

Transported Snapshot Model Order Reduction for Steady Transonic Flows Simulated On Unstructured Grids

Manthan Goyal* and Aniruddha Sinha†

Dept. of Aerospace Engineering, Indian Institute of Technology Bombay, Mumbai 400076, INDIA

We present an upgraded version of the Transported Snapshot Model Order Reduction (TSMOR) approach presented recently by Nair and Balajewicz (*International Journal for Numerical Methods in Engineering*, 117(12), 2019) for shock dominated flows. The method is based on the observation that the shocks in a flow tend to change their shape and location in a smooth manner when one or more of the operating parameters are changed. The authors validated TSMOR on Cartesian meshes only, whereas here its applicability is explored for curvilinear structured grids and unstructured grids. Furthermore, a sliding boundary procedure is implemented to more faithfully preserve wall boundary conditions while using TSMOR. The method is demonstrated for transonic inviscid flow inside a two-dimensional channel with a bump on one wall. The inlet Mach number is the sole variable parameter when constructing the training database. A start is made with a structured grid for this problem owing to the relative simplicity of implementing TSMOR on it. Subsequently, the derived learning is propagated to the case of an unstructured grid for this problem. The results are validated at several transonic Mach numbers lying within the range of the training database. The final results from both the structured and unstructured cases demonstrate encouraging agreement with the full order model in all validation cases. The ultimate aim of this work is to apply the augmented TSMOR approach to more utilitarian problems, such as flow over an airfoil (in 2-D) and wings, aircrafts and missiles (in 3-D).

I. Introduction

Computational fluid dynamics (CFD) has become an indispensable tool to analyze the flow in different fluid dynamics problems, having applications in most engineering streams. With significant advancements in computer technology and infrastructure, it is becoming easier to simulate various fluid flow problems and predict related characteristics like lift and drag. However, high fidelity simulations are still computationally challenging and remain out of routine and frequent usage. As a result, the influence of CFD on parametric and time-critical applications such as aerodynamic shape optimization has been low and is yet to be realized to its full potential. The primary reason for an extended run time of traditional CFD is the many degrees of freedom which arises from the grid used by the discretized differential equations to perform CFD. The run time of a simulation can be shortened by using methods like model order reduction (MOR) [1, 2], theoretical simplifications, AI and Machine learning, etc. This paper will majorly focus on MOR, which works on the principle of dimensionality reduction.

Significant progress has been made in applying MOR in the subsonic regime, but less work has been done when it comes to the transonic and supersonic counterparts. This is because the conventional MOR methods fail to capture discontinuities like shocks in the flow. Specific procedures have been developed to model the flow in such regimes [e.g., 3–6]. In this work, we focus on the transported snapshot MOR (TSMOR) developed by Nair and Balajewicz [7]; the authors demonstrated its superiority over other prominent alternatives like Least Squares Petrov Galerkin [5] and \mathcal{L}^1 dictionary [6] approaches. Here we focus on steady flow solutions that differ in some underlying parameters in an empirical database, and evaluate the use of MOR in predicting the solutions for new parameters.

TSMOR [7] is a nonlinear MOR where several neighbouring snapshots (in parameter space) are transported (or distorted) to create a local spatial basis for a new parameter condition; the weights of these basis functions are determined by minimization of a residual. The method is described briefly in section III, where we also highlight some of its limitations as presented by Nair and Balajewicz [7]. In this work, we present some further developments of the TSMOR method to improve its applicability in problems involving complex geometries – an aspect that was not addressed fully in the original work [7]. In particular, the twofold novelty of this work lies in (a) implementation of a wall boundary

*Graduate student

† Associate Professor, AIAA Member; Corresponding author: sinha.a@iitb.ac.in

condition in TSMOR for non-planar surfaces, and (b) application to unstructured grids as opposed to the rectangular Cartesian grids employed by Nair and Balajewicz [7]. For the purpose of this study, we have chosen the inviscid flow in a 2D channel with a bump on one wall, and analyse transonic flow inside it. The updated method is able to predict the shock location and strength with reasonable accuracy; the overall flow fields also match those in the reference ‘truth’ solutions.

II. Background

Model order reduction (MOR) methods typically have two parts – extraction of a minimal basis from the data and solution of a residual minimisation problem [1].

The basis consists of a set of vector fields obtained from a training database of steady flow solutions spanning a rich enough parameter set. It is used to reconstruct or represent the solution for any parameter (within a certain range) that was omitted from the training database. Some techniques include proper orthogonal decomposition (POD) [1, 8] and reduced basis method [9]. The basis can be global or local in parameter space. In the latter case, a subset of the training database comprising a neighbourhood in parameter space is used to obtain a local basis that is useful for reconstructing the flow solution for any other parameter restricted to the same neighbourhood.

The residual minimization process for MOR has been described by LeGresley and Alonso [10]; it uses the original governing equations (or a simplification of the same) to arrive at a more accurate prediction of the solution than is possible through simple parameter interpolation approaches. Denoting a vector field of (in general, unsteady) flow variables by $\mathbf{q}(\mathbf{x}, t; \boldsymbol{\mu})$, where \mathbf{x} , t and $\boldsymbol{\mu}$ signify the spatial coordinates, time and parameter vector respectively, a fluid flow problem can be represented as

$$\frac{\partial}{\partial t} C(\mathbf{q}(\mathbf{x}, t; \boldsymbol{\mu})) + R(\mathbf{q}(\mathbf{x}, t; \boldsymbol{\mu})) = 0 \quad \text{subject to} \quad B(\mathbf{q}(\mathbf{x}, t; \boldsymbol{\mu})) = 0 \quad \text{at boundaries.} \quad (1)$$

In the above, the (possibly primitive) flow variables are mapped to the conserved flow variables by the operator C , $R(\mathbf{q})$ represents all terms other than the time derivative, and B denotes the operator encoding the boundary conditions. For the steady problems of interest here, the partial time derivative term can be dropped from eqn. (1) yielding $R(\mathbf{q}(\mathbf{x}; \boldsymbol{\mu})) = 0$. That is, R represents the residual operator which must evaluate to zero for the exact solution \mathbf{q} . Of course, since MOR approaches provide an approximate solution $\tilde{\mathbf{q}}$, we can only hope to minimize $R(\tilde{\mathbf{q}})$.

Most existing MOR approaches [2] are suited for problems involving smooth (i.e., shock-free) flow fields, so that transonic flow problems remain difficult to address. Indeed it has been seen that when conventional MOR techniques are used for shock dominated flows, a staircase pattern is observed. Over the years, several techniques have been proposed to circumvent the issue, a few of which are discussed now.

- 1) Lucia et al. [3] and Alonso et al. [4] demonstrated that addressing the shocked and unshocked regions separately may be an option. The pre-existing MOR techniques can then be applied directly to the un-shocked area, whereas a full order simulation can be run for the shocked region. There are certain drawbacks of this approach which make it unsuitable for a large class of problems. Firstly, a full order model needs to be applied for regions with shock, which is time-consuming. Secondly, it is challenging to find an optimal size for the possibly-shocked region where the full order model needs to be applied. Moreover, to merge the boundaries of the smooth flow region with that of the shocked region, multiple iterations are required making it time-consuming again.
- 2) On the other hand, Abgrall et al. [11] and Abgrall and Crisovan [6] advocated the use of the \mathcal{L}^1 norm for residual minimization, rather than the more traditional \mathcal{L}^2 norm. Indeed, it has been seen that the latter choice of norm, which gives better performance with incompressible and subsonic flows, gives rise the oscillatory and/or staircase pattern in transonic flows. Although this helps with the problems associated with \mathcal{L}^2 norm, the underlying issue of using a linear basis remains.
- 3) Recently, Nair and Balajewicz [7] proposed a novel method called Transported Snapshot MOR (TSMOR). It uses the fact that, for a large class of transonic flow problems, shockwaves tend to change in shape and location smoothly with changes in the operating parameters. Recognizing this, the authors pursued a parameterized nonlinear spatial distortion of the training flow solutions to arrive at a local linear basis. Subsequently, the extant residual minimisation approach is used to arrive at the final solution. Crucially, the issue of the staircase pattern is alleviated by bringing in non-linearity in the construction of local basis functions.

The TSMOR approach was demonstrated by Nair and Balajewicz [7] on three problems of increasing complexity. However, all of them had simple geometry and uniform cartesian grid. After presenting a detailed description of the method in section III, we then describe efforts at extending the approach to geometries with unstructured grids.

III. Transported snapshot model order reduction (TSMOR)

This method of Nair and Balajewicz [7] is based on the idea that the flow solution for an unsampled parameter can be obtained by linearly combining transported solutions corresponding to neighbouring parameters. Snapshot transport refers to first moving the grid points (but keeping their corresponding flow field variable values the same) to get a new flow field, and then interpolating it back on to the original grid so as to be useful for the subsequent steps. These distorted snapshots form the local linear basis functions, which are then used in the residual minimization problem. The process starts with a rich training database like all other empirical data-based techniques and is then pursued in two stages – offline and online. The offline stage forms the most crucial part as it is responsible for calculating a transport field that is a smooth function of the parameters; it determines the amount of distortion of the grid to be performed when constructing the basis functions. The online stage comprises both local basis evaluation and residual minimization. The mathematical formulation for both the stages is explained separately below.

A. Offline stage

Let us consider a training database with N steady flow snapshots. Each snapshot may be uniquely identified by its parameter vector $\boldsymbol{\mu}$. That is, the j^{th} snapshot is expressed as $\boldsymbol{q}(\boldsymbol{x}; \boldsymbol{\mu}_j)$, where $j \in \{1, \dots, N\}$. The transport (or grid distortion) fields are also unique to each snapshot in the training database, as the amount and nature of distortion depends on the snapshot's location in the parameter space. Moreover, each spatial coordinate direction of the problem will have its own transport field, even for a single snapshot. The transport field for the j^{th} snapshot in the i^{th} coordinate direction is denoted as $f_i(\boldsymbol{x}; \boldsymbol{\mu}_j, \boldsymbol{\mu})$. It is a function of the spatial coordinate \boldsymbol{x} ; it is parameterized by the parameter vector of the snapshot $\boldsymbol{\mu}_j$ as well as the new parameter vector $\boldsymbol{\mu}$ that is sought to be approximated. The corresponding vector transport field is denoted by $\boldsymbol{f}(\boldsymbol{x}; \boldsymbol{\mu}_j, \boldsymbol{\mu})$. For example, in a 2D Cartesian problem, $\boldsymbol{f} := [f_1, f_2]^T$. Before we describe the composition of the transport field, we establish the manner in which it is used.

The idea of transport field is that, the j^{th} snapshot should be transported somehow so that it is able to approximately reconstruct the snapshots in its parameter neighbourhood. For this, the position vector on the transported grid is defined as

$$\hat{\boldsymbol{x}}(\boldsymbol{x}; \boldsymbol{\mu}_j, \boldsymbol{\mu}) := \boldsymbol{x} + \boldsymbol{f}(\boldsymbol{x}; \boldsymbol{\mu}_j, \boldsymbol{\mu}). \quad (2)$$

Now, we define the transported flow field corresponding to the j^{th} snapshot and the neighbourhood parameter value $\boldsymbol{\mu}$ as

$$\hat{\boldsymbol{q}}_j(\hat{\boldsymbol{x}}(\boldsymbol{x}; \boldsymbol{\mu}_j, \boldsymbol{\mu})) := \boldsymbol{q}(\boldsymbol{x}; \boldsymbol{\mu}_j). \quad (3)$$

That is, even though we transport the grid, we do not change the flow field values on the individual grid points (or grid cells); this is how the snapshot gets distorted. Next, we (linearly) interpolate (and possibly extrapolate) $\hat{\boldsymbol{q}}_j$ from the distorted grid on to the original undistorted grid of the snapshot:

$$\hat{\boldsymbol{q}}_j(\hat{\boldsymbol{x}}(\boldsymbol{x}; \boldsymbol{\mu}_j, \boldsymbol{\mu})) \xrightarrow[\text{extrapolation}]{\text{interpolation/}} \check{\boldsymbol{q}}_j(\boldsymbol{x}; \boldsymbol{\mu}). \quad (4)$$

That this transported and interpolated approximate snapshot was originated with the parameter $\boldsymbol{\mu}_j$ is indicated by the j subscript. This is important since subsequently we will approximate the solution for the same parameter $\boldsymbol{\mu}$ starting from different original parameters in the learning database.

For a scalar parameter space, we write the i^{th} component of the transport field for the j^{th} snapshot as

$$f_i(\boldsymbol{x}; \boldsymbol{\mu}_j, \boldsymbol{\mu}) = \sum_{m=1}^{N_m} \sum_{a=1}^{N_a} c_{j,i}^{m,a} g_{i,m}(\boldsymbol{x}) (\boldsymbol{\mu} - \boldsymbol{\mu}_j)^a. \quad (5)$$

Evidently, the transport field is assumed to be a continuous function of the parametric deviation. In particular, the continuous parametric variation is encoded using polynomial terms in $(\boldsymbol{\mu} - \boldsymbol{\mu}_j)$, with maximum order N_a . Moreover, the entire flow field, including the shock-waves, are also assumed to transport smoothly in space. So, the spatial variation is modeled with N_m basis functions $g_{i,m}(\boldsymbol{x})$, the notation referring to the m^{th} spatial basis function in the i^{th} direction. Nair and Balajewicz [7] propose the use of smooth sine/cosine functions or Chebyshev polynomials; we have exclusively used sine functions in our work, and found them to be adequate. On selecting these quantities, the sole remaining unknown quantities are the coefficients $c_{j,i}^{m,a}$. (The original work [7] proposes an extension to address vector parameter spaces.)

The coefficients for the j^{th} transport field are found by transporting the j^{th} snapshot in the training database such that it is able to optimally reconstruct a set of neighbouring snapshots. Let \mathcal{M}_j denote the set of parameters represented




Fig. 1 (a) Selection of $N_v = 3$ nearest neighbours of μ_5 in the offline stage for a 2-D parameter space. (b) Selection of $N_k = 3$ nearest neighbours of μ^* in the online stage for a 2-D parameter space. Here, $\mu := [\mu^1, \mu^2]^T$.

in the training database that are deemed to constitute the neighbourhood of the j^{th} snapshot in parameter space. A graphical representation of this in a 2-D parameter space is shown in fig. 1a. It has been seen that for problems involving shocks that move with changing parameter, the accuracy of reconstruction decreases with increase in parameter distance [7]. Hence, the cardinality N_v of the \mathcal{M}_j set is typically kept much smaller than N , the size of the training database. Evidently, for multi-dimensional parameter spaces, the proper normalization of the various individual parameters is essential prior to this neighbourhood determination.

With the preceding setup, we can now define the ‘offline stage’ problem. We wish to find the coefficients $c_{j,i}^{m,a}$ of the j^{th} transport field (for a particular snapshot indexed by j) that minimize the square of the \mathcal{L}^2 -norm of the flow field approximation error averaged over all the neighbours, defined as

$$\epsilon_j := \frac{1}{N_v} \sum_{\mu_k \in \mathcal{M}_j} \|\check{q}_j(\mathbf{x}; \mu_k) - \mathbf{q}(\mathbf{x}; \mu_k)\|_2^2, \quad j \in \{1, \dots, N\}. \quad (6)$$

Evidently, this minimization must be pursued independently for each of the N snapshots in the learning database. The actual error minimization problem is subject to two kinds of constraints. The first one restricts the amount of distortion of the cells of the grid in the transportation process. The second one approximately implements the boundary conditions. Both will be discussed in the context of specific problems subsequently. Standard software packages are available to perform such constrained optimization tasks.

B. Online stage

After completing the offline stage, which is a one-time process, the solution at an unsampled parameter value μ^* can be predicted. The approximate solution can be given as a linear combination of the local basis derived by transporting neighbouring snapshots. Mirroring the offline stage, let \mathcal{M}^* denote the set of parameters presented in the training database that are deemed to constitute the neighbourhood of μ^* . The cardinality of this set is denoted by N_k . Figure 1b provides a graphical representation of the selection strategy.

The approximate solution is posed as the linear combination

$$\tilde{q}(\mathbf{x}; \mu^*) = \sum_{l \in \{1, \dots, N_k\}, \mu_l \in \mathcal{M}^*} \eta_l \check{q}_l(\mathbf{x}; \mu^*), \quad (7)$$

where $\check{q}_l(\mathbf{x}; \mu^*)$ represents the transport of the snapshot corresponding to the parameter μ_l to approximate the solution for the new parameter μ^* . The unknown coefficients $\{\eta_l\}_{l=1}^{N_k}$ are found by substituting the above approximate solution in the expression for the residual introduced in section II, and subsequently minimizing its \mathcal{L}^1 -norm over the entire flow domain (or a suitable sub-domain):

$$\min_{\eta_1, \dots, \eta_{N_k}} \|R(\tilde{q}(\cdot; \mu^*))\|_1 + \sigma \|B(\tilde{q}(\cdot; \mu^*))\|_1. \quad (8)$$

Usually, a local basis is small in cardinality, making the optimization problem severely over determined. In particular, it becomes challenging to add boundary constraints. Thus, one should include them only if essential, and then too in the

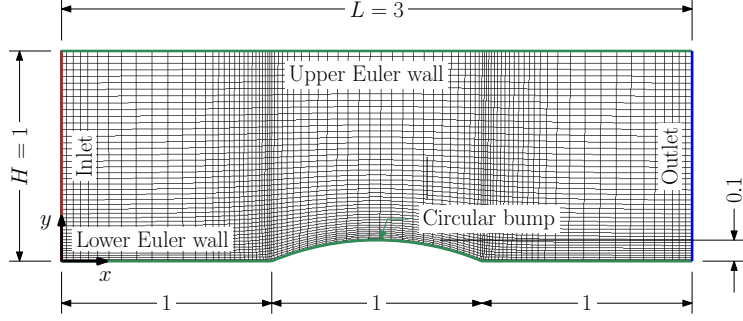


Fig. 2 Geometry of the bump problem as given in the SU2 tutorial. The structured grid topology is coarsened for representation by a factor of 3 from the one actually used (i.e., having 86 points along the flow and 44 across).

form of penalty functions. This is demonstrated above with the penalty weight σ ; the error in boundary condition is evaluated on the boundary of the domain using the \mathcal{L}^1 -norm again.

The basis functions $\{\check{q}_l(\mathbf{x}; \boldsymbol{\mu}^*)\}_{l \in \{1, \dots, N_k\}, \boldsymbol{\mu}_l \in \mathcal{M}^*}$ delivered by the offline stage approximate the flow solution for the new $\boldsymbol{\mu}^*$ already. The purpose of the online stage is to choose the best linear combination of these basis functions so as to minimize the residual. Hence, the paucity of degrees of freedom in the online stage is not a severe shortcoming.

IV. Further developments of TSMOR in the context of the channel-with-bump problem

The steady 2-D inviscid transonic channel-with-bump problem was chosen to validate some further developments of the TSMOR approach, principally addressing boundary condition implementation issues. This problem, introduced in section IV.A, is appropriate because it contains a well-defined shock whose strength, location and orientation are dependent on a single parameter – viz. the inlet Mach number. It also has a wall boundary condition that is simple to implement in TSMOR for a structured grid, but much more difficult for general unstructured grids. Thus, we first explored and firmed up ideas with a structured body-fitted curvilinear grid for this problem, as delineated in section IV.B. The lessons learned therefrom were then applied to the ultimate objective testbench involving an unstructured grid for this problem, as presented in section IV.C. Although the novel developments of TSMOR are discussed exclusively in the context of the channel-with-bump problem, it will be clear towards the end of the discussion that the proposed modifications have much wider applicability.

A. Problem Description

Figure 2 presents the geometry of the 2-D channel with a bump on the lower wall. It is adopted from a tutorial of the open-source CFD software SU2 [12]. The bump’s streamwise extent serves as the characteristic length scale of the problem; i.e., it is considered to be of unit length. The bump is a portion of a circular arc, with the height being 0.1 unit; this makes its radius 1.3 units. The channel, in whose bottom wall the bump is centered, is 3 units long (i.e., $L = 3$) and 1 unit high (i.e., $H = 1$).

We are interested in the steady, inviscid 2-D flow over the bump, which is governed by the Euler equations given by

$$\frac{\partial \mathbf{F}}{\partial x} + \frac{\partial \mathbf{G}}{\partial y} = 0, \quad \mathbf{F} := \begin{bmatrix} \rho u \\ \rho u^2 + p \\ \rho uv \\ (\rho E + p)u \end{bmatrix}, \quad \mathbf{G} := \begin{bmatrix} \rho v \\ \rho uv \\ \rho v^2 + p \\ (\rho E + p)v \end{bmatrix}, \quad (9)$$

where, ρ , p , E , are respectively the density, pressure and total energy, and u and v are the x and y -components of the velocity. The bottom wall (with the bump) and the top wall have the no-penetration wall condition imposed. The stagnation conditions are specified at the inlet, and the outlet has the static pressure specified. The problem is solved using the Euler solver of SU2 [13], adapting the configuration presented in the tutorial.

The inlet stagnation conditions are calculated with isentropic flow relations from a nominal inlet Mach number \hat{M} provided by the user, assuming the static conditions to be atmospheric. However, because of the particularities of flow choking at the bump, the actual inlet Mach number \tilde{M} (as determined by the flow solver) turns out to be lower




Fig. 3 Approximately linear relationship between actual inlet Mach number \tilde{M} and nominal one \hat{M} , in the range of interest.

than the nominal value. Also, the static density and pressure come out to be different from the assumed atmospheric values. Even though it is not achieved at the inlet, we will continue to refer to the nominal inlet Mach number \hat{M} as the variable parameter of this problem (i.e., the applicable scalar μ). Indeed, the training database has six snapshots at \hat{M} of 0.7, 0.73, 0.76, 0.79, 0.82, and 0.85, that correspond to the shock foot moving gradually from the leading corner to the trailing corner of the bump. For this small range of \hat{M} , the corresponding actual inlet Mach numbers (i.e., \tilde{M} 's) appear to be linearly related, as demonstrated in fig. 3; this relation will be exploited later. The procedure is validated at intermediate \hat{M} 's of 0.72, 0.75, 0.78, 0.81 and 0.84.

The flow variables are normalized by the assumed inlet static conditions. Thus, velocity is normalized by the assumed inlet speed of sound, density by the assumed inlet static density, and pressure by assumed inlet static pressure times the specific heat ratio γ (assumed to be 1.4). Since the actual static conditions differ from these assumed ones as discussed above, the normalized inlet density is different from unity and the normalized inlet pressure is different from $1/\gamma$. The flow variable vector for this 2D Euler problem is $\mathbf{q} = [\rho, p, u, v]^T$.

Of relevance in our discussions is a measure of the discrepancy between two flow solutions, say \mathbf{q}_1 and \mathbf{q}_2 . These may represent, for example, an ‘actual’ solution and its TSMOR-approximated counterpart. We measure such discrepancies using the square of the \mathcal{L}^2 -norm of the flow vector field error (which was also used in the offline stage in eqn. (6)):

$$\varepsilon(\mathbf{q}_1, \mathbf{q}_2) := \int_A \|\mathbf{q}_1(\mathbf{x}) - \mathbf{q}_2(\mathbf{x})\|^2 dA. \quad (10)$$

In the above, the Euclidean 2-norm of the flow vector field error is first evaluated in each grid cell. For this, we simply add the squares of the errors in density, pressure and the two velocity components. This is meaningful since ρ , p and \mathbf{u} have each been normalized to be of order unity and dimensionless. (If the flow solution is available on grid nodes, then the solution on a cell is approximated by averaging over its nodal values.) Next, the area of the grid cell is multiplied to this error, and the results are summed over all grid cells (to approximate the area integral). With the above definition, we can rewrite the offline stage cost function of eqn. (6) as

$$\epsilon_j := \frac{1}{N_v} \sum_{\mu_k \in \mathcal{M}_j} \varepsilon(\check{\mathbf{q}}_j(\cdot; \mu_k), \mathbf{q}(\cdot; \mu_k)), \quad j \in \{1, \dots, N\}.$$

B. Implementation of TSMOR model on a structured grid

It is clear from the description of the offline stage in section III.A that the snapshot transportation requires interpolation of the flow field on to the original grid. If the grid is distorted in the primitive $x - y$ Cartesian domain, significant extrapolation will be needed near the bump wall. However, a workaround exists for the simple problem at hand; indeed, the problem was chosen because the workaround is available in its case. We first pursue TSMOR with a structured grid for the bump-in-channel flow problem, and subsequently proceed to an unstructured grid. In particular, the domain is meshed with 256 grid points along the primary flow direction and 128 grid points across it. This is the same mesh as presented in the SU2 tutorial [12]; we have performed grid convergence studies to ascertain its adequacy. A three-times coarsened mesh with the same topology is represented in fig. 2.

Prior to the offline stage, the problem is converted from the original $x - y$ domain to the underlying unit-spacing rectangular Cartesian $x' - y'$ domain. The new length of the domain is $L' = 256$ and $H' = 128$. The bump is flattened out in the transformed coordinates. The issue of extrapolation is resolved by constraining the points on the bump wall




Fig. 4 (a) Orthogonal projection of velocity onto the local y' direction of the underlying structured grid, with the remainder denoted as the x'' component, all in the $x - y$ domain. (b) Grid transport in $x' - y'$ domain, with velocity components preserved at grid points. (c) Interpolation to $\hat{x} - \hat{y}$ domain, $\hat{u}_{y''}$ and $\hat{u}_{x''}$ assumed to be along the transported local y' and x' directions, respectively.

(as well as on the top wall) to move along the wall boundary only; we term this as a ‘sliding boundary’. To effect the sliding boundary, the snapshots must not be transported in y' at $y' = 0$ and $y' = H'$. This can be achieved by restricting the basis functions for the y' -transport field to sine functions with nodes at $y' = 0$ and H' . In fact, the spatial basis functions are chosen as

$$g_1(x', y') \in \left\{ 1, \sin\left(\frac{\pi x'}{L'}\right), \sin\left(\frac{2\pi x'}{L'}\right), \sin\left(\frac{3\pi x'}{L'}\right) \right\}, \quad g_2(x', y') \in \left\{ \sin\left(\frac{\pi y'}{H'}\right), \sin\left(\frac{2\pi y'}{H'}\right), \sin\left(\frac{3\pi y'}{H'}\right) \right\}. \quad (11)$$

The unity basis function allows a bodily movement of the grid in x' , which necessitates extrapolation. But this is straightforward as the original linear inlet and outlet edges remain linear after the transport. Also, notice that there are no cross terms – i.e., y' -dependent functions in x' distortion, and vice versa. This implies that the original rectangular Cartesian $x' - y'$ grid remains rectangular after transportation, although it loses the unit-spacing character. The transported $x' - y'$ grid is linearly interpolated to the $x - y$ domain, and the flow variables are assumed to be transported unchanged to this distorted $x - y$ grid. Finally, another round of interpolation and extrapolation (the latter only at the inlet and outlet) are used to map the flow solution back on to the original $x - y$ grid.

Working in the $x' - y'$ domain and sliding the wall boundaries together ensure that the distorted grid respects the geometry of the walls. However, more work is necessary to maintain the no-penetration Euler wall condition on the bump. For this, we take advantage of an excellent quality – viz. the approximate orthogonality – of the topology of the curvilinear body-fitted structured grid adopted from SU2’s tutorial for this problem; this may be observed in fig. 2 also. So, prior to the offline stage, the flow velocity vector in the $x - y$ domain is resolved into two orthogonal directions such that one component is along the local x' grid line in the $x - y$ domain and the other component is orthogonal to it. This is graphically represented in fig. 4 where $u_{x'}$ is the projection of \mathbf{u} onto the local x' axis, and y'' is the corresponding orthogonal direction (which is almost along the local y' axis). These are the two components of velocities that are transported. Following the interpolation (and extrapolation) on to the original $x - y$ grid (as mentioned in the previous paragraph), these components are used to retrieve the x - and y -components of the velocity, at which stage it is assumed that $\hat{u}_{y''}$ is along the transported local y' direction. On the bump wall, $u_{y''}$ is zero to start with, and it remains zero throughout the snapshot transport process; this ensures the automatic satisfaction of the no-penetration condition in the offline stage. Moreover, since all transported snapshots respect the Euler wall condition, it is automatically enforced in the online stage too.

For the one-dimensional parameter space of this problem, the neighbourhood set cardinality is set to $N_v = 2$ in the online stage (refer to eqn. (6)). That is, the $\hat{M} = 0.7$ snapshot is transported to approximate the $\hat{M} = 0.73$ and 0.76 snapshots, the $\hat{M} = 0.73$ snapshot is transported to approximate the $\hat{M} = 0.7$ and 0.76 snapshots, and so on. This follows the choice shown to be appropriate by Nair and Balajewicz [7] for the three problems that they addressed.

A preliminary study is performed to determine the appropriate order of the polynomial involving parameter deviations, i.e., N_a of eqn. (5). The error in the offline stage is calculated for all snapshots attempting to approximate all their respective neighbouring snapshots, and the average is taken. This study is pursued for increasing values of N_a , and fig. 5 shows that $N_a = 2$ is sufficient, which is adopted for further results presented here.

In the online stage, we again chose the neighbourhood set cardinality N_k to be 2. That is, any new nominal inlet Mach




Fig. 5 Error in offline stage vs. degree N_a of polynomial in $(\mu - \mu_j)$, averaged over all snapshots in the learning database with structured grid.

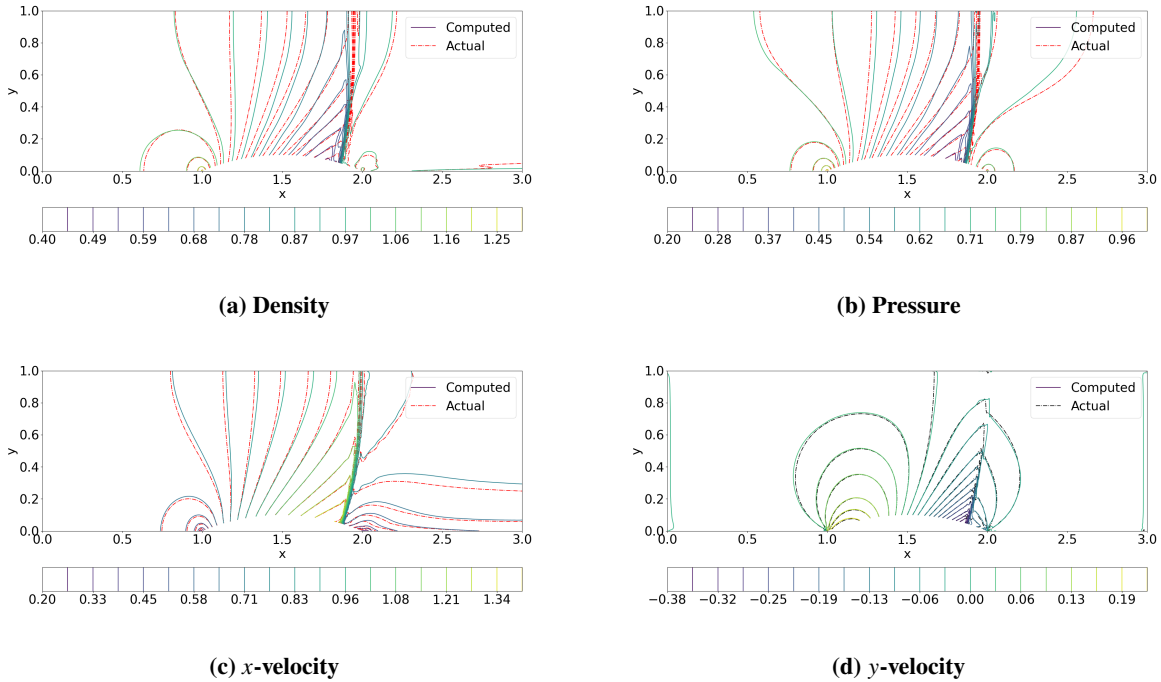


Fig. 6 Flow field from TSMOR ('computed') and Euler solver ('actual') at $\hat{M} = 0.75$, using the structured grid.

number \hat{M}^* is solved for using basis functions derived by transporting its immediately lower and immediately higher \hat{M} neighbours from the learning database. This is again in line with the practice of Nair and Balajewicz [7]. It means that there are only two degrees of freedom in the optimization pursued in the online stage. During the online stage, the inlet boundary condition is enforced by severely penalising the residual with the error in the inlet stagnation conditions. However, the stagnation conditions depend only on the nominal inlet Mach number \hat{M} (which is the only parameter that is specified explicitly) that in turn is linearly related to the actual inlet Mach number \bar{M} . Thus, corresponding to a specified \hat{M}^* , we interpolate the learning database to determine its approximate \bar{M}^* , find the deviation from it of the TSMOR-predicted flow field's inlet Mach number, and use it as the penalty term in the cost function (see eqn. (8)). The lower and upper wall boundary conditions are enforced implicitly by the preceding steps. Finally, since all snapshots have the same outlet boundary condition (static pressure is atmospheric), it is automatically satisfied in the online stage if the linear combination's weight coefficients (i.e., η 's) add up to unity. It was found in practice that the explicit implementation of this constraint is unnecessary – it was automatically satisfied. Strict enforcement would of course have brought the degree of freedom of the problem down to one, which is too restrictive.

Figure 6 presents a comparison of the TSMOR-predicted flow field with its 'actual' counterpart for a nominal inlet Mach number of $\hat{M} = 0.75$, which falls within the range included in the learning database, but does not belong to it. The position, strength and shape of the shock are reproduced well. Importantly, one can remark the absence of the staircase

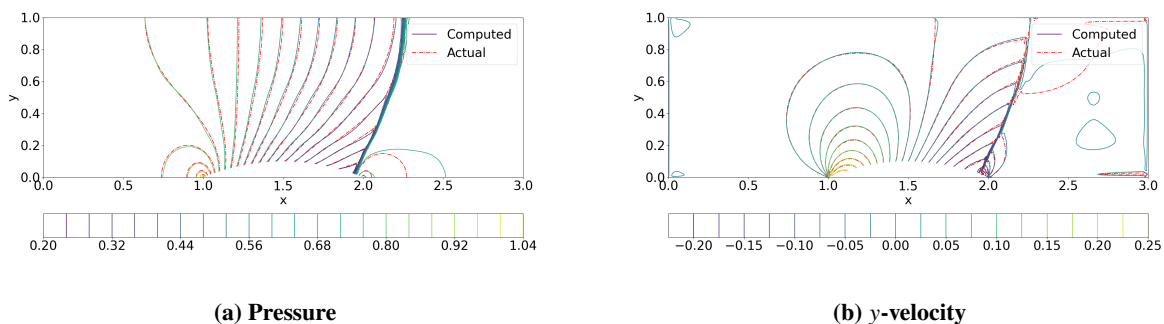


Fig. 7 Selected flow field variables from TSMOR ('computed') and Euler solver ('actual') at $\hat{M} = 0.81$, using the structured grid.




Fig. 8 TSMOR-predicted and actual pressure distributions on the lower wall of the channel at two \hat{M} 's, using the structured grid.

artifact that bedevils linear approaches like POD. Other portions of the flow field are also estimated with reasonable fidelity. This is true for all the four flow variables. Evidently, pressure contours are quite similar to the contours of density as well as x -velocity; so we will only depict contours of pressure and y -velocity in subsequent comparison figures.

Similar results for another nominal inlet Mach number, viz. $\hat{M} = 0.81$ are shown in fig. 7. The agreement between the TSMOR predictions and the 'actual' solution is even better than the $\hat{M} = 0.75$ case. The entire curved shock structure is captured very well. This improved performance is probably an outcome of the cross-direction grid lines aligning particularly well with the shocks in the neighbourhood of this Mach number – i.e., for $\hat{M} \in \{0.79, 0.81, 0.82\}$. For the previous $\hat{M} = 0.75$ case, this kind of serendipitous alignment is probably not attained, thereby resulting in the lack of agreement towards the upper end of the shock wave. This may be ameliorated by introducing cross terms in the grid transport basis set. Besides this, an adaptive sampling strategy [14] can also be used to capture the regions in parameter space where changes are much more rapid. Both these improvements are pursued in the context of the approach involving the unstructured grid that is described now.

Of greatest interest in our studies are the aerodynamic forces on external walls. For this, fig. 8 presents the pressure on the lower wall of the channel, comparing the actual values with the TSMOR-computed ones. We see an almost exact match, with shock strength and location being predicted very accurately.

In the foregoing, we have demonstrated that the TSMOR approach with the proposed novel modifications is able to approximate the flow at new parameter conditions with very agreeable accuracy. This was shown in flow field contour plot overlays and in flow variable distributions on the most relevant lower wall with the bump. Now we pursue a more quantitative validation. In particular, the fidelity of reproduction of the flow field over the entire domain is investigated using the square of the \mathcal{L}^2 -norm of the flow vector field error introduced in eqn. (10). That is, for any test parameter μ^* , we evaluate $\varepsilon(\tilde{q}(\cdot; \mu^*), q(\cdot; \mu^*))$. Table 1 demonstrates that this domain-integrated error is very small for the first four test cases, with the maximum being less than 0.0012. For context, a value of ≈ 0.002 was found in fig. 5 as the average

\hat{M}	0.72	0.75	0.78	0.81	0.84
Domain error $\varepsilon \times 10^3$	0.575	1.173	0.660	0.188	47.303

Table 1 Quantitative comparison of five structured grid test cases in terms of the domain-integrated error metric ε defined in eqn. (10).




Fig. 9 Schematic showing how a point P' that is transported from its original location on a wall is constrained to remain within a distance \hat{d} of the wall in case of an unstructured grid (i.e., $d \leq \hat{d}$).

error of the offline stage. Two things are significant in this comparison. Firstly, the current *online* stage outcome is faring well vis-à-vis the previous *offline* stage result that already has information about the reference snapshots that are being approximated. Furthermore, the *maximum* error in the current context is less than the *average* discrepancy in the analogue. There is a drastic shift in the shock shape on moving from the nominal inlet Mach number $\hat{M} = 0.82$ to 0.85, which introduced large errors in this \hat{M} range; we will address this issue subsequently when using the unstructured grid.

These results motivate the pursuit of the next part of this work, which seeks to adapt the TSMOR method for an unstructured grid; this is the primary objective of the paper.

C. Implementation of TSMOR on unstructured grid

In the implementation of TSMOR on a structured grid (described in section IV.B), the imposition of boundary condition on the walls – be they curved or flat – was relatively straightforward. The ‘sliding’ wall boundary condition could be effected by the suitable choice of the spatial basis functions for the grid transport. This becomes impossible in case of an unstructured grid for an arbitrary geometry. Instead, extra constraints have to be added in the offline stage to prevent large departures of walls normal to themselves during the grid transport. Apart from this, we also describe the inclusion of cross terms in the basis functions for the grid transport. We also discuss the inclusion of cross terms in the spatial basis functions for the grid transport, since most shocks will be curved instead of being aligned with one or the other coordinate direction. In the case of a structured grid, the curvilinear grid could be designed so that shocks were approximately aligned with one of the (curved) grid directions, thereby obviating the need for cross terms. Both these enhancements are discussed now in the context of solving for the flow in the channel-with-bump using an unstructured grid.

In the offline stage, we use the following spatial basis functions for the transport of the unstructured grid:

$$g_1 \in \left\{ 1, \sin\left(\frac{\pi x}{L}\right), \sin\left(\frac{2\pi x}{L}\right), \sin\left(\frac{\pi y}{2H}\right), \sin\left(\frac{\pi y}{H}\right) \right\}, \quad g_2 \in \left\{ 1, \sin\left(\frac{\pi y}{H}\right), \sin\left(\frac{2\pi y}{H}\right) \right\}. \quad (12)$$

There are several changes vis-à-vis the structured grid counterpart in eqn. (11). First of all, the basis functions are in the $x - y$ domain, since there is no underlying rectangular Cartesian $x' - y'$ domain for an unstructured grid. Next, the y -transport now includes a unity basis function to allow bodily movement of the grid in this direction; this was avoided to engineer the sliding wall boundary condition in case of the structured grid, but is essential in the present case. Furthermore, we restrict the basis to two sine terms in the self directions instead of three used in the structured grid case, to keep the degrees of freedom manageable in the offline stage optimization. However, we include two cross terms now – viz. $\sin(ky)$ terms – in the x -transport basis; this choice will be discussed subsequently. Finally, a linear polynomial in parameter difference is chosen for the transport field (i.e., $N_a = 1$ in eqn. (5)), in contrast to a quadratic one chosen for the structured case. This is done to keep the optimization degrees of freedom viable (i.e., eight instead of sixteen), while ensuring minimal sacrifice of fidelity.

To minimize the need for extrapolation near no-penetration walls, one would want to deter excessive transport of the wall grid normal to itself, while allowing tangential transport. To encode this in the offline stage’s optimization

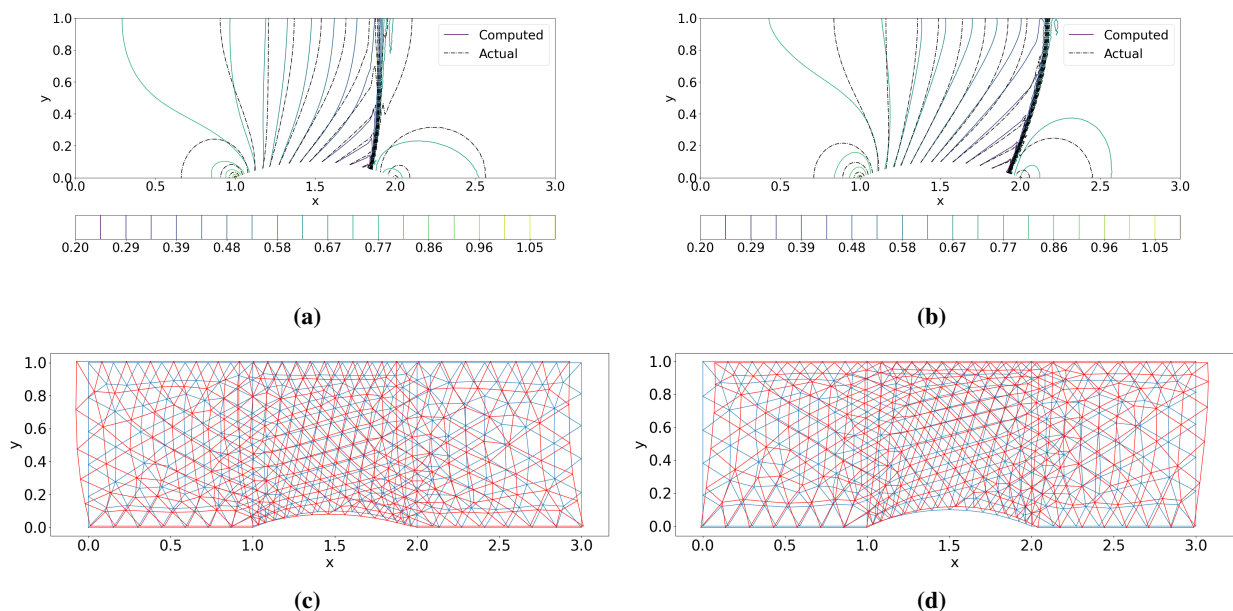


Fig. 10 Results from the offline stage for reference snapshot at $\hat{M} = 0.76$ with an unstructured grid, when it is transported to approximate (a) & (c) the $\hat{M} = 0.73$ flow, and (b) & (d) the $\hat{M} = 0.79$ solution. (a-b) Pressure contours ‘computed’ by transporting the snapshot compared with their ‘actual’ counterparts. (c-d) Baseline unstructured grid (blue) overlaid by the distorted counterparts (red) after transportation, all being 10 times coarser than the production grid.

problem, we impose an additional constraint as depicted in fig. 9. Consider a point P on the boundary wall, and let it be transported to P' . We measure the perpendicular distance d of the transported boundary point P' from the boundary wall, and constrain it to be less than some limit, say \hat{d} . Through trial and error, we set $\hat{d} = 0.0015$; i.e., wall normal movement is limited to .15% of the bump length, which is the length scale of the problem.

The other constraint imposed on the offline stage optimization is to ensure that the grid topology does not change; i.e., grid cells should not shrink too much. To enforce this, we start by ascertaining the signed area of all cells in the undistorted grid. For each candidate transport field, we also find the signed area of the cells. The magnitude of area is constrained to be greater than one-hundredth of the area of the smallest undistorted cell. Moreover, the sign of the area of each cell is constrained to remain same through the distortion. This is a more complicated formulation of its counterpart in the structured grid case, which was propounded in the original TSMOR work by Nair and Balajewicz [7].

The unstructured grid used for the problem had about 24,000 triangular cells and half as many nodes. Along the lower wall, there are 189 edges; the corresponding numbers for the upper wall, inlet and outlet are 171, 35 and 35, respectively. Although the blue grid in fig. 10c is coarsened by a factor of 10 from this production grid, it still gives an idea of the topology. A grid convergence study was performed to ensure that the results are independent of the chosen grid resolution.

Figure 10 presents a representative result from the offline stage, where the $\hat{M} = 0.76$ snapshot is being transported to approximate its two neighbours – viz. $\hat{M} = 0.73$ and 0.79 solutions. The comparison of the pressure field contours demonstrate that the shock is captured well in both transports, in terms of its strength, position and shape (i.e., curvature). The corresponding distortions of the grid are also depicted in fig. 10.

We notice that there is minimal overall movement of the grid at the lower wall, although there is some distortion near the bump that allows the approximation of the movement of the shock foot. However, there is significant x -movement of the grid at the upper wall, to account for larger movement of the shock head across this range of \hat{M} 's. It is the $\sin(0.5\pi y/H)$ cross term in the x -transport basis set of eqn. (12) that affords this y -dependent x -transport of the grid. The other cross term – viz. $\sin(\pi y/H)$ is useful for fine-tuning the distortion in the interior of the domain (i.e., in $0 < y < H$). Indeed, fig. 11 demonstrates the degradation of fidelity in the offline stage if the cross terms are excluded from the basis set.




Fig. 11 Pressure field from the offline stage for reference snapshot at $\hat{M} = 0.76$ with an unstructured grid, when it is transported to approximate the $\hat{M} = 0.73$ flow, but with a basis lacking in cross terms; this should be compared with fig. 10a which has the cross terms.

\hat{M}		0.72	0.75	0.78	0.81	0.83
c_l	Actual	0.687	0.679	0.654	0.624	0.600
	TSMOR	0.680	0.671	0.649	0.621	0.599
	% error	1.14	1.18	0.75	0.63	0.13
c_d	Actual	0.051	0.077	0.100	0.117	0.130
	TSMOR	0.052	0.078	0.100	0.119	0.130
	% error	1.73	0.95	0.39	1.49	0.03
Domain error $\varepsilon \times 10^3$		1.983	0.293	0.575	1.173	0.293

Table 2 Quantitative comparison of five unstructured grid test cases in terms of the lift and drag coefficients for the bump, as well as the domain-integrated error metric ε defined in eqn. (10).

It was remarked in the context of the structured grid that the shock shape changes significantly between $\hat{M} = 0.82$ to 0.85, which introduces large errors in the offline stage error in this \hat{M} range. A manual variant of adaptive sampling [14] was used to overcome this. Namely, another snapshot was introduced at $\hat{M} = 0.84$ in the learning database. This halved the corresponding offline stage error.

The online stage remains unchanged from the structured grid case. Its performance is demonstrated in fig. 12. In particular, as for the structured grid case, we approximate the flow solution for $\hat{M} = 0.75$ by transporting the $\hat{M} = 0.73$ and 0.76 snapshots. Similarly, the $\hat{M} = 0.81$ flow is predicted by transporting the $\hat{M} = 0.79$ and 0.84 instances. The pressure and y -velocity fields show excellent match with the actual solutions. In particular, the slope of the contours on the bump are reproduced very well, which validates our implementation of the wall boundary condition.

As in the structured grid case, we demonstrate in fig. 13 that the pressure distribution on the lower wall, and on the bump in particular, is predicted very well by the TSMOR approach.

Just as in the context of the structured grid, we also pursue a quantitative validation exercise now, only more extensive this time. The lift and drag forces (per unit span) on the bump are appropriate stand-in metrics for the aerodynamic quantities of interest in these kinds of problems. These are computed by integrating the pressure component of the TSMOR solution, and compared with the reference ‘actual’ flow field. For this, we assume that the other side of the wall – the one which is not exposed to the flow – is at the nominal inlet static pressure (which is atmospheric). Moreover, the forces per unit span are normalized by the dynamic head computed from the nominal inlet flow condition and the ‘chord’ length of bump (which is the length scale of the problem). The results are presented in table 2. The test \hat{M} ’s are chosen to be in the interstices of the learning database cases, including the $\hat{M} = 0.75$ and 0.81 cases that have been discussed exclusively till now. It is observed that both the lift and drag coefficients on the bump are predicted with appreciable accuracy, with the error being less than 3% across all cases tested.

Although lift and drag on the bump are indeed the quantities of greatest interest, the fidelity of reproduction of the flow field over the entire domain is also relevant. As in table 1, this is investigated using the square of the \mathcal{L}^2 -norm of the flow vector field error introduced in eqn. (10). Table 2 demonstrates that this domain-integrated error is also very

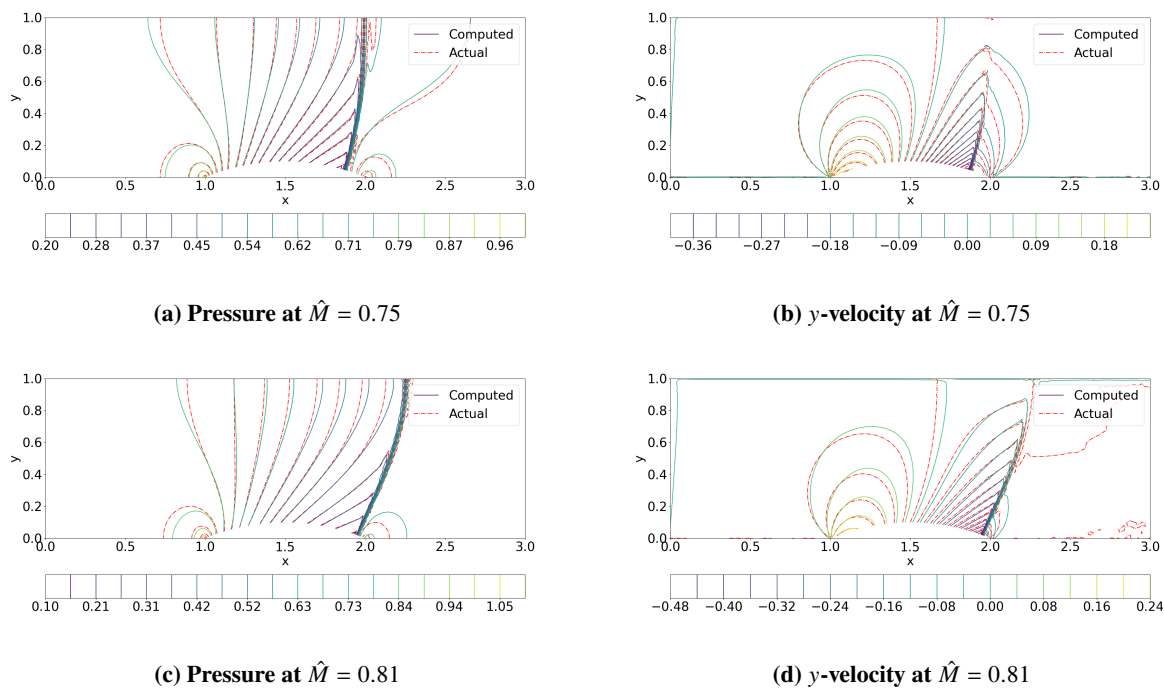


Fig. 12 Selected flow field variables from TSMOR ('computed') and Euler solver ('actual') at two values of nominal inlet Mach no. \hat{M} , using the unstructured grid.

small for all the five test cases, with the maximum being less than 0.002. This is slightly higher than in the structured grid case, but nevertheless very much acceptable. Moreover, the crude adaptive sampling strategy is also working well apparently, and we see low error in the final \hat{M} case. Evidently, the TSMOR approach with its implemented modifications is able to predict the flow field over the entire domain with acceptable fidelity.

V. Conclusion

Model order reduction (MOR) for shock-dominated flows has been an active area of research and several discontinuity-aware MOR methods have been developed. In this paper we extend and expand upon one such method, the transported snapshot model order reduction (TSMOR) approach proposed recently by Nair and Balajewicz [7]. The method was demonstrated to outperform other MOR techniques in its class, when applied to three benchmark problems. However, all these problems relied on Cartesian grids. Here, we augment TSMOR to make it more amenable for curvilinear structured grids as well as unstructured grids. Further, we propose a 'sliding boundary' concept to better preserve the wall boundary condition that was not enforced strictly in the reference paper.

This augmented TSMOR is applied on a two-dimensional steady transonic inviscid channel-with-bump problem, where the sole variable parameter is the inlet Mach number. Initially, a body-fitted structured curvilinear grid is used, since the no-penetration wall boundary condition is simpler to enforce in TSMOR in this context. The insights derived from this first exercise are leveraged to apply TSMOR to an unstructured grid for this same problem. Very encouraging agreement is found in qualitative and quantitative comparisons of the extended TSMOR method's predictions with those from a full order model (i.e., inviscid CFD), for both classes of grids.

The channel-with-bump problem is simple, but it is selected here to act as a stepping stone towards development of MOR approaches for shock-dominated flows over more complicated geometries. One such problem of more practical relevance is the flow over an airfoil. It is envisaged that the augmentations made to TSMOR for enforcing the wall boundary condition should carry over in a straightforward manner to this problem. Just as in the present problem, working with a structured grid (which is not uncommon in airfoil problems) should render the adaptation of TSMOR particularly simple. The developments made towards applying TSMOR to unstructured grids may also be extended to the airfoil problem with some more work. Apart from this, the method will need to be extended to work with multiple




Fig. 13 TSMOR-predicted and actual pressure distributions on the lower wall of the channel at two values of nominal inlet Mach no. \hat{M} , using the unstructured grid.

parameters (viz. freestream Mach number and angle of attack). Further improvements may involve an automated dynamic sampling method to select the most optimum training database which would help in delivering ideal results along with reduction in number of training snapshots. The ultimate objective of this line of research is to enhance TSMOR to work with more complex geometries in three-dimensional problems.

References

- [1] Taira, K., Brunton, S. L., Dawson, S. T. M., Rowley, C. W., Colonius, T., McKeon, B. J., Schmidt, O. T., Gordeyev, S., Theofilis, V., and Ukeiley, L. S., “Modal analysis of fluid flows: An overview,” *AIAA Journal*, Vol. 55, No. 12, 2017, pp. 4013–4041.
- [2] Yano, M., “Model reduction in computational aerodynamics,” *Model order reduction. Volume 3: Applications*, edited by P. Benner, S. Grivet-Talocia, A. Quarteroni, G. Rozza, W. Schilders, and L. M. Silveira, De Gruyter, 2021, Chap. 6, pp. 201–236.
- [3] Lucia, D. J., King, P. I., Beran, P. S., and Oxley, M. E., “Reduced order modelling for a one-dimensional nozzle flow with moving shocks,” *15th AIAA Computational Fluid Dynamics Conference, Paper 2602*, 2001.
- [4] Alonso, D., Vega, J. M., and Velazquez, A., “Reduced order model for viscous aerodynamic flow past an airfoil,” *AIAA Journal*, Vol. 48, No. 9, 2010, pp. 1946–1958.
- [5] Carlberg, K., Choi, Y., and Sargsyan, S., “Conservative model reduction for finite-volume models,” *Journal of Computational Physics*, Vol. 371, 2018, pp. 280–314.
- [6] Abgrall, R., and Crisovan, R., “Model reduction using L^1 -norm minimization as an application to nonlinear hyperbolic problems,” *International Journal for Numerical Methods in Fluids*, Vol. 87, No. 12, 2018, pp. 628–651.
- [7] Nair, N. J., and Balajewicz, M., “Transported snapshot model order reduction approach for parametric, steady-state fluid flows containing parameter-dependent shocks,” *International Journal for Numerical Methods in Engineering*, Vol. 117, No. 12, 2019, pp. 1234–1262.
- [8] Lumley, J. L., “The structure of inhomogeneous turbulent flows,” *Atm. Turb. and Radio Wave Prop.*, edited by A. M. Yaglom and V. I. Tatarsky, Nauka, Moscow, 1967, pp. 166–178.
- [9] Boyaval, S., Bris, C., Lelièvre, T., Maday, Y., Nguyen, N., and Patera, A., “Reduced basis techniques for stochastic problems,” *Archives of Computational Methods in Engineering*, Vol. 17, 2010, pp. 435–454.
- [10] LeGresley, P., and Alonso, J. J., “Investigation of non-linear projection for POD based reduced order models for aerodynamics,” *39th AIAA Aerospace Sciences Meeting, Paper 926*, 2001.
- [11] Abgrall, R., Amsallem, D., and Crisovan, R., “Robust model reduction by L^1 -norm minimization and approximation via dictionaries: application to nonlinear hyperbolic problems,” *Advanced Modeling and Simulation in Engineering Sciences*, Vol. 3, No. 1, 2016, pp. 1–16.
- [12] “Inviscid bump in channel – an SU2 tutorial problem,” https://su2code.github.io/tutorials/Inviscid_Bump, n.d.

- [13] Palacios, F., Colonno, M. R., Aranake, A. C., Campos, A., Copeland, S. R., Economon, T. D., Lonkar, A. K., Lukaczyk, T. W., Taylor, T. W. R., and Alonso, J. J., “Stanford University Unstructured (SU2): an open-source integrated computational environment for multi-physics simulation and design,” *51st AIAA Aerospace Sciences Meeting, Paper 287*, 2013.
- [14] Paul-Dubois-Taine, A., and Amsallem, D., “An adaptive and efficient greedy procedure for the optimal training of parametric reduced-order models,” *International Journal for Numerical Methods in Engineering*, Vol. 102, No. 5, 2015, pp. 1262–1292.