measurement and dividing it by the sum of the weights themselves. In practice, this recipe offers accurate estimates if the spatial distribution of the measured quantity is uniform throughout the measurement plane. For spatially nonuniform flows, the validity of this approach hinges on the circumferential placement of the rakes and the harmonic content of the signal. Should all the rakes be placed so as to capture the trough of the waveforms, then the sector area-average will likely underestimate the *true* area-average. A similar argument holds if the rakes are placed so as to capture only the peaks of the circumferential pattern (Seshadri et al., 2020b). It is therefore common to use empirical corrections to account for the uncertainty in such measurements, however, these corrections may introduce biases. It should be noted that in-flight engines may only be fitted with one or two rakes, which may warrant additional corrections. This is different from test-bed (simulated altitude) which often has more rakes along the same axial plane. Additionally, test-bed engines may also have more axial measurement stations.

### **1.2. Limitations with computational fluid dynamics**

A salient point to note here concerns the use and limitations (see Denton, 2010) of a strictly computational approach to estimate engine pressures and temperatures. Today, aeroengine computational fluid dynamics (CFD) flow-field approximations via Reynolds averaged Navier Stokes (RANS), large eddy simulations (LES) (Gourdain et al., 2014), and in some cases, via direct numerical simulations (DNS) (Wheeler et al., 2016) are being increasingly adopted to gain insight into both component- and subsystem-level design. These CFD solvers with varying fidelities of underpinning equations and corresponding domain discretizations have found success—balancing simulation accuracy with simulation cost—in understanding the flow-physics in the numerous subsystems of an aeroengine (see figure 11 in Tyacke et al., 2019). However, in most cases, CFD-experimental validation is carried out using scaled experimental rigs which typically isolate one subsystem or a few stages (rows of rotors and stators) in an engine. Although there has been a tremendous body of work dedicated to incorporating real-engine effects through aleatory (Seshadri et al., 2014, 2015; Montomoli, 2015) and epistemic uncertainty quantification (Emory et al., 2016) studies, as a community, we are still far from being able to replicate the aerothermal environment in engines: it is incredibly complex. For instance, the hub and casing are never perfectly circular owing to variability in thermal and fatigue loads; engine structural components introduce asymmetries into the flow that can propagate far downstream into the machine, leading to flow-field distortions; and the pressure and temperature variations induced by bleeds and leakage flows are not circumferentially uniform. The presence of these *engine modes* (also termed engine wave numbers) makes it challenging to use CFD in isolation to calculate aeroengine performance.

### 1.3. Coverage versus accuracy

Before we delve into the main ideas that underpin this article, it will be helpful to understand the experimental coverage versus accuracy trade-off. Sensor placement in an engine is tedious: there are stringent space constraints on the number of sensors, the dimensions of each sensor, and its ancillary equipment, along with its axial, radial, and tangential location in the engine. However, engine measurements offer the most accurate representation of engine flow physics. Scaled rigs, on the other hand, offer far greater flexibility in sensor number, type, and placement, and consequently yield greater *measurement coverage*. While they are unable to capture the engine modes—and thus are limited in their ability to emulate the engine—they offer an incredibly rich repository of information on the *blade-to-blade* modes. These modes include those associated with periodic viscous mixing (such as from blade tip vortices), overturning boundary layers between two adjacent blades, and the periodic inviscid wakes (Sanders et al., 2002; Mailach et al., 2008). Although present in the engine environment too, engines have insufficient measurement coverage to capture these blade-to-blade modes. One can think of the spatial distribution of pressure or temperature as being a superposition of such blade-to-blade modes (visible in rig experiments), engine modes (visible in engine tests), and noise (see Figure 3). Succinctly stated, our best window on flow into an aeroengine—and, in consequence, its composite temperatures and pressures—stems from real engine measurements themselves. The challenge is that they are few and far between.

In summary, to compute important engine performance metrics such as component efficiencies, pressure ratios, EGT, and thrust—all of which are 1D performance metrics—we need to spatially average sparse pressure and temperature measurements that arise from engine tests. However, this averaging needs to be done in a rigorous manner to account for possible nonuniformity in the flow, limited spatial measurements, and uncertainty in each measurement itself.

### 1.4. State of the art

While publicly available work in the areas of measurement aeroengine metrology (Saravanmuttoo, 1990; SAE International, 2017), averaging (Greitzer et al., 2004; Cumpsty and Horlock, 2006), and spatial field approximation (Seshadri et al., 2020a,b) are prevalent, there is no unifying framework for these related concerns. In other words, there is no established workflow that stems from measurements to spatial field approximation to averaging, whilst rigorously accounting for all the sources of uncertainties. There are isolated estimates of uncertainties tailored for specific cases. For instance, Bonham et al. (2017) state that, for compressors, at least seven measurements are required in the radial direction, and at least five measurements in the circumferential direction to resolve the flow. This is a heuristic, based on the negligible change in isentropic efficiency if more measurements are taken. It should be noted that this assessment is not based on a spatial model, but rather on experimental observations for a compressor with an inlet stagnation temperature of 300 K and a polytropic efficiency of 85% at three different pressure ratios. It is difficult to generalize this across all compressors.

In Seshadri et al. (2020a,b), the authors draw our attention to the lack of a comprehensive averaging and uncertainty assessment strategy in literature—especially for spatially nonuniform flows. They articulate the limitations of widely adopted existing experimental measurement, data processing, and uncertainty guides, such as the International Organization for Standardization (ISO) guide of uncertainty in measurements, the SAE (formerly the Society of Automotive Engineers) Aerospace Information Report AIR1419C (SAE International, 2017), and the American Society of Mechanical Engineers (ASME) performance test code (PTC) 19.1 (Dieck et al., 2005). A key point the authors argue is that computing 1D metrics from experimental data should be directly based on some spatial representation of the experimental data. To this end, they present a regularized linear least squares strategy for estimating the spatial flow field from a grid of measurements formed by radial and circumferentially placed probes. Their data-driven model represents the spatial flow-field in the circumferential direction via a Fourier series expansion, while capturing flow in the radial direction using a high-degree polynomial. Although an improvement in the state of the art (Lou and Key, 2021), their model does have limitations. For instance, the placement of probes may lead to Runge's phenomenon (see Chapter 13 in Trefethen, 2013) in

the radial direction, while the harmonic content is set by the Nyquist condition (see Chapter 4 in Strang, 2012) in the circumferential direction. Another hindrance, one not systemic to their work, but one mentioned in several texts (see 8.1.4.4.3 in Saravanmuttoo, 1990 and in Pianko and Wazelt, 1983), is the definition of the uncertainty associated with insufficient spatial sampling and that associated with the imprecision of each sensor. This decomposition of the overall uncertainty is important as it informs aeroengine manufacturers whether they need more measurement sensors or whether they need to improve the precision of existing sensor systems. At present, there are no rigorously derived metrics for this.

The “so what?” and overarching motivation for this article are succinctly summarized by the following two facts. First, engine manufacturers spend millions of dollars toward incremental gains in efficiency. This is because even a 1% increase in the efficiency of a subsystem can have a sizeable reduction in fuel consumption. Second, existing methods for averaging and delivering uncertainty assessments are provably inadequate and likely too conservative. For instance, for a modern turbine using existing measurement practices leads to a 1.5–2.5 K uncertainty in stagnation temperature measurements. This results in a 1% uncertainty in efficiency (see 2.2 in Seshadri et al., 2020a). More accurate approaches for averaging and uncertainty quantification for temperature (and pressure) measurements will lead to more accurate efficiency estimates.

### 1.5. Paper outline

In this article, we argue that an assessment of the area average and a decomposition of the overall uncertainty is only possible with a priori knowledge of the spatial flow field. Thus, we frame part of our scope as follows. Given an array of engine sensor measurements at single or multiple axial stations, our goal is to formulate techniques to:

- construct a spatial model to approximate the flow-field at an axial station given the inherent uncertainty in the measurements and certain physical assumptions (see Section 2.2);
- compute the area-average of the stagnation pressure and temperature based on this model (Section 4.1);
- distinguish between uncertainty in the spatial model (and its averages) induced by sensor imprecision, and insufficient spatial sampling (Section 5);
- quantify the dominant circumferential harmonics leveraging some notion of sparsity (Section 3.2);
- develop methodologies that can transfer information from relatively more heavily instrumented test-bed engines to very sparsely instrumented flight engines at the same plane; and
- foster the transfer of information between adjacent planes in an engine with the intention of reducing uncertainty (see Section 2.2).

These latter two aims will be addressed using *transfer learning*—an emerging subdiscipline of machine learning that seeks to transfer information between tasks, intelligently (Skolidis, 2012) especially when one task is afforded more information than the other. In this article, we explore the topics of spatial field estimation, area averaging, instrumentation sampling versus precision uncertainty estimation, and transfer learning with Gaussian processes (GPs; Rasmussen and Williams, 2006).

## 2. GP Aeroengine Model

GPs provide a powerful framework for nonparametric regression, where the regression function is modeled as a random process, such that the distribution of the function evaluated at any finite set of points is jointly Gaussian. GPs are characterized by a mean function and a two-point covariance function. GPs have been widely used to model spatial and temporal varying data since their first application in modeling ore reserves in mining (Krige, 1951), leading to a method for spatial interpolation known as kriging in the geostatistics community (Stein, 2012; Cressie, 2015). The seminal work of Kennedy and O’Hagan (2001) provides a mature Bayesian formulation which forms the underpinnings of the approach

adopted in this article. Emulation methods based on GPs are now widespread and find uses in numerous applications ranging from computer code calibration (Higdon et al., 2004) and uncertainty analysis (Oakley and O'Hagan, 2002) to sensitivity analysis (Oakley and O'Hagan, 2004). Since then GP regression has enjoyed a rich modern history within uncertainty quantification, with increasingly sophisticated extensions beyond the classical formulation, including latent space models (Chen et al., 2015), coregional models (Alvarez et al., 2012) convolutional processes (Higdon, 2002; Álvarez and Lawrence, 2011), multi-task processes (Bonilla et al., 2008), and GPs with incorporated dimension reduction (Liu and Guillas, 2017; Seshadri et al., 2019).

In a multi-task GP, one is given similar but distinct multiple input–output data sets—each referred to as a task. Rather than train a single model for each task (single-task), the idea is to train a single model for all the tasks simultaneously. The advantage is that by constructing the latter, information can be readily shared across tasks in a meaningful manner, thereby aiding in improved inference. This implies, either implicitly or explicitly, that there are features of the model that are either hierarchical or that define common structure across the different tasks. As Skolidis (2012) remarks, multi-task GPs can be advantageous when compared to single-task GPs when there is insufficient data to infer all the model's parameters. It is expected that multi-task GPs would exploit the common structure prevalent across all tasks for improved parameter inference. Practically, one approach is to express the covariance function as the Kronecker product of a task-based covariance function and a data-based covariance function (see 54 in Skolidis, 2012). While a Kronecker product-based definition of the multi-task kernel does have computational advantages, it restricts one to using the same set of radial and circumferential measurements at each measurement plane.

We end this brief literature survey with a remark on subtlety between multi-task models and models with transfer learning. All transfer learning models are inherently multi-task, however, not all multi-task models are transfer learning models. The key distinction lies in whether any information is actually transferred across the tasks, and whether that transfer leads to a more well-defined model.

### 2.1. Preliminaries and data

In this subsection, we present a GP aeroengine spatial model—designed to emulate the steady-state temperature and pressure distributions at multiple axial planes. Given the complexity of the flow, our aim is to capture the *primary aerothermal features* rather than resolve the flow field to minute detail. One can think of the primary aerothermal features as being the engine modes in the circumferential direction. In what follows we detail our GP regression model; our notation closely follows the GP exposition of Rogers and Girolami (see chapter 8 in Rogers and Girolami, 2016).

Let us assume that we have sensor measurement location and sensor reading pairs  $(\mathbf{x}_i, f_i)$  for  $i = 1, \dots, N$ , and  $M$  locations at which we would like to make reading predictions

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_1 \\ \vdots \\ \mathbf{x}_N \end{bmatrix} \quad \mathbf{f} = \begin{bmatrix} f_1 \\ \vdots \\ f_N \end{bmatrix} \quad \text{and} \quad \mathbf{X}^* = \begin{bmatrix} \mathbf{x}_1^* \\ \vdots \\ \mathbf{x}_M^* \end{bmatrix} \quad \mathbf{f}^* = \begin{bmatrix} f_1^* \\ \vdots \\ f_M^* \end{bmatrix}, \quad (1)$$

where the superscript  $(*)$  denotes the latter. Here  $\mathbf{x}_i \in \mathbb{R}^3$ , thus  $\mathbf{X} \in \mathbb{R}^{N \times 3}$ . Without loss in generality, we assume that  $\sum_i^N f_i = 0$ , so that the components correspond to deviations around the mean; physically, being either temperature or pressure measurements taken at the locations in  $\mathbf{X}$ . Let the values in  $\mathbf{f}$  be characterized by a symmetric *measurement covariance* matrix  $\Sigma \in \mathbb{R}^{N \times N}$  with diagonal measurement variance terms  $\sigma_m^2$ , that is,  $\Sigma = \sigma_m^2 \mathbf{I}$ . In practice,  $\Sigma$ , or at least an upper bound on  $\Sigma$ , can be determined from the instrumentation device used and the correlations between measurement uncertainties, which will be set by an array of factors such as the instrumentation wiring, batch calibration procedure, data acquisition system, and filtering methodologies. Thus, the true measurements  $\mathbf{t} \in \mathbb{R}^N$  are corrupted by a zero-mean Gaussian noise,  $\mathbf{f} = \mathbf{t} + \mathcal{N}(\mathbf{0}, \Sigma)$  yielding the observed sensor values. This noise model, or likelihood, induces a probability distribution  $\mathbb{P}(\mathbf{f} | \mathbf{t}, \mathbf{X}) = \mathcal{N}(\mathbf{f}, \Sigma)$  around the true measurements  $\mathbf{t}$ .

In the absence of measurements, we assume that  $\mathbf{f}$  is a Gaussian random field with a mean of  $\mathbf{0}$  and has a two-point covariance function  $k(\cdot, \cdot)$ . The joint distribution of  $(\mathbf{f}, \mathbf{f}^*)$  satisfies

$$\begin{bmatrix} \mathbf{f} \\ \mathbf{f}^* \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{K}_{\bullet\bullet} + \Sigma & \mathbf{K}_{\bullet*} \\ \mathbf{K}_{\bullet*}^T & \mathbf{K}_{**} \end{bmatrix}\right), \tag{2}$$

where the Gram matrices are given by  $[\mathbf{K}_{\bullet\bullet}]_{(i,j)} = k(\mathbf{x}_i, \mathbf{x}_j)$ ,  $[\mathbf{K}_{\bullet*}]_{(i,l)} = k(\mathbf{x}_i, \mathbf{x}_l^*)$ , and  $[\mathbf{K}_{**}]_{(l,m)} = k(\mathbf{x}_l^*, \mathbf{x}_m^*)$ , for  $i, j = 1, \dots, N$  and  $l, m = 1, \dots, M$ . From (2), we can write the predictive posterior distribution of  $\mathbf{f}^*$  given  $\mathbf{f}$  as  $\mathbb{P}(\mathbf{f}^* | \mathbf{f}, \mathbf{X}^*, \mathbf{X}) = \mathcal{N}(\boldsymbol{\mu}^*, \boldsymbol{\Psi}^*)$  where the conditional mean is given by

$$\begin{aligned} \boldsymbol{\mu}^* &= \mathbf{K}_{\bullet*}^T (\mathbf{K}_{\bullet\bullet} + \Sigma)^{-1} \mathbf{f} \\ &= \mathbf{K}_{\bullet*}^T \mathbf{G}^{-1} \mathbf{f}, \end{aligned} \tag{3}$$

with  $\mathbf{G} = (\mathbf{K}_{\bullet\bullet} + \Sigma)$ ; the conditional covariance is

$$\boldsymbol{\Psi}^* = \mathbf{K}_{**} - \mathbf{K}_{\bullet*}^T \mathbf{G}^{-1} \mathbf{K}_{\bullet*}. \tag{4}$$

**2.2. Defining the covariance kernels**

As our interest lies in applying GP regression over  $P$  engine planes, our inputs  $\mathbf{x}_i \in \{(r_i, \theta_i, \rho_i) : r_i \in [0, 1], \theta_i \in [0, 2\pi), \rho_i \in \{1, \dots, P\}\}$  can be parameterized as

$$\mathbf{X} = \begin{bmatrix} r_1 & \theta_1 & \rho_1 \\ \vdots & \vdots & \vdots \\ r_N & \theta_N & \rho_N \end{bmatrix} = [\mathbf{r} \ \boldsymbol{\theta} \ \boldsymbol{\rho}], \text{ and } \mathbf{X}^* = \begin{bmatrix} r_1 & \theta_1 & \rho_1 \\ \vdots & \vdots & \vdots \\ r_M & \theta_M & \rho_M \end{bmatrix} = [\mathbf{r}^* \ \boldsymbol{\theta}^* \ \boldsymbol{\rho}^*]. \tag{5}$$

In most situations under consideration, we expect that  $\mathbf{X} = \{(r_i, \theta_j, \rho_l), r_i \in \mathbf{r}, \theta_j \in \boldsymbol{\theta}, \rho_l \in \boldsymbol{\rho}\}$  where  $\mathbf{r}$  is a set of  $L$  radial locations,  $\boldsymbol{\theta}$  is a set of  $O$  circumferential locations and  $P$  is the number of measurement planes, such that  $N = L \times O \times P$ , assuming the measurements across the  $P$  planes are taken at the same locations. We define the spatial kernel to be a product of a Fourier kernel  $k_c$  in the circumferential direction, a squared exponential kernel  $k_r$  in the radial direction, and a planar kernel  $k_p$  along the discrete  $P$  different planes

$$\begin{aligned} k(\mathbf{x}, \mathbf{x}') &= k((\mathbf{r}, \boldsymbol{\theta}, \boldsymbol{\rho}), (\mathbf{r}', \boldsymbol{\theta}', \boldsymbol{\rho}')) \\ &= k_r(\mathbf{r}, \mathbf{r}') \odot k_c(\boldsymbol{\theta}, \boldsymbol{\theta}') \odot k_p(\boldsymbol{\rho}, \boldsymbol{\rho}'), \end{aligned} \tag{6}$$

where the symbol  $\odot$  indicates a Hadamard (element-wise) product.<sup>2</sup>

Along the radial direction, the kernel has the form

$$k_s(\mathbf{r}, \mathbf{r}') = \sigma_f^2 \exp\left(-\frac{1}{2l^2}(\mathbf{r} - \mathbf{r}')^T(\mathbf{r} - \mathbf{r}')\right), \tag{7}$$

where  $\sigma_f$  is the signal variance and  $l$  is the length-scale—two hyperparameters that need to be computationally ascertained.

Our primary mechanism for facilitating transfer learning is via the planar kernel. For this kernel, we define  $\mathbf{s} \in \mathbb{Z}_+^P$  to be a *similarity* vector of length  $P$  comprised of strictly positive integers. Repetitions in  $\mathbf{s}$  are permitted and are used to indicate which planes are similar. For instance, if we set  $\mathbf{s} = (1, 1, 2)$ , this indicates that the first two planes are similar. We will use the notation  $\mathbf{s}(\rho)$  to select the similarity value corresponding to a specific plane  $\rho$ . The number of unique integers in  $\mathbf{s}$  may be thought of as the number of independent planes; let this be given by  $Q$ , implying  $Q \leq P$ .

Seeing as there are  $Q$  independent planes, we require a metric that serves to *correlate* the different independent plane combinations. To this end, consider a symmetric matrix  $\mathbf{S} \in \mathbb{R}^{Q \times Q}$  with a diagonal

<sup>2</sup> For computational efficiency, the Kronecker product can also be used in cases where there are no missing entries, that is, sensor values can be obtained from a grid of measurements.

formed of  $\eta$  values. The  $\eta$  values denote the correlation between planes that are similar, and by construction it is a tunable hyperparameter. In practice, unless the planes are identical, their correlation will be less than unity, that is,  $\eta < 1$ . Next, set  $W = Q(Q - 1)/2$ , corresponding to the number of upper (or lower) triangular off-diagonal elements in a  $Q \times Q$  matrix. As each off-diagonal entry represents a pairwise correlation between two independent planes, it needs to be represented via another appropriate hyperparameter. Let  $\xi = (\xi_1, \dots, \xi_W)$  be this hyperparameter, yielding

$$S = \begin{bmatrix} \eta & \xi_1 & \dots & \xi_{Q-1} \\ \xi_1 & \ddots & \dots & \vdots \\ \vdots & \dots & \ddots & \xi_W \\ \xi_{Q-1} & \dots & \xi_W & \eta \end{bmatrix}. \tag{8}$$

Then, the planar kernel us given by

$$k_p(\rho_i, \rho'_j) = \begin{cases} 1 & \text{if } \rho_i = \rho'_j \\ [S]_{s(\rho_i), s(\rho'_j)} & \text{otherwise.} \end{cases} \tag{9}$$

In summary, the planar kernel establishes the correlation between all the  $P$  measurement planes. It is invariant to the radial and circumferential values and is only dependent upon the planes chosen. We remark here that the type of transfer learning facilitated by (9) is inherently inductive as we are not reusing parameters from a prior regression task—typically seen across many deep learning approaches.

Prior to defining the kernel along the circumferential direction, a few additional definitions are necessary. Let  $\omega = (\omega_1, \dots, \omega_K)$  indicate the  $K$  wave numbers present along the circumferential direction for a given plane. These can be a specific set, that is,  $\omega = (1, 4, 6, 10)$ , or can be all wave numbers up to a particular cut-off, that is,  $\omega = (1, 2, \dots, 25)$ . We define a *Fourier design matrix*  $F \in \mathbb{R}^{(2K+1) \times N}$ , the entries of which are given by

$$F_{ij}(\theta) = \begin{cases} 1 & \text{if } i = 1 \\ \sin(\omega_i \pi \theta_j / 180^\circ) & \text{if } i > 1 \text{ when } i \text{ is even} \\ \cos(\omega_{i-1} \pi \theta_j / 180^\circ) & \text{if } i > 1 \text{ when } i \text{ is odd} \end{cases}. \tag{10}$$

Note that the number of columns in  $F$  depends on the size of the inputs  $\theta$ . To *partially* control the amplitude and phase of the Fourier modes and the value of the mean term, we introduce a set of diagonal matrices  $\mathcal{D} = (\mathbf{D}_1, \dots, \mathbf{D}_Q)$ . Each matrix has dimension  $\mathbb{R}^{(2K+1) \times (2K+1)}$ , with entries  $\mathbf{D}_i = \text{diag}(\lambda_{i,1}^2, \dots, \lambda_{i,2K+1}^2)$  for  $i = 1, \dots, Q$ . Note, we use the word *partially*, as these hyperparameters are not indicative of the amplitude or phase directly, as they depend on the measured data too. The hyperparameters themselves are variances, denoted using the squared terms  $\lambda_{i,j}^2$  along the diagonal in  $\mathbf{D}$ . Furthermore, note that the matrices in  $\mathbf{D}_i$ , and thus number of tunable hyperparameters, scale as a function of the number of independent planes  $Q$  and not by the total number of planes  $P$ . The kernel in the circumferential direction may then be written as

$$k_c((\theta, \rho_i), (\theta', \rho'_j)) = \mathbf{F}(\theta)^T \sqrt{\mathbf{D}_{s(\rho_i)}} \sqrt{\mathbf{D}_{s(\rho'_j)}} \mathbf{F}(\theta'), \tag{11}$$

where the notation  $\mathbf{D}_{s(\rho_i)}$  corresponds to the diagonal matrix index by  $s(\rho_i)$ . We remark here that as written in (11) the Fourier modes across all the  $P$  planes are fixed, though the amplitudes and phases can vary.

Having established the definition of the radial, planar and circumferential kernels, it is worthwhile to take stock of our aim. We wish to represent the primary aerothermal attributes using radial, circumferential and planar kernels. While the focus of this article is on engine and rig test data, a few comments regarding transfer learning with engine measurements and high-fidelity CFD is in order. Should temperature or pressure values across the annulus—or a part thereof—be available from RANS, mean unsteady RANS, or even time-averaged LES, we can still use the radial and

circumferential kernels on that data. Alternatively, as the spatial resolution of the CFD data will be far greater than the experimental one, a standard Matern kernel function along both the radial and circumferential directions  $k_{\text{CFD}}(\mathbf{x}, \mathbf{x}')$  may also be used. For a single plane, we can then define additive kernels of the form  $k(\mathbf{x}, \mathbf{x}') = k_c(\mathbf{r}, \mathbf{r}') \odot k_r(\mathbf{r}, \mathbf{r}') + k_{\text{CFD}}(\mathbf{x}, \mathbf{x}')$ . This idea can in practice capture the superposition is shown in Figure 3, where the CFD is used solely to resolve the higher frequency blade-to-blade modes. In terms of extending the current framework to temporally (or unsteady) problems, we note that this will require the development of a temporal kernel  $k_{\text{time}}(t, t')$  defined over the times  $t$ .

### 3. Priors

Using the GP regression framework implies that our model prior is Gaussian  $\mathbb{P}(\mathbf{t}|\mathbf{X}) = \mathcal{N}(\mathbf{0}, \mathbf{K}_{\bullet\bullet})$ . We have already established that our likelihood function is also Gaussian. The central objective of our effort is to determine the posterior  $\mathbb{P}(\mathbf{f}|\mathbf{X}, \sigma_f^2, l^2, \lambda_{i,1}^2, \dots, \lambda_{i,2K+1}^2, \eta, \xi_1, \dots, \xi_W)$ .

In this section, we impose priors on the hyperparameters in (6). Priors for the squared exponential kernel are given by

$$\begin{aligned} \sigma_f &\sim \mathcal{N}^+(0, 1), \\ l &\sim \mathcal{N}^+(0, 1), \end{aligned} \tag{12}$$

where  $\mathcal{N}^+$  represents a half-Gaussian distribution. For the planar kernel, in this article we set  $\eta \sim \mathcal{U}[0.8, 1.0]$ , where  $\mathcal{U}$  represents a uniform distribution. This range is chosen to foster a strong positive correlation between the two planes deemed similar. Note that this prior structure can be augmented to have each diagonal term in (8) have its own prior—an idea we do not pursue in this article. Finally, we assign

$$\xi_i \sim \mathcal{U}[-1, 1], \tag{13}$$

for  $i = 1, \dots, W$ . Priors for the Fourier kernel are detailed below.

#### 3.1. Simple prior

There are likely to be instances where the harmonics  $\omega$  are known, although this is typically the exception and not the norm. In such cases, the Fourier priors for a given plane index  $i$  may be given by  $\lambda_{i,j}^2 \sim \mathcal{N}^+(0, 1)$ , for  $j = 1, \dots, 2K + 1$ .

#### 3.2. Sparsity promoting prior

In the absence of further physical knowledge, we constrain the posterior by invoking an assumption of sparsity, that is, the spatial measurements can be adequately explained by a small subset of the possible harmonics. This is motivated by the expectation that a sparse number of Fourier modes contribute to the spatial pattern in the circumferential direction. In adopting this assumption, we expect to reduce the variance at the cost of a possible misfit. Here, we engage the use of sparsity promoting priors, which mimic the shrinkage behavior of the least absolute shrinkage and selection operator (LASSO) (Tibshirani, 1996; Bühlmann and Van De Geer, 2011) in the fully Bayesian context.

A well-known shrinkage prior for regression models is the spike-and-slab prior (Ishwaran and Rao, 2005), which involves discrete binary variables indicating whether or not a particular frequency is employed in the regression. While this choice of prior would result in a truly sparse regression model, where Fourier modes are selected or deselected discretely, sampling methods for such models tend to demonstrate extremely poor mixing. This motivates the use of continuous shrinkage priors, such as the horseshoe (Carvalho et al., 2009) and regularized horseshoe (Pironen and Vehtari, 2017) prior. In both of these a global scale parameter  $\tau$  is introduced for promoting sparsity; large values of  $\tau$  will lead to diffuse priors and permit a small amount of shrinkage, while small values of  $\tau$  will shrink all of the hyperparameters toward 0. The regularized horseshoe is given by

$$\begin{aligned}
c &\sim \mathcal{IG}\left(\frac{\gamma}{2}, \frac{\gamma s^2}{2}\right), \\
\tilde{\lambda}_{i,j} &\sim \mathcal{C}^+(0, 1), \\
\lambda_{i,j}^2 &= \frac{c \tilde{\lambda}_{i,j}^2}{c + \tau^2 \tilde{\lambda}_{i,j}^2}, \text{ for } j = 1, \dots, 2K + 1, \text{ and } i = 1, \dots, Q,
\end{aligned} \tag{14}$$

where  $\mathcal{C}^+$  denotes a half-Cauchy distribution;  $\mathcal{IG}$  denotes an inverse gamma distribution, and where the constants

$$\tau = \frac{\beta \sigma_m}{(1 - \beta) \sqrt{N}}, \gamma = 30 \text{ and } s = 1.0. \tag{15}$$

Hyperparameters  $\tilde{\lambda}_{i,j}$  are assigned half-Cauchy distributions that have thick tails so they may allow a fraction of the Fourier  $\lambda_{i,j}$  hyperparameters to avoid the shrinkage, while the remainder are assigned very small values. These hyperparameters indirectly control the amplitude and phase of the Fourier series representation, as mentioned before.

The scale parameter  $c$  is set to have an inverse gamma distribution—characterized by a light left tail and a heavy right tail—designed to prevent probability mass from aggregating close to 0 (Piironen and Vehtari, 2017). This parameter is used when a priori information on the scale of the hyperparameters is not known; it addresses a known limitation in the horseshoe prior where hyperparameters whose values exceed  $\tau$  would not be regularized. Through its relationship with  $\tilde{\lambda}_{i,j}$ , it offers a numerical way to avoid shrinking the standard deviation of the Fourier modes that are far from 0. Constants  $\gamma$  and  $s$  are used to adjust the mean and the variance of the inverse gamma scale parameter  $c$ , while constant  $\beta$  controls the extent of sparsity; large values of  $\beta$  imply that more harmonics will participate in the Fourier expansion, while smaller values of  $\beta$  would offer a more parsimonious representation.

There is one additional remark regarding the hierarchical nature of the priors above. If two measurement planes are similar as classified by  $\mathbf{s}$ , then they have the same set of Fourier hyperparameters. Note that having the same Fourier hyperparameters does not imply that the planes have the same circumferential amplitudes and phases. While we assume that all planes share the base harmonics  $\omega$ , the model above has sufficient flexibility to have multiple planes with distinct dominant harmonics—provided the overall extent of sparsity remains approximately similar. This property is useful in an aeroengine as dominant harmonics upstream may not be dominant downstream, owing to changes in vane counts, flow diffusion, the introduction of cooling flows, struts, and bleeds, among other componentry.

## 4. Posterior Inference

We generate approximate samples from the posterior distribution jointly on  $\mathbf{f}^*$  and the hyperparameters using Hamiltonian Monte Carlo (HMC) (Duane et al., 1987; Horowitz, 1991). In this work, we specifically use the No-U-Turn (NUTS) sampler of Hoffman and Gelman (2014), which is a widely adopted extension of HMC. The main advantage of this approach is that it mitigates the sensitivity of sampler performance on the HMC step size and the number of leapfrog steps.

### 4.1. Predictive posterior inference for the area average

The analytical area-weighted average of a spatially varying temperature or pressure function  $y(\mathbf{x})$  at an isolated measurement plane indexed by  $l \in [1, P]$ , where  $r \in [0, 1]$  and  $\theta \in [0, 2\pi)$ , is given by

$$\mu_{\text{area},l} = v_l \int_0^1 \int_0^{2\pi} T(r, \theta) h(r) dr d\theta, \tag{16}$$

where  $T$  represents the spatially varying temperature or pressure at a given axial measurement plane, and

$$v_l = \frac{r_{\text{outer},l} - r_{\text{inner},l}}{\pi(r_{\text{outer},l}^2 - r_{\text{inner},l}^2)} \text{ and } h_l(r) = r(r_{\text{outer},l} - r_{\text{inner},l}) + r_{\text{inner},l}, \tag{17}$$

where  $r_{\text{inner},l}$  is the inner radius and  $r_{\text{outer},l}$  the outer radius for plane  $l$ . For the joint distribution (2) constructed across  $P$  axial planes, one can express the area average as

$$\left[ \begin{array}{c} \mathbf{f}(\mathbf{X}) \\ \int f(\mathbf{z}) \mathbf{h}(\mathbf{z}) d\mathbf{z} \cdot \mathbf{v} \end{array} \right] \sim \mathcal{N} \left( \mathbf{0}, \left[ \begin{array}{c} \mathbf{K}_{\bullet\bullet} + \Sigma \\ \int \mathbf{K}(\mathbf{z}, \mathbf{X}) \odot \mathbf{h}(\mathbf{z}) d\mathbf{z} \cdot \mathbf{v} \\ \int \mathbf{K}(\mathbf{z}, \mathbf{X}) \odot \mathbf{h}(\mathbf{z}) d\mathbf{z} \cdot \mathbf{v} \left( \mathbf{v}^T \int \int \mathbf{K}(\mathbf{z}, \mathbf{z}) \odot \mathbf{h}^2(\mathbf{z}) d\mathbf{z} d\mathbf{z} \cdot \mathbf{v} \right) \end{array} \right] \right), \tag{18}$$

where  $\mathbf{v} = (v_1, \dots, v_P)$ ,  $\mathbf{h} = (h_1, \dots, h_P)$  and  $\mathbf{z} \in \{(r, \theta) : r \in [0, 1], \theta \in [0, 2\pi), \rho \in \{1, \dots, P\}\}$ . Through this construction, we can define the area-average spatial quantity as multivariate Gaussian distribution with mean

$$\mu_{\text{area}}[f] = \left( \mathbf{v} \int \mathbf{K}(\mathbf{z}, \mathbf{X}) \odot \mathbf{h}(\mathbf{z}) d\mathbf{z} \right) \mathbf{G}^{-1} \mathbf{f}, \tag{19}$$

where  $\mu_{\text{area}} \in \mathbb{R}^P$ . The posterior here is obtained by averaging over the hyperparameters; the covariance is given by

$$\Sigma_{\text{area}}^2[f] = \left( \mathbf{v}^T \int \int \mathbf{K}(\mathbf{z}, \mathbf{z}) \odot \mathbf{h}^2(\mathbf{z}) d\mathbf{z} d\mathbf{z} \cdot \mathbf{v} \right) - \left( \int \mathbf{K}(\mathbf{z}, \mathbf{X}) \odot \mathbf{h}(\mathbf{z}) d\mathbf{z} \cdot \mathbf{v} \right) \cdot \mathbf{G}^{-1} \cdot \left( \int \mathbf{K}(\mathbf{X}, \mathbf{z}) \odot \mathbf{h}(\mathbf{z}) d\mathbf{z} \cdot \mathbf{v} \right). \tag{20}$$

We remark that although the integral of the harmonic terms is 0, the hyperparameters associated with those terms do not drop out and thus do contribute to the overall variance.

### 5. Decomposition of Uncertainty

To motivate this section, we consider the following questions:

- Can we ascertain whether the addition of instrumentation will alter the area-average of a single measurement plane (and its uncertainty)?
- How do we determine whether we require more sensors of the present variety, or higher precision sensors at present measurement locations at a given plane?
- In the case of the former, can we determine where these additional sensors should be placed?

As instrumentation costs in aeroengines are expensive, statistically justified reductions in instrumentation can lead to substantial savings per engine test. Thus, the answers to the questions above are important. At the same time, greater accuracy in both the spatial pattern and its area-average can offer improved aerothermal inference. To aid our mathematical exposition, for the remainder of the methodology section of this article, we restrict our analysis to a single measurement plane. In other words,  $P = 1$  and thus the planar kernel does not play a role in the Gaussian random field.

5.1. Spatial field covariance decomposition

To offer practical solutions to aid our inquiry, we utilize the *law of total covariance* which breaks down the total covariance into its composite components  $\text{cov}[\mathbb{E}(\mathbf{f}^*|\mathbf{f}, \mathbf{X})]$  and  $\mathbb{E}(\text{cov}[\mathbf{f}^*|\mathbf{f}, \mathbf{X}])$ . These are given by

$$\text{cov}[\mathbb{E}(\mathbf{f}^*|\mathbf{f}, \mathbf{X})] = \mathbf{K}_{**}^T \mathbf{K}_{**}^{-1} \mathbf{\Psi}_f \mathbf{K}_{**}^{-1} \mathbf{K}_{**} \tag{21}$$

and

$$\mathbb{E}(\text{cov}[\mathbf{f}^*|\mathbf{f}, \mathbf{X}]) = \mathbf{K}_{**} - \mathbf{K}_{**}^T \mathbf{K}_{**}^{-1} \mathbf{K}_{**}, \tag{22}$$

where

$$\mathbf{\Psi}_f = (\mathbf{K}_{**}^{-1} + \mathbf{\Sigma}^{-1})^{-1}; \tag{23}$$

once again we are marginalizing over the hyperparameters. We term the uncertainty in (21) the *impact of measurement imprecision*, that is, the contribution owing to measurement imprecision. Increasing the precision of each sensor should abate this uncertainty. The remaining component of the covariance is given in (22), which we define as *spatial sampling uncertainty*, that is, the contribution owing to limited spatial sensor coverage (see Pianko and Wazelt, 1983). Note that this term does not have any measurement noise associated with it. Adding more sensors, particularly in regions where this uncertainty is high, should diminish the contribution of this uncertainty.

5.2. Decomposition of area average uncertainty

Extracting 1D metrics that split the contribution of the total area-average variance into its composite spatial sampling  $\sigma_{\text{area,s}}^2$  and impact of measurement imprecision  $\sigma_{\text{area,m}}^2$  is a direct corollary of the law of total covariance, that is,

$$\sigma_{\text{area,s}}^2 = \left( v_l^2 \int \int \mathbf{K}(\mathbf{z}, \mathbf{z}) h_l^2(\mathbf{z}) d\mathbf{z} d\mathbf{z} \right) - \left( v_l \int \mathbf{K}(\mathbf{z}, \mathbf{X}) h(\mathbf{z}) d\mathbf{z} \right) \cdot \mathbf{K}_{**}^{-1} \cdot \left( v_l \int \mathbf{K}(\mathbf{X}, \mathbf{z}) h_l(\mathbf{z}) d\mathbf{z} \right) \tag{24}$$

and

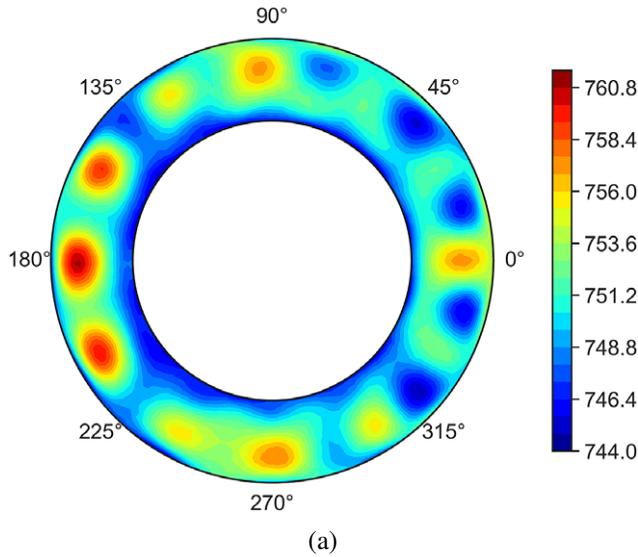
$$\sigma_{\text{area,m}}^2 = \left( v_l \int \mathbf{K}(\mathbf{z}, \mathbf{X}) h_l(\mathbf{z}) d\mathbf{z} \right) \cdot \mathbf{K}_{**}^{-1} \mathbf{\Psi}_f \mathbf{K}_{**}^{-1} \cdot \left( v_l \int \mathbf{K}(\mathbf{X}, \mathbf{z}) h_l(\mathbf{z}) d\mathbf{z} \right), \tag{25}$$

where  $v_l$  and  $h_l$  were defined previously in (17). We remark here that whole-engine performance analysis tools usually require an estimate of sampling and measurement uncertainty—with the latter often being further decomposed into contributions from static calibration, the data acquisition system, and additional factors. Sampling uncertainty has been historically defined by the sample variance (see 8.1.4.4.3 in Saravanmuttoo, 1990). We argue that our metric offers a more principled and practical assessment.

Guidelines on whether engine manufacturers need to (a) add more instrumentation, or (b) increase the precision of existing measurement infrastructure can then follow, facilitating a much-needed step-change from prior efforts (Pianko and Wazelt, 1983; Saravanmuttoo, 1990).

6. Isolated Plane Studies with the Simple Prior

To set the stage for an exposition of our formulations and algorithms, we design the spatial temperature distribution shown in Figure 4. This field comprises of five circumferentially varying harmonics  $\omega = (1,4,7,12,14)$  that have different amplitudes and phases going from the hub to the casing. This synthetic data was generated by radially interpolating four distinct circumferentially varying Fourier series’ expansions—all with the same harmonics  $\omega$ . This interpolation was done using a cubic polynomial to



**Figure 4.** Ground truth spatial distribution of temperature.

**Table 1.** Summary of sampling locations for the default test case.

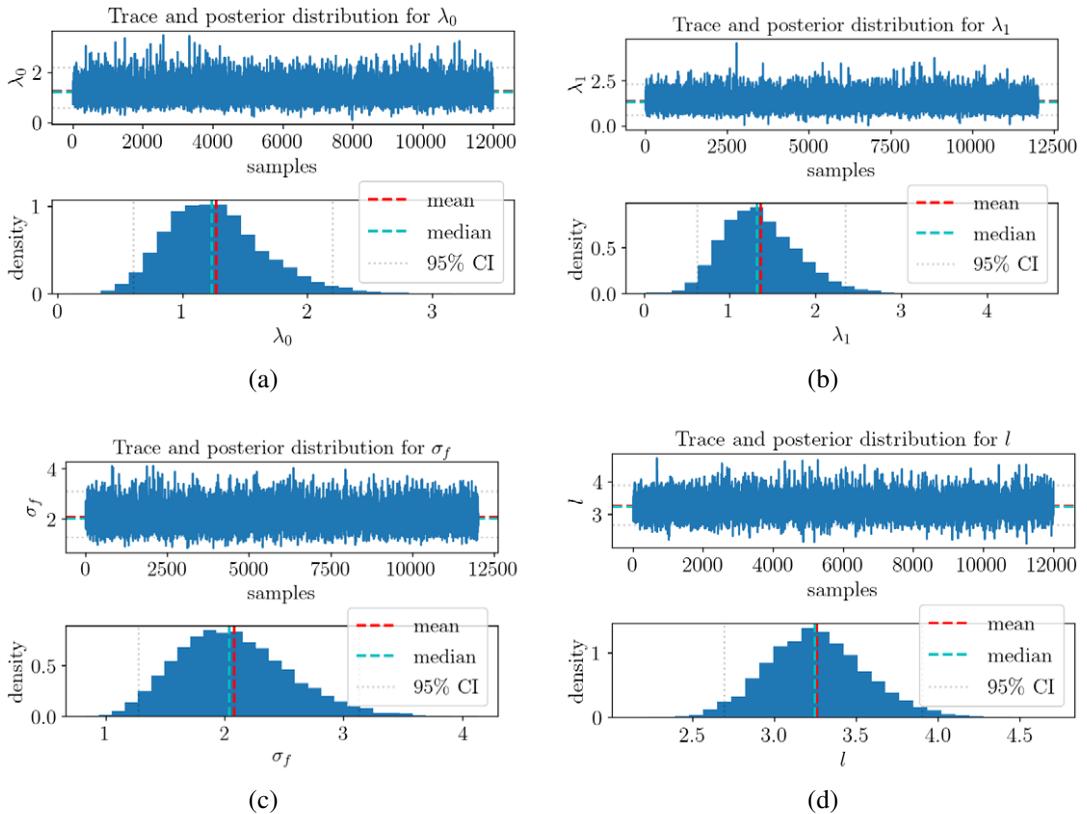
Property name	Symbol	Value(s)
Rake arrangement	$\theta$	(12°, 55°, 97°, 170°, 215°, 305°)
Probe locations (non-dimensional)	$\mathbf{r}$	(0.07,0.2,0.35,0.5,0.66,0.8,0.95)

extrapolate the field across the annulus. A small zero-mean Gaussian noise with a standard deviation of 0.1 Kelvin (K) is added to the spatial field. The computed area average mean of the field is 750.94 K.

**6.1. Spatial field estimation**

Consider a six-rake arrangement given by instrumentation placed as per Table 1, representative of certain planes in an engine. Note that rake arrangements in engines are driven by structural, logistical (access), and flexibility constraints, and thus, it is not uncommon for them to be periodically positioned. As will be demonstrated, the rake arrangements have an impact on the spatial random field and the area average.

We set our *simple priors* (non-sparsity promoting) as per (12); harmonics to  $\omega = (1,4,7,12,14)$ , and extract training data from the circumferential and radial locations provided in Table 1. Traceplots for the NUTS sampler for hyperparameters  $\lambda_0, \lambda_1, \sigma_f$  and  $l$  are shown in Figure 5 for the chosen rake placement. Note that these plots exclude the first few burn-in samples and are the outcome of four parallel chains. The visible stationarity in these traces, along with their low autocorrelation values give us confidence in the convergence of NUTS for this problem. The Gelman-Rubin statistic for all hyperparameters above was found to be 1.00; the Geweke z-scores were found to be well within the two standard deviation limit. Figure 6a plots the mean of the resulting spatial distribution (ensemble averaged), while Figure 6b plots its standard deviation. In comparing Figure 6a with Figure 4, we note that in addition to adequately approximating the radial variation (cooler hub and warmer casing), our methods are able to delineate the relatively hotter left half-annulus and its three hot spots at 150°, 180°, and 210°. This is especially surprising given the fact that we have 5 spatial harmonics and only 6 and 7 rakes, and not the 11 needed as per the Nyquist bound. A circumferential slice of these plots is shown in Figure 6c at a radial height of



**Figure 5.** Trace plots for the MCMC chain for some of the hyperparameters (a)  $\lambda_0$ ; (b)  $\lambda_1$ ; (c)  $\sigma_f$ ; and (d)  $l$ .

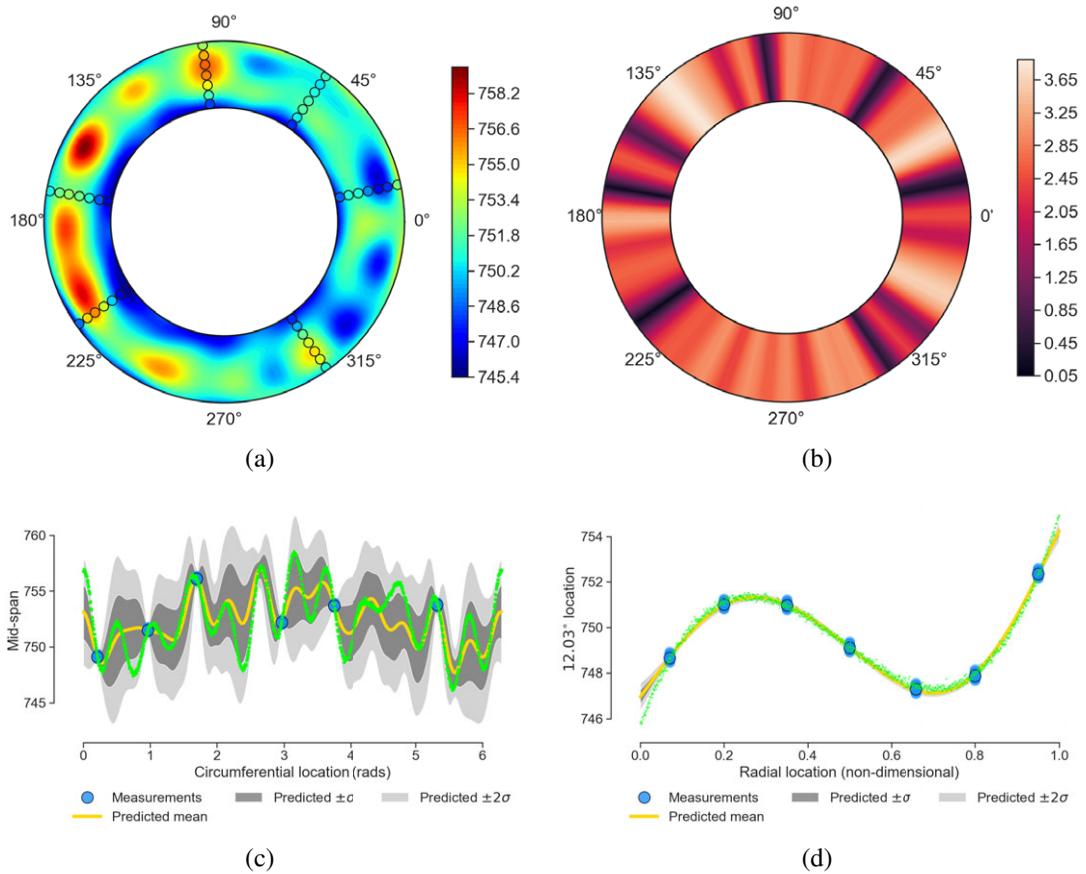
0.5 mm; a radial slice is shown in Figure 6d at a circumferential location of 0.21 radians. Here, we note that the true spatial variation (shown as a green line) lies within the standard deviation intervals in the circumferential direction, demonstrating that our approach is able to provide sufficiently accurate uncertainty estimates in this case.

For completeness, we plot the decomposition of the uncertainty in Figure 7, where the contribution of *impact of measurement imprecision* is, on average, an order of magnitude lower than that of *spatial sampling*. When inspecting these plots one can state that reductions in the overall uncertainty can be obtained by adding additional rakes at 215° and 300° (see Figure 7b).

## 6.2. Spatial field uncertainty variations

To assist in our understanding of the spatial uncertainty decompositions above, we carry out a study varying the number of rakes and their spatial locations. Figure 8 plots the two components of uncertainty for 1, 2, and 3 rakes, while Figure 9 plots them for 9, 10, and 11 rakes. There are several interesting observations to report.

First, the impact of measurement uncertainty deviates from the location of the sensor with the accumulation of more rakes. For instance, in the case with one rake in Figure 8a, light blue and red contours can be found near each sensor measurement. However, as we add more rakes, there seems to be a phase shift that is introduced to this pattern. This is because the measurement uncertainty will not necessarily lie around the rakes themselves—especially if knowledge about a sensor's measurement can be obtained from other rakes—but rather be in regions that are most sensitive to that particular sensor's value. Furthermore, in the case with an isolated rake, the impact of measurement imprecision locally will



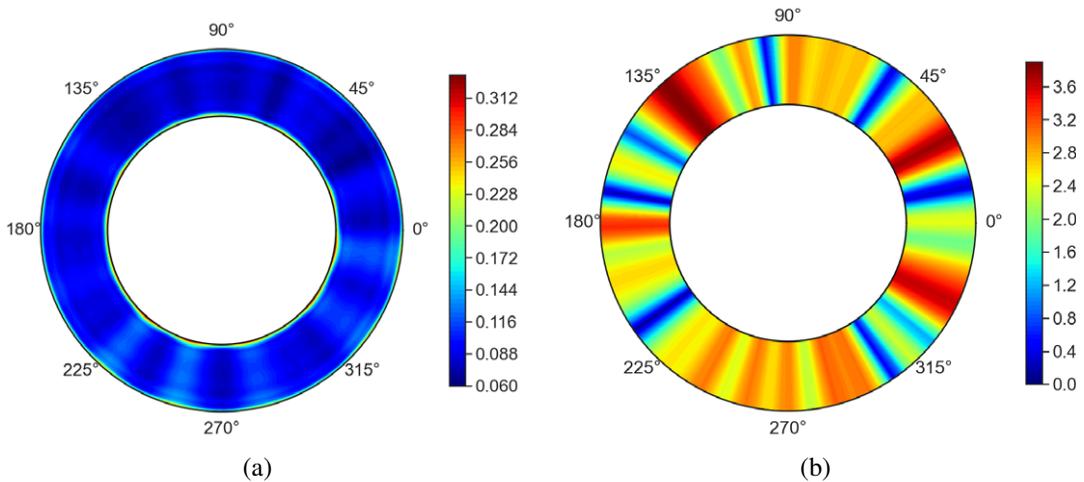
**Figure 6.** Spatial distributions for (a) the mean and (b) the standard deviation, generated using an ensemble average of the iterates in the MCMC chain (accepted samples with burn-in removed plus thinning, across four chains), and a circumferential slice at (c) mid-span and a radial slice at (d) 12.03°. Green circular markers are the true values for this synthetic case.

be very close to the  $\sigma_m$  value assigned as the measurement noise. However, with the addition of more instrumentation, the impact of measurement imprecision will increase, as observed in Figure 8g–i, before decreasing again once the spatial pattern is fully known (see Figure 9g–i).

Second, when the number of rakes is equal to 11, the spatial sampling uncertainty in the circumferential direction will not vary, and thus the only source of spatial sampling uncertainty will be due to having only seven radial measurements. The former is due to the fact that with five harmonics, we have 11 circumferential unknowns. This is clearly seen in Figure 9f. It should be noted that the position of the rakes can abate the uncertainties observed. This raises a very important point concerning experimental design, and how within a Bayesian framework, sampling uncertainty can be significantly reduced when the rakes are accordingly positioned.

### 6.3. Bayesian area average

Bayesian area average estimates are obtained by integrating the spatial approximation (as per (18)) at each iteration of the previously presented MCMC chain—ignoring the burn-in samples—and averaging over sample realizations. The deficiency of the sector area-average compared to the Bayesian area average is apparent when one studies its convergence.



**Figure 7.** Decomposition of the standard deviations in the temperature: (a) impact of measurement imprecision, and (b) spatial sampling.

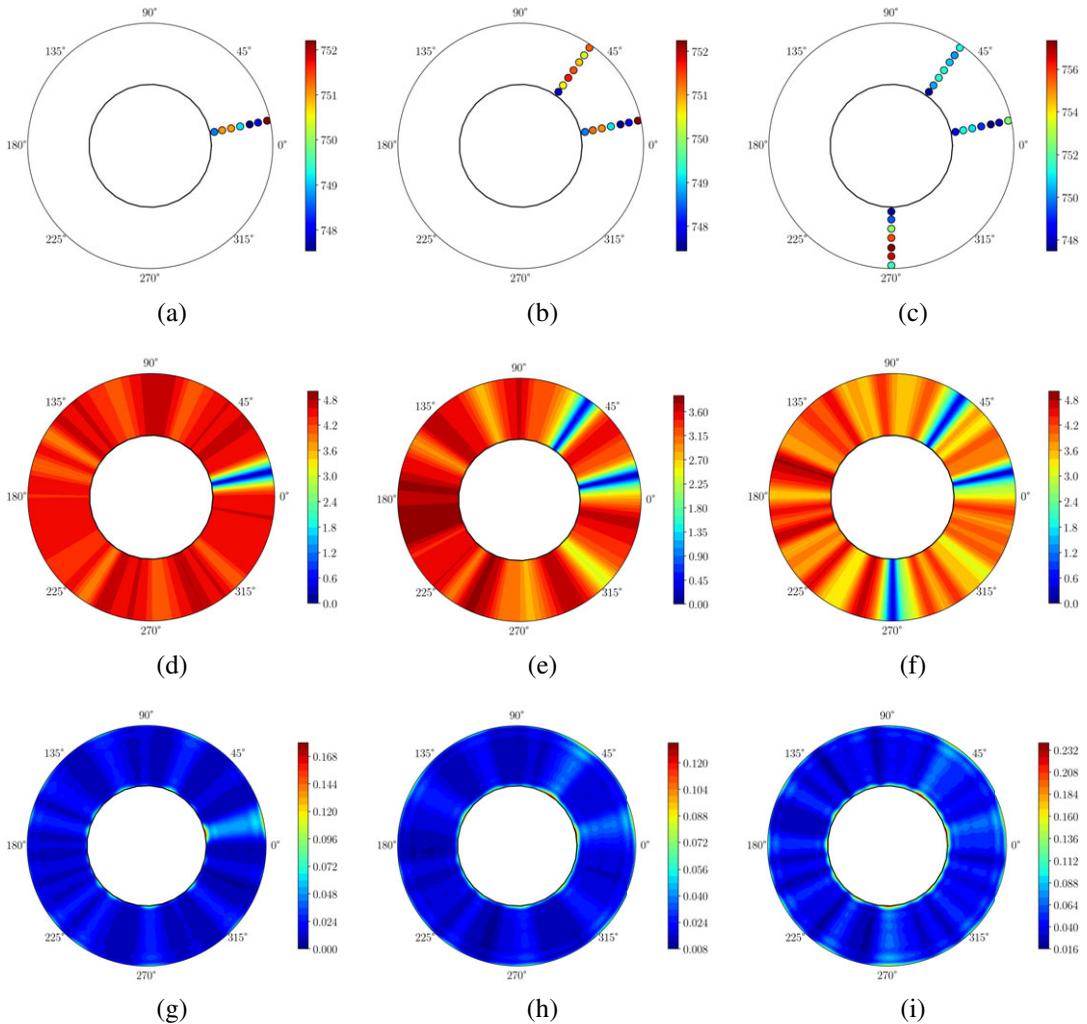
To do this, we sample our true spatial distribution at 40 different randomized circumferential locations for different numbers of rakes, while maintaining the number of radial probes and their locations. The circumferential locations are varied by randomly selecting rake positions between  $0^\circ$  and  $355^\circ$  inclusive, in increments of  $5^\circ$ . Figure 10a plots the resulting sector area-average. The yellow line represents the true area-average and the shaded gray intervals around it reflect the measurement noise. It is clear that the addition of rakes does not necessarily result in any convergence of the area-average temperature. Furthermore, the reported area-average is extremely sensitive to the placement of the rakes; in some cases, a  $\pm 2$  K variation is observed. In Figure 10b we plot the reported mean for each randomized trial using our Bayesian framework. Not only is the scatter less, but, in fact, after 10 rakes we see that reported area-averages lie within  $\pm 2\sigma_m$ , where  $\sigma_m$  is the measurement noise.

At this stage, it is worth reemphasizing how this translates to efficiency. Following the analysis of Seshadri et al. (2020a), the impact of this uncertainty can be propagated through to subsystem efficiencies. Consider an engine representative low-pressure isentropic turbine. Working with the average uncertainty in temperature for a given number of rakes—obtained from Figure 10a,b—the average uncertainty in efficiency can be computed. Table 2 contrasts these uncertainties for the sector weighted area average and proposed Bayesian area average. It is clear that the latter metric offers more representative efficiency estimates. Beyond efficiency, we reiterate that temperature measurements themselves form the backbone of EGT (see Figure 1), which feeds into remaining useful life estimates. Large uncertainties in EGT will likely impact maintenance and overhaul decisions.

This makes a compelling case for replacing the practice for computing area-averages in turbomachinery via sector weights with the proposed Bayesian treatment. As a side note, the idea of computing a Bayesian area average across an annulus has motivated the development of a more physically representative Bayesian *mass* average (see Seshadri et al., 2022).

Next, we study the decomposition of the area-average variance in these randomized experiments and plot their *spatial sampling* and *impact of measurement imprecision* components (see (24) and (25)) in Figure 11. As before, the measurement noise is demarcated as a solid yellow line. There are interesting observations to make regarding these results.

First, the impact of measurement uncertainty increases with more instrumentation, till the model is able to adequately capture all the Fourier harmonics (after 11 rakes); we made an analogous finding when studying the spatial decomposition plots. This intuitively makes sense, as the more instrumentation we



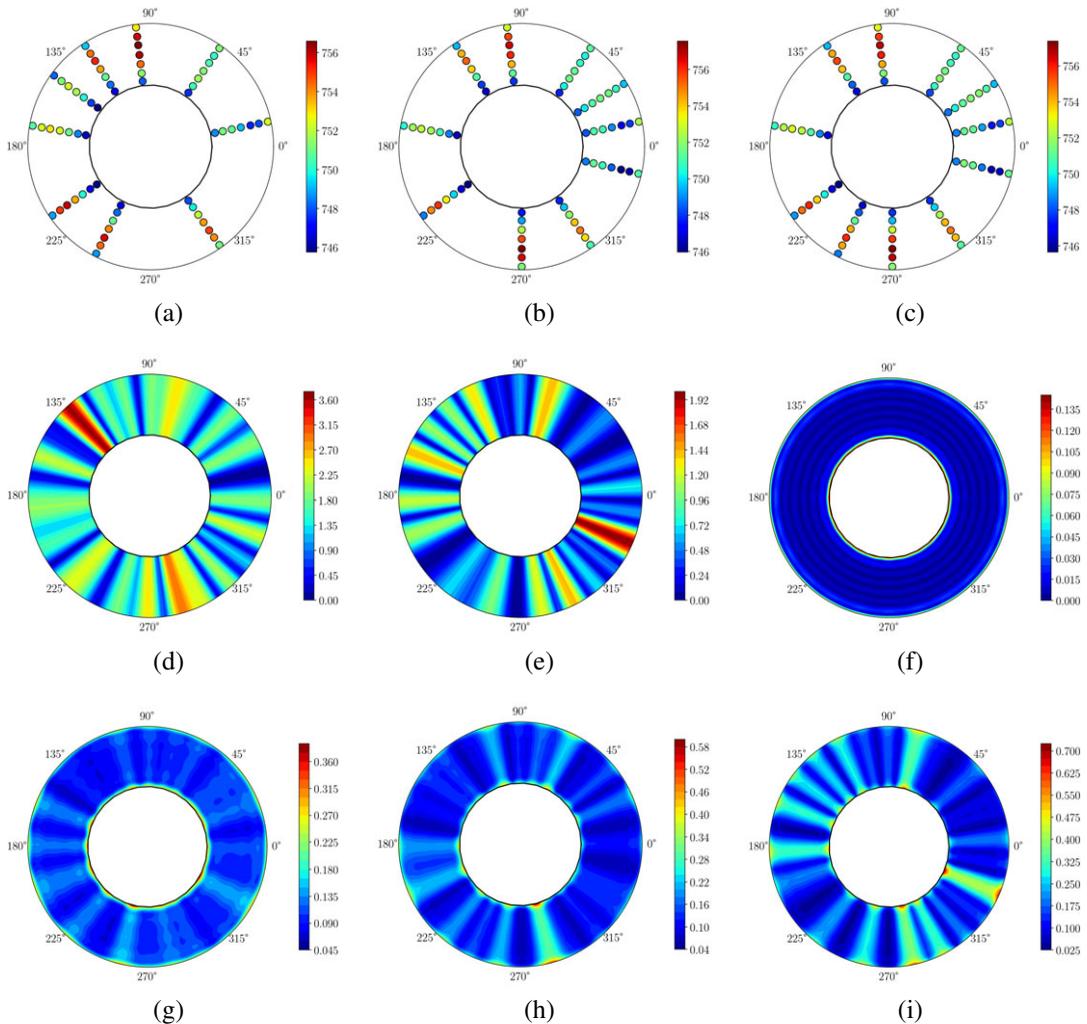
**Figure 8.** Decomposition of the standard deviations in the temperature for different number of rakes where the top row shows the measurement locations, the middle row illustrates the spatial sampling uncertainty, and the bottom row shows the impact of measurement imprecision. Results are shown for (a,d,g) one rake; (b,e,h) two rakes; and (c,f,i) three rakes.

add, the greater the impact of measurement uncertainty. It is also worth noting that numerous rake arrangements can be found that curtail this source of uncertainty, many far below the threshold associated with the measurement noise.

Second, across the 40 rake configurations tested, spatial sampling uncertainty contributions were found to be very similar when using only two to three rakes. The variability in spatial sampling uncertainty decreases significantly when the number of rakes is sufficient to capture the circumferential harmonics. Thereafter, it is relatively constant, as observed by the collapsing of the red circles in Figure 11.

### 7. Transfer Learning Results with the Sparsity Promoting Prior

In this results section, we present the results of our inductive transfer learning framework with sparsity-promoting priors.

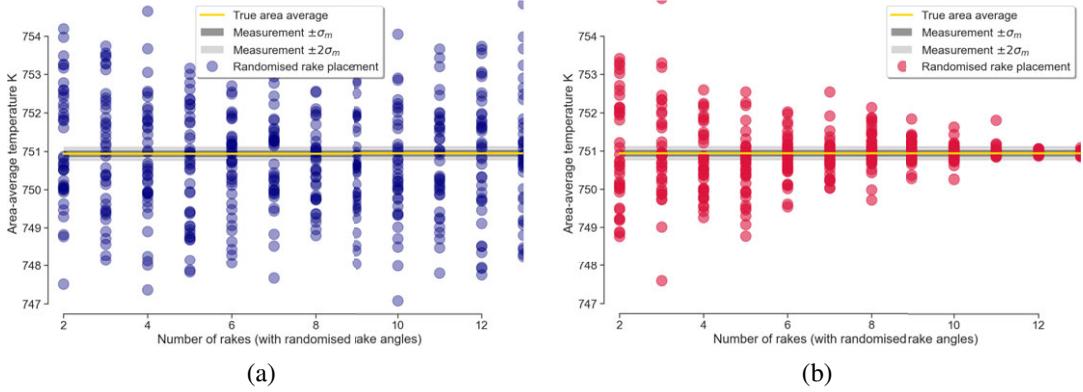


**Figure 9.** Decomposition of the standard deviations in the temperature for different number of rakes where the top row shows the measurement locations, the middle row illustrates the spatial sampling uncertainty, and the bottom row shows the impact of measurement imprecision. Results are shown for (a,d,g) 9 rakes; (b,e,h) 10 rakes; and (c,f,i) 11 rakes.

### 7.1. Transfer learning by splitting instrumentation

In this first case study, we consider traverse temperature measurements taken from a research turbine rig. Figure 12a shows the traverse locations at a temperature station, while Figure 12b shows the resulting steady-state temperature field. A fast Fourier transform was carried out on the temperature field at the hub, mid-span, and tip along the circumferential direction; the resulting amplitudes are captured in Figure 12c–e. It is clear that wave numbers 1, 12, and 24 are dominant. Additionally, we note that the signal is generally sparse, and thus utilization of the aforementioned sparsity-promoting priors seems like a sensible decision.

Let us assume that we can sample this spatial field using only 4 circumferential rakes, each fitted with 6 probes. Assume further that we are permitted to do this twice, with different circumferential rake placements. In both cases the rakes are clocked with respect to the upstream components, that is, they are aligned to be at the same pitchwise location, so as not to capture any upstream wakes.

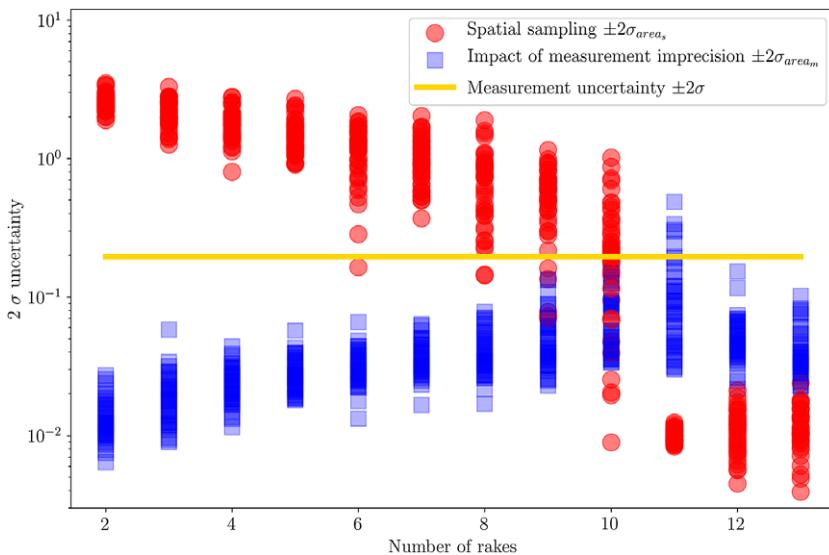


**Figure 10.** Convergence of (a) the sector weighted area-average and (b) the Bayesian area-average (only mean reported) for 40 randomized arrangements of rake positions.

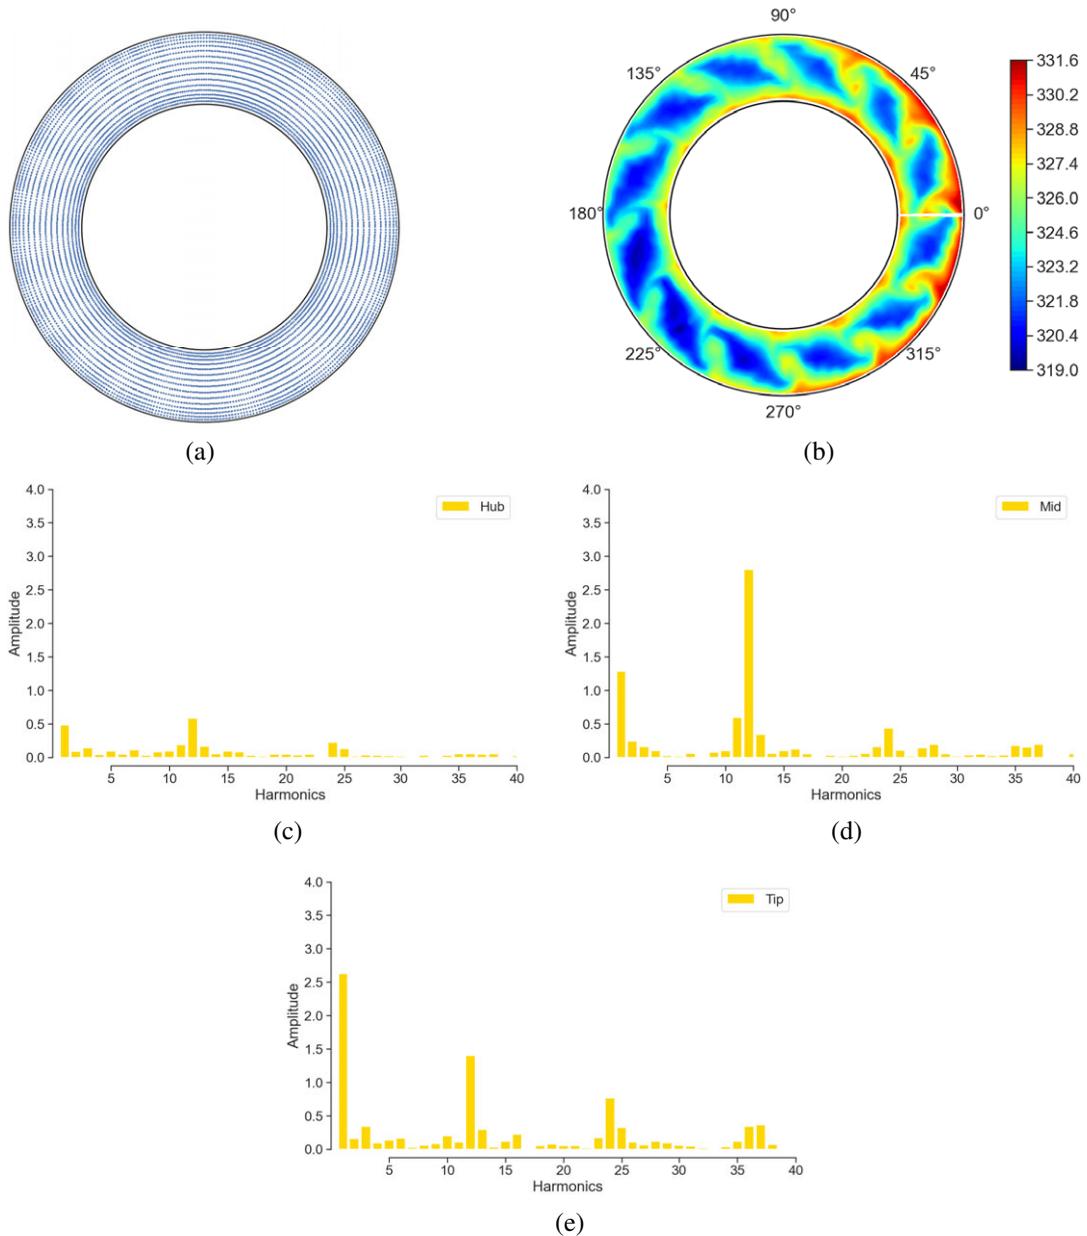
**Table 2.** Sample back-of-the-envelope uncertainty calculations for a representative isentropic turbine based on assuming both inlet and exit planes have the same uncertainty in stagnation temperature; stagnation pressures are assumed constant.

No. of rakes	Sector area-weighted average		Bayesian average	
	Temperature (K)	Efficiency (%)	Temperature (K)	Efficiency (%)
2	$\pm 1.53$	$\pm 0.62$	$\pm 1.52$	$\pm 0.61$
5	$\pm 1.51$	$\pm 0.61$	$\pm 0.81$	$\pm 0.32$
10	$\pm 1.74$	$\pm 0.69$	$\pm 0.21$	$\pm 0.08$
12	$\pm 1.63$	$\pm 0.65$	$\pm 0.04$	$\pm 0.01$

Note: All reported values are standard deviations based on Figure 10a,b.



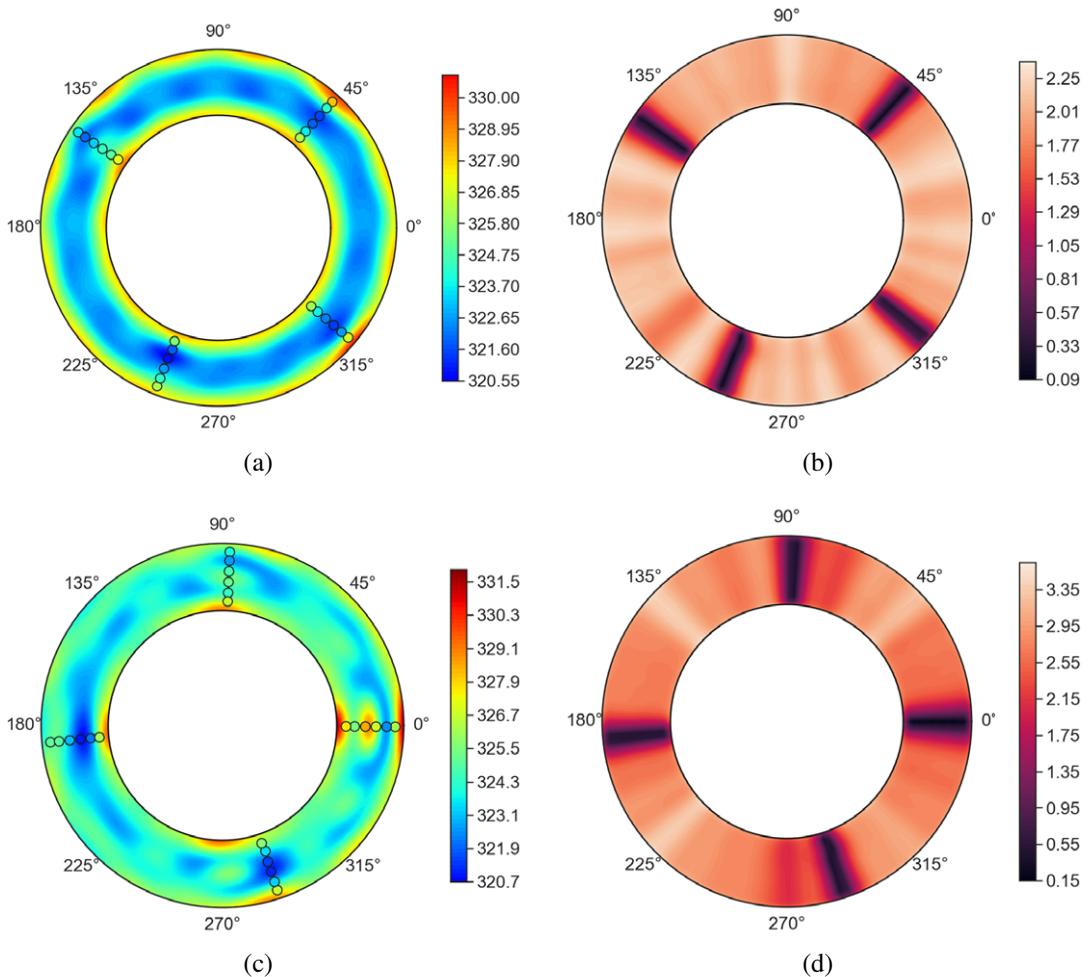
**Figure 11.** Decomposition of area-average spatial sampling and impact of measurement imprecision area-average values for 40 randomized arrangements of rake positions.



**Figure 12.** Experimental data from an exit station in a high-pressure turbine test rig: (a) traverse locations; (b) true temperature; Fourier amplitudes at the (c) hub, (d) mid-span, and (e) tip.

Running the isolated plane model—that is, with no planar kernel—with sparsity promoting priors with  $\omega = (1, 2, \dots, 15)$  for the first of the chosen rake arrangements, we obtain annular mean and standard deviation plots as shown in Figure 13a,b. While the posterior Gaussian random field does interpolate the measurements, by inspection it is readily apparent that the spatial pattern in Figure 13a does not match the truth in Figure 12b. Results run for the second rake arrangement are shown in Figure 13c,d.

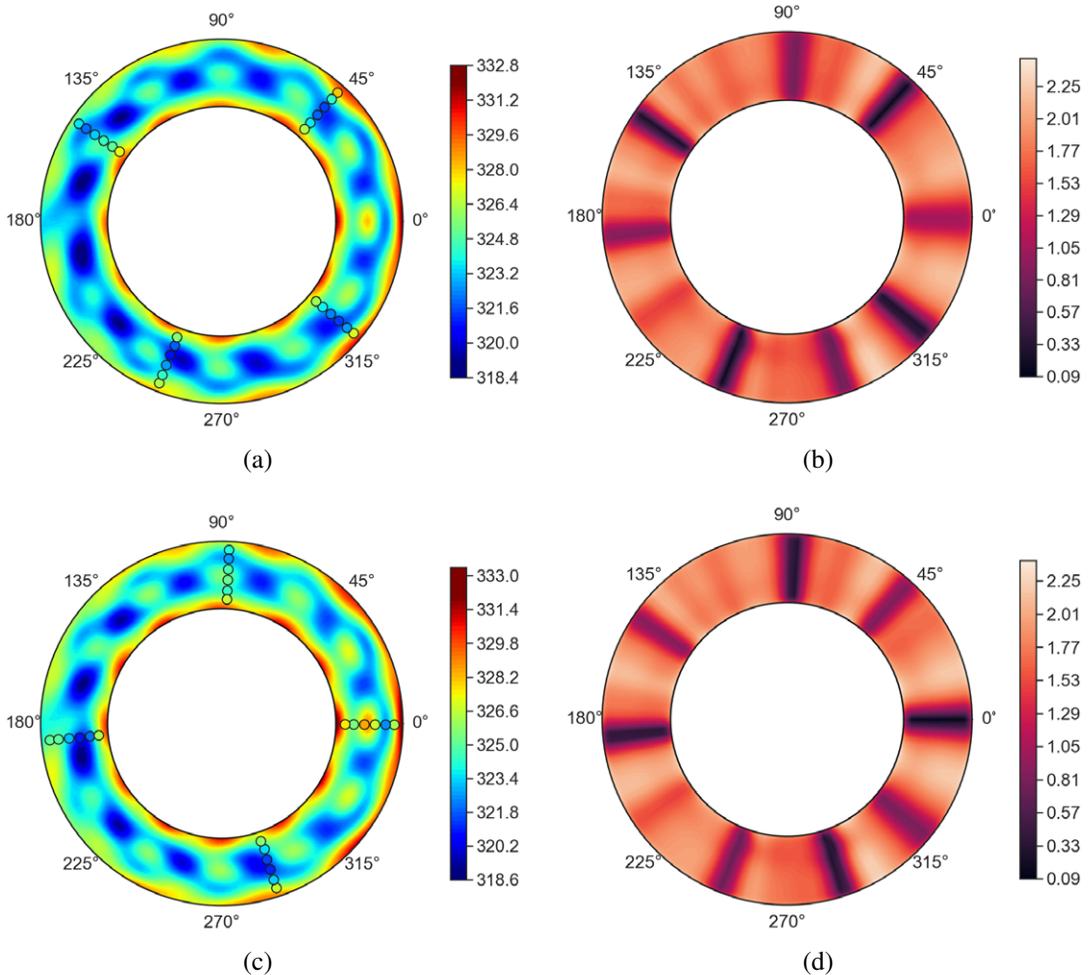
To ascertain if the transfer learning approach works, we pass these two rake arrangements as two separate measurement planes in the multi-plane model. As these measurements are from the same



**Figure 13.** Single plane calculations for the first rake arrangement (top row) and the second rake arrangement (bottom row). Posterior annular mean in (a,c); standard deviation in (b,d).

physical measurement station, they are both assigned the same value in the similarity vector, that is,  $\mathbf{s} = (1, 1)$ . Once again, we resort to sparsity-promoting priors for inference. The results are shown in Figure 14. Beyond the greater resemblance to the truth in Figure 12b, a slight reduction in the spatial uncertainty is also observed. It is clear that the model has successfully transferred information across the two measurement planes to arrive at a more representative estimate of the temperature distribution. At the same time, the model still has sufficient flexibility to offer slightly different temperature distributions for each plane individually; as we will see in the next case study, this is an extremely useful characteristic.

Circumferential plots of the single and the multi-plane yielded posterior distributions are contrasted in Figure 15 at the mid-span location; single plane results are shown in (a,c,e), while the multi-plane (inductive transfer learning) results are shown in (b,d,f). Subfigures (e) and (f) show the Fourier series amplitudes at the first plane only—these are very similar for the second plane. While it is apparent that in both cases the true pattern (shown with green circular markers) is well-captured within two standard deviations, in (b) and (d) the uncertainty is significantly reduced partly owing to the improved prediction of the mean. This comparison is important to emphasize, as it demonstrates the accuracy of the model’s predictions.



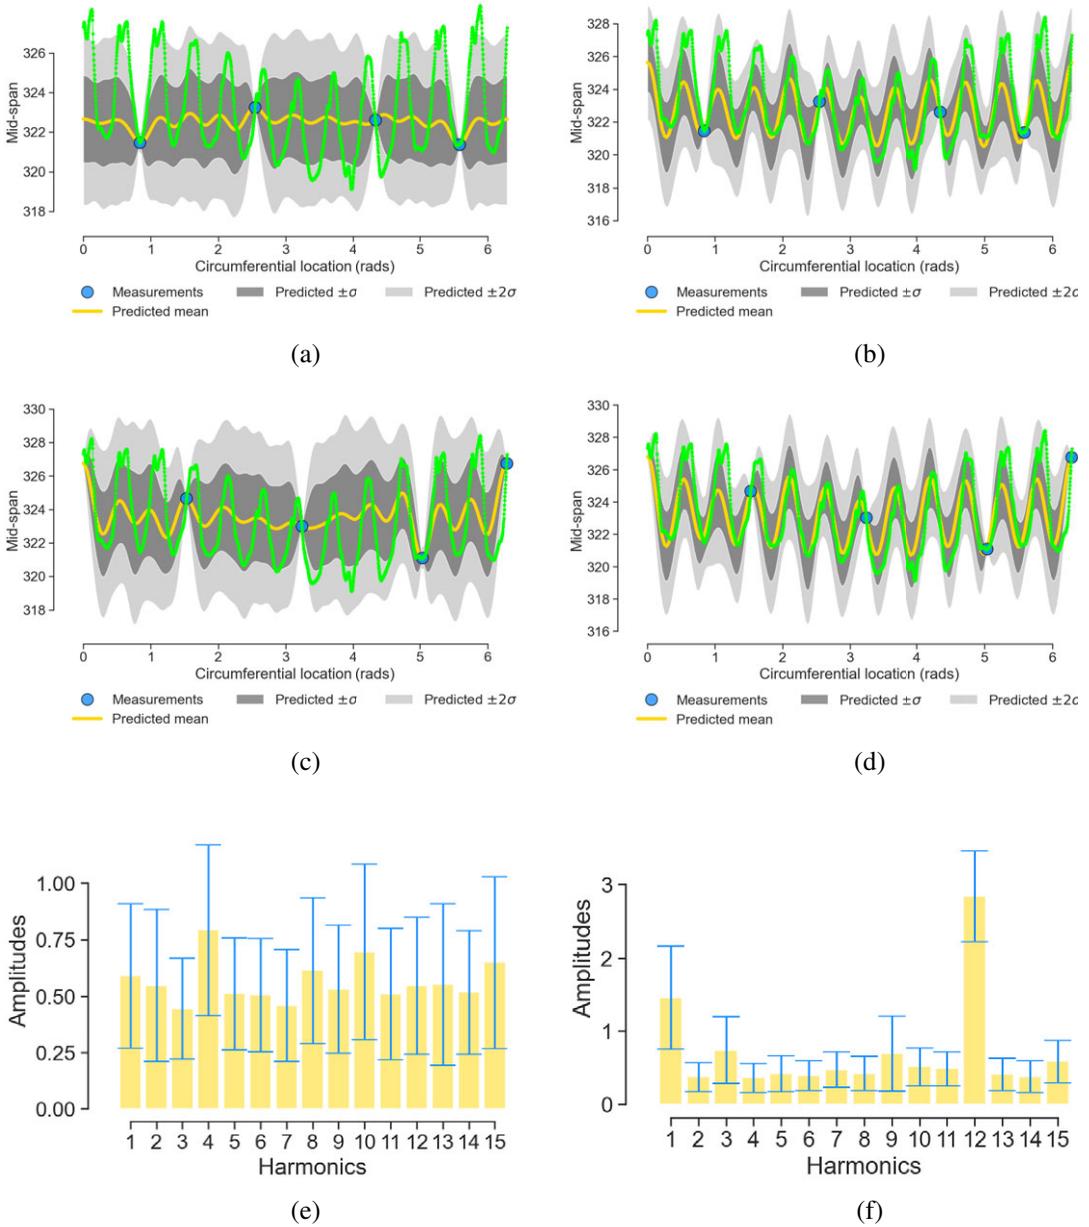
**Figure 14.** Multi-plane calculations for the first rake arrangement (top row) and the second rake arrangement (bottom row). Posterior annular mean in (a,c); standard deviation in (b,d).

### 7.2. Transfer learning with two adjacent planes from the same research engine

Next, we consider two temperature measurement planes located axially adjacent to each other in a research aeroengine. The first plane comprises 7 rakes each fitted with 7 radial temperature probes. The second plane comprises 24 thermocouples all placed at mid-span. As there are no rotating components between these two measurement stations, and owing to the fact that the flow is predominantly axial, it is hypothesized that they should have very similar temperature behavior.

The results of evaluating each measurement plane in isolation are captured in Figure 16 with circumferential distributions at mid-span for each plane. For these results, the sparsity priors were used with wave numbers  $\omega = (1, 2, \dots, 9)$ . This choice was set by the fact that the inclusion of wave numbers above 9 in the first plane leads to aliasing as the minimum angular distance between probes is  $36^\circ$ .

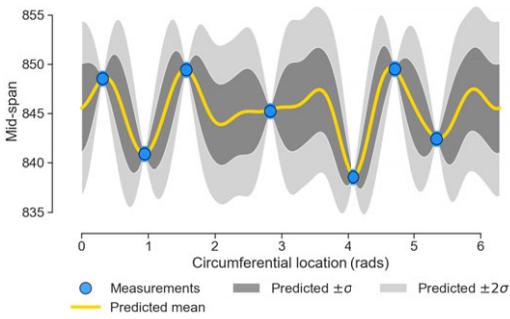
It is clear that owing to the number of measurements in the second plane, there is little uncertainty in the overall circumferential distribution. The same cannot be said for the upstream stator plane in (a). Thus, the goal here is to explore whether the transfer learning enabled multi-plane model can (a) reduce



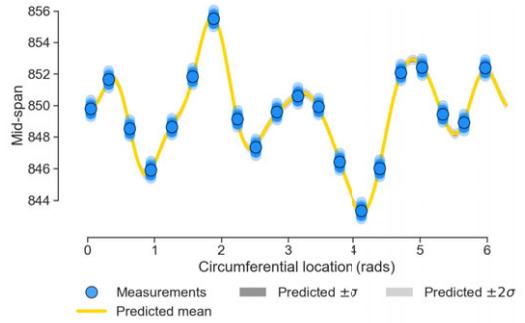
**Figure 15.** Comparison between the (a,c,e) single plane model and (b,d,f) the multi-plane transfer learning model at the mid-span location. Note that the amplitudes in (e) and (f) are only shown for the first planes (a) and (b). Green circular markers are the true values from the rig; blue markers represent a subset of four rakes.

the circumferential uncertainty in the first plane, whilst (b) reducing the radial uncertainty in the second plane.

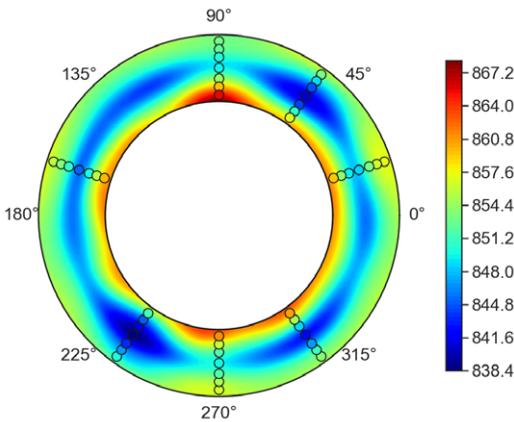
As before, for the transfer learning model we set  $s = (1, 1)$ . Note that this explicitly assumes that both planes have the same set of wave numbers, although their precise amplitudes and phases may moderately differ, as stated before.



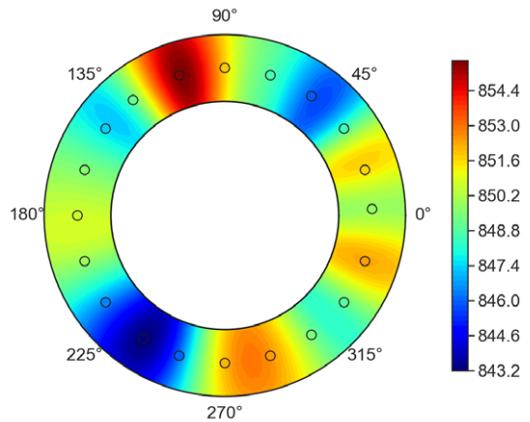
(a)



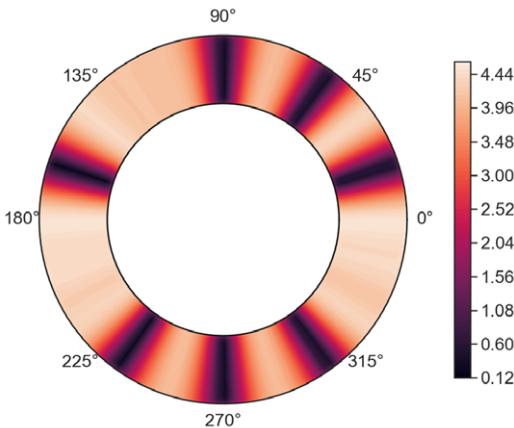
(b)



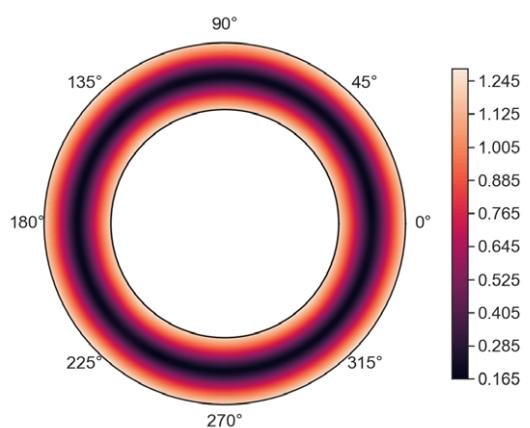
(c)



(d)



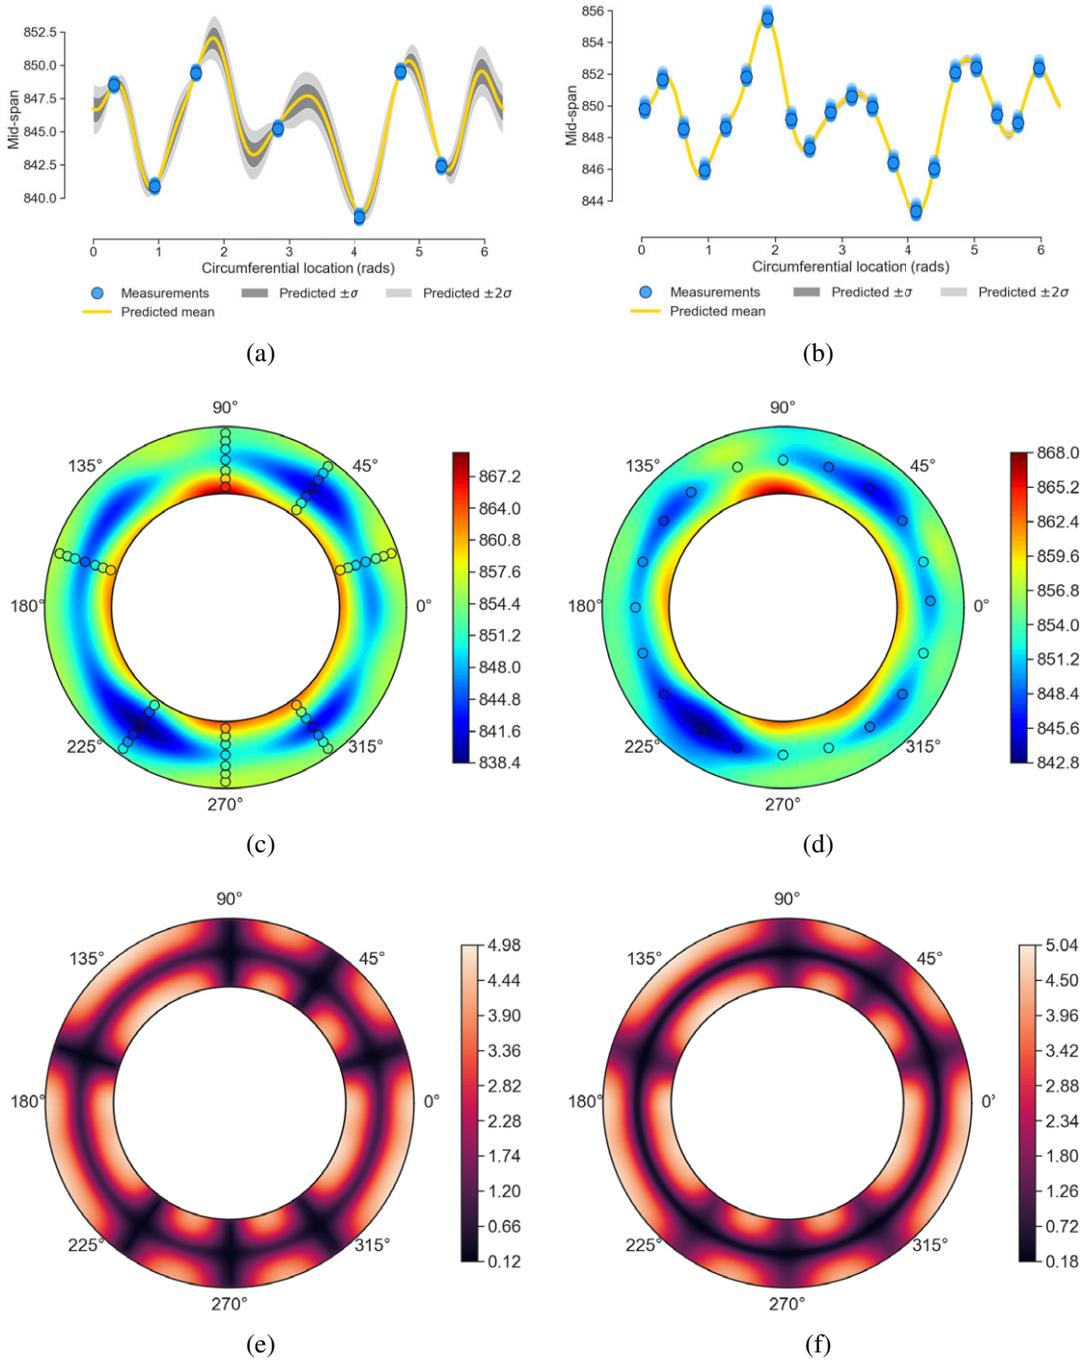
(e)



(f)

**Figure 16.** Single model results for the first plane in (a,c,e) and the second plane in (b,d,f); here each plane was run individually.

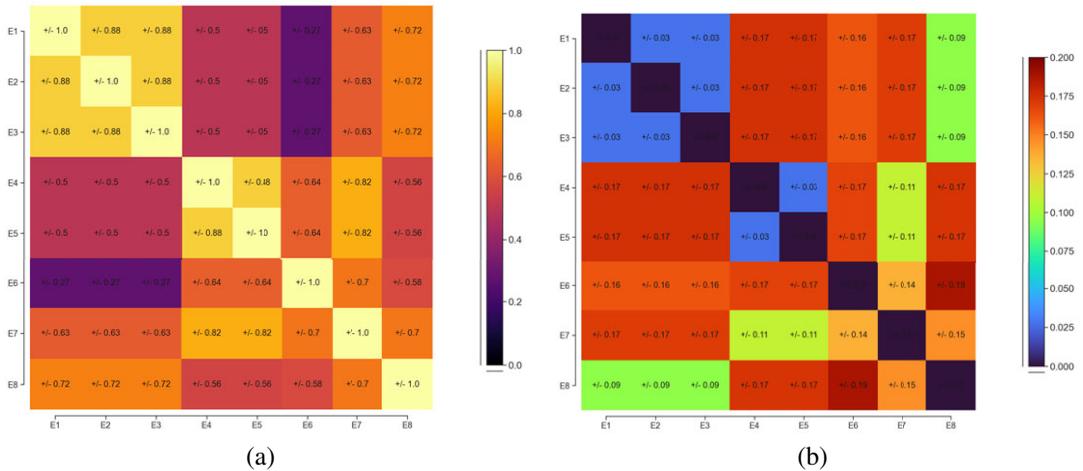
Figure 17 shows the results of the proposed model. It is clear that there is a reduction in the uncertainties in the radial direction in the burner plane, corresponding to the rake locations in the stator plane. There is also a significant reduction in the circumferential direction at mid-span region in the stator plane. Additionally, note how the radial distribution of temperature in the second plane resembles that seen on the first plane.



**Figure 17.** Multi-plane model results for the first plane in (a,c,e) and the second plane in (b,d,f).

### 7.3. Transfer learning across a fleet

One criticism of the work thus far is the reliance on  $s$ . Whilst in many cases, it is easy to establish whether two sets of measurements are similar, there may be equally many instances where such connections are difficult to draw. Ideally in such scenarios, it will be useful if the model itself can shed some light on the relative similarity between measurement planes, by virtue of radial and circumferential characteristics.



**Figure 18.** A planar correlation plot for the posterior distributions of the parameters in  $\mathcal{S}$ : (a) mean; (b) standard deviation of MCMC samples (with burn-in removed).

In this last example, we study the results of the multi-plane model on eight planes. The data chosen for this study corresponds to the temperature measurements taken from the same measurement at approximately the same throttle setting for eight different research engines. Planes E1–E3 belong to the same family, and planes E4, E5, and E7 belong to another family. Plane E8 is similar to E1–E3, but does have a different blade numbers. Additionally, planes E6 and E7 are more closely related to E4 and E5 than to planes E1–E3.

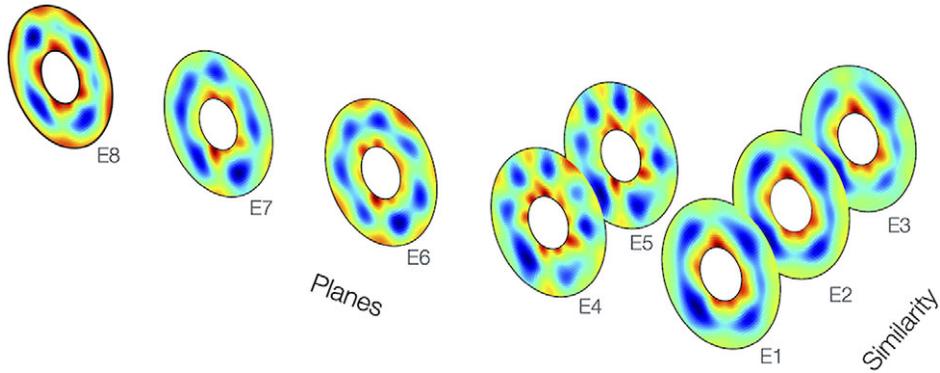
Rather than encode all these relationships in  $\mathbf{s}$ , we intentionally capture only the first few and set  $\mathbf{s} = (1, 1, 1, 2, 2, 3, 4, 5)$ . From the resulting posterior distributions of  $\eta$  and  $\xi_1, \dots, \xi_W$  and their placement in  $\mathcal{S}$ , we can construct the correlation matrix shown in Figure 18 by taking the mean of all the relevant hyperparameters  $\xi$ . Subfigure (a) plots the mean of the correlation parameters, whilst (b) plots the standard deviation—bearing in mind that each iterate of the MCMC chain will yield a correlation matrix of this form. To reiterate, these correlation values stem from the constants in (9).

From this correlation plot, we observe that many of the relationships previously mentioned but not captured in  $\mathbf{s}$  are apparent. For instance, E6 is observed to be more closely related to E4 and E5 respectively, compared to E1 (and by extension E2 and E3) with a value of 0.27. The model also rated E8’s similarity to E1 at 0.72, which seems reasonable given that both have a dominant mode four pattern. This value is higher than the correlation between E1, E2, and E3 and any of the other engines, which also aligns with our expectations. For completeness, we include the posterior mean distributions in Figure 19, where qualitatively one can observe a difference between engines that have a four-lobe versus a six-lobe pattern.

## 8. Running Times and Code

All the results shown in this section and the next one were generated from a 2.6 GHz MacBook Pro with a dual core Intel i7 chip and with 64 GB of random access memory (RAM). All run time estimates provided are based on this machine configuration.

The underpinning GP framework was coded in python 3.8 using numpy (Harris et al., 2020), scipy (Virtanen et al., 2020), and pymc3 (Salvatier et al., 2016). The eight-plane transfer learning model took approximately 5 hr, while the two-plane transfer learning results took roughly 10 min. Both models assumed sparse priors. Single plane sparse MCMC results took under 5 min, while single plane simple prior (nonsparse) results took less than a minute. For all the MCMC results, we ran four chains in parallel.



**Figure 19.** Posterior spatial means of the different planes.

## 9. Conclusions

Understanding the spatial annular pattern born from engine measurements provides valuable aerothermal insight. This article presents a transfer learning model suited for engine temperature and pressure measurements. It represents a step-up from prior averaging, uncertainty assessment, and spatial extrapolation works. Central to our contribution is the ability to transfer information across planes with a planar kernel and a user-defined input on the similarity between the different measurement planes. Beyond the results presented in this article, the proposed model has been extensively tested on measurement planes with 1–2 rakes of instrumentation when paired with planes with 6–7 rakes of instrumentation—with the goal of improving the spatial prediction even with 1–2 rakes. Across all cases, the multi-plane model yielded improved predictions. With respect to the proposed transfer learning model, future work could leverage Dirichlet prior models to learn  $\mathbf{S}$  from data directly as part of the inference process.

Given the utility of the proposed model to test and measurement, future engine design, and engine health monitoring programmes, we anticipate many forthcoming advances within this modeling paradigm, especially for prognostic and diagnostic efforts.

**Acknowledgment.** The authors are grateful to Rolls-Royce plc for permission to publish this article.

**Competing Interests.** The authors declare no competing interests exist.

**Data Availability Statement.** The data that support the findings of this study is available from Rolls-Royce. Restrictions apply to the availability of this data.

**Author Contributions.** Conceptualization: P.S., A.B.D., G.T.; Data curation: P.S., G.T.; Data visualization: P.S.; Methodology: P.S., A.B.D., G.T.; Writing—original draft: P.S., A.B.D., G.T., G.P., R.V.D., M.G. All authors approved the final submitted draft.

**Funding Statement.** The work was partly funded by the Fan and Nacelle Future Aerodynamic Research (FANFARE) project under grant number 113286, which receives UK national funding through the Aerospace Technology Institute (ATI) and Innovate UK, together with work funded by Rolls-Royce plc.

## References

- Álvarez MA and Lawrence ND (2011) Computationally efficient convolved multiple output Gaussian processes. *The Journal of Machine Learning Research* 12, 1459–1500.
- Alvarez MA, Rosasco L and Lawrence ND (2012) *Kernels for Vector-Valued Functions: A Review*. Foundations and Trends in Machine Learning, Boston, Massachusetts, USA.
- Aust J and Pons D (2019) Taxonomy of gas turbine blade defects. *Aerospace* 6(5), 58.
- Bonham C, Thorpe SJ, Erlund MN and Stevenson R (2017) Combination probes for stagnation pressure and temperature measurements in gas turbine engines. *Measurement Science and Technology* 29(1), 015002.
- Bonilla EV, Chai KM and Williams C (2008) Multi-task Gaussian process prediction. *Advances in Neural Information Processing Systems* 20, 153–160.

- Bonnet A** (2007) Avoiding high speed rejected takeoffs due to EGT limit exceedance. *Safety First* 4, 8–13.
- Bühlmann P and Van De Geer S** (2011) *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer Science & Business Media, Berlin, Germany.
- Carvalho CM, Polson NG and Scott JG** (2009) Handling sparsity via the horseshoe. In *Artificial Intelligence and Statistics*. Clearwater Beach, FL: PMLR, pp. 73–80.
- Chen P, Zabarav N and Bilionis I** (2015) Uncertainty propagation using infinite mixture of Gaussian processes and variational Bayesian inference. *Journal of Computational Physics* 284, 291–333.
- Cressie N** (2015) *Statistics for Spatial Data*. John Wiley & Sons, New York, USA.
- Cumpsty NA and Horlock JH** (2006) Averaging nonuniform flow for a purpose. *Journal of Turbomachinery* 128(1), 120–129.
- Denton JD** (2010) Some limitations of turbomachinery CFD. In *ASME Turbo Expo 2010: Power for Land, Sea, and Air*. American Society of Mechanical Engineers Digital Collection, Glasgow, UK, pp. 735–745.
- Dieck RH, Steele W and Osolobe G** (2005) *Test Uncertainty PTC 19.1*. American Society of Mechanical Engineers, New York, USA.
- Duane S, Kennedy AD, Pendleton BJ and Roweth D** (1987) Hybrid Monte Carlo. *Physics Letters B* 195(2), 216–222.
- Emory M, Iaccarino G and Laskowski GM** (2016) Uncertainty quantification in turbomachinery simulations. In *ASME Turbo Expo 2016: Turbomachinery Technical Conference and Exposition*. American Society of Mechanical Engineers Digital Collection, Seoul, South Korea.
- Flightradar24** (2021) On the flight deck to Mallorca: A special trip aboard the Lufthansa A350.
- Francis ST and Morse IE** (1989) *Measurement and Instrumentation in Engineering: Principles and Basic Laboratory Experiments*, Vol. 67. Boca Raton, FL: CRC Press.
- Gourdain N, Sicot F, Duchaine F and Gicquel L** (2014) Large eddy simulation of flows in industrial compressors: A path from 2015 to 2035. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 372, 20130323.
- Greitzer EM, Tan CS and Graf MB** (2004) *Internal Flow: Concepts and Applications*. Cambridge Engine Technology Series. Cambridge: Cambridge University Press.
- Harris CR, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, Wieser E, Taylor J, Berg S, Smith NJ, Kern R, Picus M, Hoyer S, van Kerkwijk MH, Brett M, Haldane A, del Río JF, Wiebe M, Peterson P, Gérard-Marchant P, Sheppard K, Reddy T, Weckesser W, Abbasi H, Gohlke C and Oliphant TE** (2020) Array programming with NumPy. *Nature* 585(7825), 357–362.
- Higdon D** (2002) Space and space-time modeling using process convolutions. In *Quantitative Methods for Current Environmental Issues*. Springer, London, pp. 37–56.
- Higdon D, Kennedy M, Cavendish JC, Cafo JA and Ryne RD** (2004) Combining field data and computer simulations for calibration and prediction. *SIAM Journal on Scientific Computing* 26(2), 448–466.
- Hoffman MD and Gelman A** (2014) The no-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research* 15(1), 1593–1623.
- Horowitz AM** (1991) A generalized guided Monte Carlo algorithm. *Physics Letters B* 268(2), 247–252.
- Ishwaran H and Rao JS** (2005) Spike and slab variable selection: Frequentist and Bayesian strategies. *The Annals of Statistics* 33(2), 730–773.
- Kennedy MC and O'Hagan A** (2001) Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 63(3), 425–464.
- Krige DG** (1951) A statistical approach to some basic mine valuation problems on the witwatersrand. *Journal of the Southern African Institute of Mining and Metallurgy* 52(6), 119–139.
- Liu X and Guillas S** (2017) Dimension reduction for gaussian process emulation: An application to the influence of bathymetry on tsunami heights. *SIAM/ASA Journal on Uncertainty Quantification* 5(1), 787–812.
- Lou F and Key NL** (2021) Reconstructing compressor non-uniform circumferential flow field from spatially undersampled data—Part I: Methodology and sensitivity analysis. *Journal of Turbomachinery* 143(8), 081002.
- Mailach R, Lehmann I and Vogeler K** (2008) Periodical unsteady flow within a rotor blade row of an axial compressor—Part I: Flow field at midspan. *Journal of Turbomachinery* 130(4), 041004.
- Marinai L** (2004) Gas-path diagnostics and prognostics for aero-engines using fuzzy logic and time series analysis.
- Montomoli F** (2015) *Uncertainty Quantification in Computational Fluid Dynamics and Aircraft Engines*. Springer, Heidelberg.
- Oakley J and O'Hagan A** (2002) Bayesian inference for the uncertainty distribution of computer model outputs. *Biometrika* 89(4), 769–784.
- Oakley JE and O'Hagan A** (2004) Probabilistic sensitivity analysis of complex models: A Bayesian approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 66(3), 751–769.
- Pianko M and Wazelt F** (1983) *Suitable Averaging Techniques in Non-Uniform Internal Flows*. Technical Report AGARD-AR-182. Advisory Group for Aerospace Research & Development, Loughton, Essex, UK.
- Piironen J and Vehtari A** (2017) Sparsity information and regularization in the horseshoe and other shrinkage priors. *Electronic Journal of Statistics* 11(2), 5018–5051.
- Rasmussen CE and Williams CK** (2006) *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, Massachusetts, USA.
- Rogers S and Girolami M** (2016) *A First Course in Machine Learning*. Chapman and Hall/CRC, Boca Raton, Florida, USA.

- SAE International** (2017) Inlet total-pressure-distortion considerations for gas-turbine engines. *Aerospace Information Report*, AIR1419(C).
- Salvatièr J, Wiecki TV and Fønnesbeck C** (2016) Probabilistic programming in python using PyMC3. *PeerJ Computer Science* 2, e55.
- Sanders A, Papalia J and Fleeter S** (2002) Multi-blade row interactions in a transonic axial compressor: Part I—Stator particle image velocimetry (PIV) investigation. *Journal of Turbomachinery* 124(1), 10–18.
- Saravanamuttoo H** (1990) *Recommended Practices for Measurement of Gas Path Pressures and Temperatures for Performance Assessment of Aircraft Turbine Engines and Components*. Technical Report AGARD-AR-245. Advisory Group for Aerospace Research & Development, Loughton, Essex, UK.
- Seshadri P, Duncan A, Simpson D, Thorne G and Parks G** (2020a) Spatial flow-field approximation using few thermodynamic measurements—Part II: Uncertainty assessments. *Journal of Turbomachinery* 142(2), 021007.
- Seshadri P, Duncan A and Thorne G** (2022) Bayesian mass averaging in rigs and engines. *Journal of Turbomachinery* 144(8), 081004.
- Seshadri P, Parks GT and Shahpar S** (2015) Leakage uncertainties in compressors: The case of rotor 37. *Journal of Propulsion and Power* 31(1), 456–466.
- Seshadri P, Shahpar S and Parks GT** (2014) Robust compressor blades for desensitizing operational tip clearance variations. In *ASME Turbo Expo 2014: Turbine Technical Conference and Exposition*. American Society of Mechanical Engineers Digital Collection, Düsseldorf, Germany.
- Seshadri P, Simpson D, Thorne G, Duncan A and Parks G** (2020b) Spatial flow-field approximation using few thermodynamic measurements—Part I: Formulation and area averaging. *Journal of Turbomachinery* 142(2), 021006.
- Seshadri P, Yuchi S and Parks GT** (2019) Dimension reduction via gaussian ridge functions. *SIAM/ASA Journal on Uncertainty Quantification* 7(4), 1301–1322.
- Skolidis G** (2012) Transfer learning with Gaussian processes.
- Stein ML** (2012) *Interpolation of Spatial Data: Some Theory for Kriging*. Springer Science & Business Media, New York, USA.
- Stoll F, Tremback JW and Arnaiz HH** (1979) *Effect of Number of Probes and Their Orientation on the Calculation of Several Compressor Face Distortion Descriptors*. Technical Report 72859. National Aeronautics and Space Administration, Edwards, California, USA.
- Strang G** (2012) *Computational Science and Engineering*, 2 Edn. Wellesley-Cambridge Press, Wellesley, Massachusetts, USA.
- Tibshirani R** (1996) Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* 58(1), 267–288.
- Trefethen LN** (2013) *Approximation Theory and Approximation Practice*. SIAM, Pennsylvania, USA.
- Tyacke J, Vadlamani N, Trojak W, Watson R, Ma Y and Tucker P** (2019) Turbomachinery simulation challenges and the future. *Progress in Aerospace Sciences* 110, 100554.
- Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, Burovski E, Peterson P, Weckesser W, Bright J, van der Walt SJ, Brett M, Wilson J, Millman KJ, Mayorov N, Nelson ARJ, Jones E, Kern R, Larson E, Carey CJ, Polat İ, Feng Y, Moore EW, VanderPlas J, Laxalde D, Perktold J, Cimrman R, Henriksen I, Quintero EA, Harris CR, Archibald AM, Ribeiro AH, Pedregosa F, van Mulbregt P and SciPy 1.0 Contributors** (2020) SciPy 1.0: Fundamental algorithms for scientific computing in python. *Nature Methods*, 17, 261–272.
- Wheeler AP, Sandberg RD, Sandham ND, Pichler R, Michelassi V and Laskowski G** (2016) Direct numerical simulations of a high-pressure turbine vane. *Journal of Turbomachinery* 138(7), 071003.