

RECONSTRUCTION METHODS IN ADJOINT-BASED ERROR ESTIMATION

Thomas P. Hunter*, Arnish Sitaram* and Steven Hulshoff*

* Delft University of Technology
Faculty of Aerospace Engineering
Kluyverweg 2, 2629 HS Delft, The Netherlands
e-mail: t.p.hunter@tudelft.nl

Key words: Super-Resolution Neural Networks, Proper Orthogonal Decomposition, Adjoint Method, Output Error Estimation, Goal-Oriented Mesh Adaptation

Summary. Numerical simulations and optimisation methods, such as mesh adaptation, rely on the accurate and inexpensive use of error estimation methods. Adjoint-based error estimation is the most accurate method, and generally the most costly. A strong contributor to this cost is the need to compute a higher resolution adjoint solution, using time dependent information. Here, reconstruction methods applied to the primal and adjoint solutions are proposed to alleviate both the storage footprint of the primal problem and the adjoint computational cost. The method is compared to reference error estimators on an unsteady Burgers' equation using the method of manufactured solutions. Two reconstruction methods, a proper orthogonal decomposition and a static convolutional neural network were used to demonstrate both the computational cost reduction and the potential for the reduction of the storage footprint of the primal problem. When reconstruction methods are applied to the primal problem, one can use both proposed approaches to reduce the footprint of the solution and reconstruct the effectively compressed primal solution to be recalled for the adjoint solution. The second approach consists in solving an adjoint solution from a coarse primal solution and using reconstruction methods to obtain a higher resolution adjoint solution, necessary for output error estimation and mesh adaptation. The obtained results give great confidence in the use of reconstruction methods for the reduction of both computational cost and storage requirements of adjoint-based error estimation, and goal-oriented mesh adaptation.

1 INTRODUCTION

Computational Fluid Dynamics (CFD) has become a vital tool of the modern engineer in the analysis of fluid flows. The increase in availability of computational power and corresponding decrease in computational cost of numerical simulations in the past decades has enabled the widespread use of CFD from the aerospace industry to healthcare [1]. However, the growing need for ever more accurate and complex fluid dynamic analyses has meant that the use of CFD has remained very specialised and user-intensive. This

is emphasised by the large investment required to conduct design and analysis tasks on a regular basis. In particular, large transient turbulent flow problems, such as the ones analysed using Large Eddy Simulation (LES), are affected by prohibitively large computational costs limiting their use as a design tool.

Numerical simulations require the generation of a finite-dimensional solution space that is used to solve fluid equations specific to the problem, for example the Navier-Stokes equations. The number of elements that are used to discretise a computational domain can be seen as a design variable. An increase in the number of elements constituting this computational domain increases the accuracy of modeling of the phenomena of interest. However, this also results in an increase in the cost required to solve the simulation. Moreover, the elements in the computational domain do not contribute equally to the inaccuracy of a simulation in the prediction of a chosen Quantity of Interest (QoI). Thus, the computational cost could be reduced by refining the element distribution only in the regions where it will help the simulation to become more accurate.

Mesh adaptation aims to tackle this issue by maximising the accuracy of a numerical solution for a given computational cost. The accuracy is here defined as the error in the QoI, also referred to as output error. In the domain of fluid dynamics, this could be lift, drag, or a combination of the two. The maximisation of the accuracy can thus be described as an optimisation, which can be resolved using goal-oriented mesh adaptation. A key component of goal-oriented mesh adaptation is the identification of the local contributions of each element to the error in the discrete solution of the QoI. These are referred to as error indicators. The sum of the error indicators then results in the error in QoI, referred to as the output error. The identification of error indicators thus enables the determination of the optimal number and size of elements to minimise the output error for a given computational effort.

The computational overhead required for the use of adjoint-based error estimation stems from the cost of the adjoint solution and the amount of information required to be stored for the solution of unsteady adjoint problems. On the one hand, the former is the main issue for steady problems and reflects the need to solve the adjoint problem on a refined finite dimensional solution space with respect to the solution of the fluid equations. While in practice this is avoided through the use of prolongation operators to partially reconstruct the fine solution [2], this reduces the accuracy of the error estimation. Therefore in this paper, the proposed method is compared to the adjoint-based error estimate obtained from the refined space.

On the other hand, the amount of storage required for the solution of the unsteady adjoint problem is a direct result of the adjoint problem representing the sensitivity of a problem. Just as the solution at a given time-step depends on the preceding time-step, the adjoint solution must be solved backwards in time in order to determine the sensitivity of a solution to the discretisation of the computational domain. As a result, the instantaneous primal solutions must be stored during the solution of the CFD simulation, and recalled during the solution of the backwards in time adjoint problem. This leads to extremely large storage requirements, with current methods writing the primal to disk [3] and loading the instantaneous primal at each time-step, slowing down the solution of

the adjoint. Checkpointing can be used to rely solely on available memory but greatly increases the complexity of the solution process [4]. The issue related to the storage footprint is particularly relevant in the application of adjoint-based error estimation to LES. LES solutions are characterised by a large range of spatial scales, represented by a reference length L , necessitating high-dimensional meshes, described by a mesh spacing h , ($L/h \gg 1$). Moreover, the physical phenomena of interests are described by a time T corresponding to a large number of time steps δt ($T/\delta t \gg 1$). This leads to intractable storage requirements.

2 PROPOSED METHOD

2.1 Output Error Estimation

We consider the unsteady one-dimensional (1D+t) Burgers' equation over the domain that spans the interval $\Omega : [0, 1] \times I : [0, T]$ in space and time respectively, as in [5].

$$N(u) := \frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} - \nu \frac{\partial^2 u}{\partial x^2} = f \quad (1)$$

where $N(u)$ is a non-linear operator and u is the velocity field with boundary conditions $u(0, t) = u(1, t) = 0$ and initial condition $u(x, 0) = u_0$. The kinematic viscosity ν and the forcing term f are chosen *a priori*. In this study, they are chosen to be $Re = 1/\nu = 100$ and the forcing term corresponds to a manufactured solution $u = \sin(\pi x) \sin^2(\pi t)$ [5].

The derivation of the continuous adjoint equation starts with the definition of the continuous residual $r(u)$, which is equal to zero when the primal Partial Differential Equation (PDE) is satisfied: $r(u) = Lu - f$.

The adjoint operator L^* is defined as the relation between the primal operator and an adjoint set of solutions Ψ through the relation $(Lu, \Psi) = (L^*\Psi, u)$. The adjoint solution can then be defined as the sensitivity of a chosen output $J(u)$ with respect to an infinitesimal perturbation in the primal residual requiring:

$$J'(\delta u) = \int_{\Omega} \Psi^T r'(\delta u) d\Omega \quad , \quad \forall \delta u \quad (2)$$

Equation 2 is referred to as the generalised form of the continuous adjoint equation. Deriving the primal residual sensitivity $r'(\delta u)$, then substituting it into Equation 2 and performing integration by parts of each term of the primal PDE enables the derivation of the adjoint operator L^* . The boundary conditions are then determined by associating the various known boundary conditions from the primal problem to the boundary condition terms obtained through integration by parts.

The adjoint equation, about the linearised state u_0 , can be shown to be:

$$-\frac{\partial \Psi}{\partial t} - u_0 \frac{\partial \Psi}{\partial x} + \Psi \frac{\partial u_0}{\partial x} - \nu \frac{\partial^2 \Psi}{\partial x^2} = \frac{dJ}{du_0} \quad (3)$$

The output error can then computed from the adjoint solution on a refined solution

space using the adjoint-weighted residual method:

$$\delta J \approx \delta J_{est} = - \int_{\Omega} \Psi_h^T r_h(u_h^H) d\Omega \quad (4)$$

The local contribution to the output error can be obtained by performing the local integral over the solution space H .

Limitations of this approach The main issues related to the computational and storage overhead required for the adoption of adjoint-based error estimation in mesh adaptation can be traced to the derived set of equations. In Equation 3, one can see that both primal values and derivatives need to be stored in order to allow for the solution of the adjoint problem. This issue is compounded with the potentially wide range of spatial and temporal scales of the problem considered.

Moreover, the solution of the adjoint problem can lead to a significant computational cost overhead due to the requirement for a refined adjoint solution, as shown in Equation 4. This means that a design choice must be made, to either approximate the refined solution, compute two resolutions of primal solutions (one for the residual term and one for the adjoint solution), or use a constrained subspace of the primal solution for the adjoint solution. In this paper, the preferred approach consists in computing the primal solution on both H and h . The former is then used to compute the primal residual term $r_h(u_h^H)$ and the latter is used to solve for the refined adjoint solution Ψ_h .

2.2 Primal Storage Footprint Reduction

The first proposed method aimed strictly at reducing the storage footprint of the primal solution is the use of reduced order representations of the primal solution. Specifically, POD is proposed due to its more widespread adoption in a variety of fields. Whereas its solution can be ill-posed [6] and performance sub-optimal for solutions presenting discontinuities such as shockwaves [6], it is readily implementable due to the number of existing subroutines and libraries.

The approach consists in taking snapshots of the solution and applying a SVD analysis to the solution matrix. This analysis gives the POD modes ϕ and the coefficients α resulting in the expression for the reduced order solution:

$$u(x, t_i) = \bar{u}(x) + \sum_{j=1}^M \alpha_j(t_i) \phi^{(j)}(x) \quad (5)$$

Where \bar{u} is the mean solution, and $\phi^{(j)}$ represent the POD modes, of which M are selected based on a criterion of 99% of energy being reconstructed. Due to the nature of the chosen manufactured solution, a single mode is sufficient in respecting this criterion.

The reduction in the footprint of the primal solution can then be computed based on the number of modes required to represent accurately the primal solution. This results in the compression of the primal being proportional to the number of points in the primal solution.

2.3 Adjoint Computational Cost Reduction

The second proposed method consists in reconstructing of the fine adjoint solution Ψ_h with an ANN from a coarse adjoint solution computed with the Finite Element Method (FEM). The present analysis seeks to explore two aspects of the proposed method. The first, using a SRNN of upscaling factor 2 aims to show the potential computational cost reduction associated with solving the adjoint problem on a coarse discrete solution space. It is referred to as $2 \times CNN$ since it uses the fully convolutional architecture as in [7].

The second expands on this idea by subsampling the primal solution u_H onto a constrained solution subspace u_{2H} and upscaling by a factor 4 ($4 \times CNN$) the solution of the adjoint Ψ_{2H} onto the discrete solution space characterised by the mesh spacing $h = H/2$, as shown in Figure 1. This reduces by a factor 2 the storage footprint of the primal solution required for the solution of the transient adjoint problem. In three dimensional problems, the storage would be correspondingly reduced by a factor of 8.

The second expands on this idea by subsampling the primal solution u_H onto a constrained solution subspace u_{2H} and upscaling by a factor 4 ($4 \times CNN$) the solution of the adjoint Ψ_{2H} onto the discrete solution space characterised by the mesh spacing $h = H/2$, as shown in Figure 1. This reduces by a factor 2 the storage footprint of the primal solution required for the solution of the transient adjoint problem. In three dimensional problems, the storage would be correspondingly reduced by a factor of 8.




Figure 1: Spatial resolution definitions and proposed method

3 RESULTS

3.1 Primal solution

As discussed previously, due to the nature of the chosen manufactured solution $u = \sin(\pi x)\sin^2(\pi t)$, the truncated POD can reconstruct the exact primal solution with a single mode.

3.2 Adjoint Solution

The reconstruction of the adjoint solutions are shown here. These are shown for the same starting resolution $H = 16$ for the $2 \times CNN$ and $2H = 16$ for the $4 \times CNN$. The obtained solutions for the adjoint solutions using the truncated POD of the primal is also shown and has excellent agreement with the one using the full primal solution.




Figure 2: Original adjoint solution (top), super-resolved adjoint solution using a $2 \times CNN$ (middle), Adjoint solution based on the POD reconstructed primal (bottom) at a resolution $H = 16$




Figure 3: Original adjoint solution on top and super-resolved adjoint solution using a $4 \times CNN$ at a resolution $H = 16$

3.3 Error Estimation

However, the accurate adjoint reconstruction is not the sole criterion for the accurate output error computation. The performance of the POD reconstruction and adjoint super-resolution is shown here. As can be seen in those figures, the output error is well reconstructed for all proposed methods.

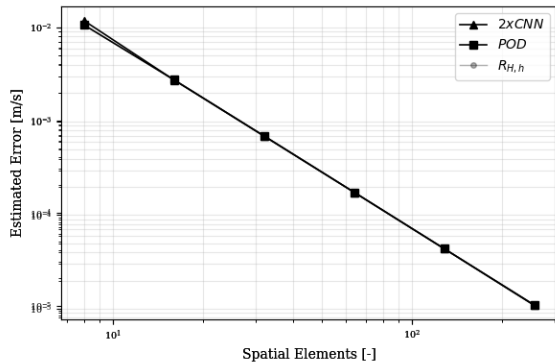


Figure 4: Output error estimate for the primal truncated POD and $2 \times CNN$ versus uniform spatial refinement H

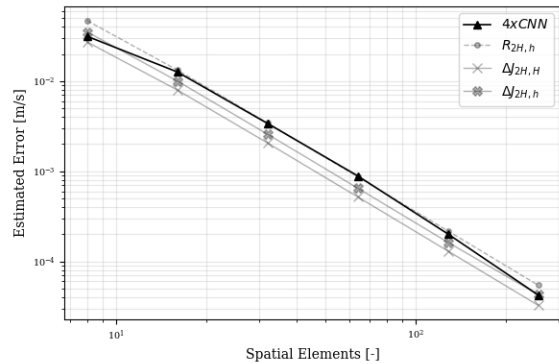


Figure 5: Output error estimate for $4 \times CNN$ versus uniform spatial refinement H

3.4 Error Indicators

Similarly to the preceding subsections, the error indicators are well reconstructed for both the truncated POD of the primal and the super-resolved adjoint. This is very en-

couraging since this would indicate that both reconstruction methods are able to identify error contributing elements. In the context of mesh adaptation this would enable the proposed methods to refine the correct elements to minimise the output error.

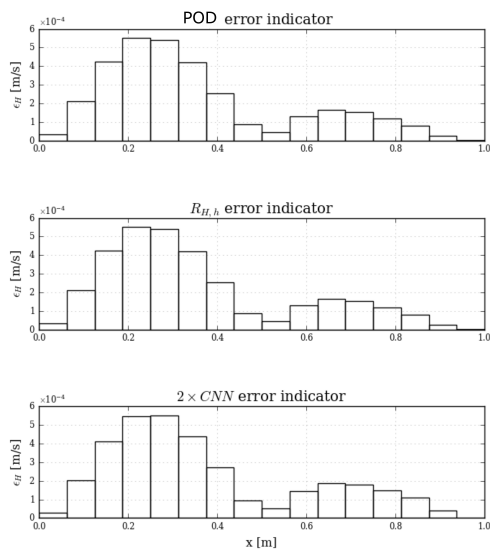


Figure 6: Error indicators for the primal truncated POD (top), reference error estimator (middle), and the $2 \times CNN$ (bottom)

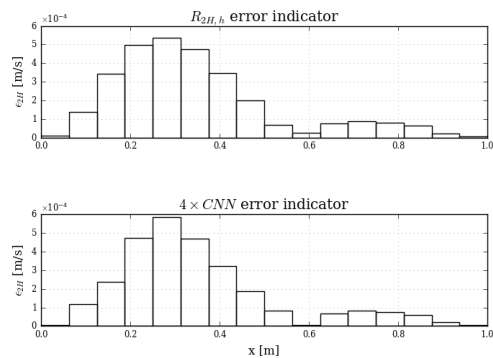


Figure 7: Error indicators for the reference error estimator (top), and the $4 \times CNN$ (bottom)

4 CONCLUSION

This paper was able to highlight the benefits of reconstruction methods applied to the computation of the primal and the adjoint solutions for output error prediction. On the one hand, in the case of the primal, the main issue remains the storage footprint of the time resolved primal solution for the solution of the unsteady adjoint equation. On the other hand, the reconstruction applied to the adjoint solution enables the reduction of the storage requirements and the computational cost of the adjoint solution. This is enabled by the storage of a sub-sampled primal solution and the computation of a coarse adjoint solution. The adjoint solution can then be super-resolved in order to enable the computation of the output error. Both methods were applied to a manufactured solution of the unsteady Burgers' equation successfully.

The primal reconstruction was implemented using a truncated POD of the fine primal solution. For this manufactured solution, this enabled a compression ratio scaling with the number of elements.

The adjoint super-resolution halved the storage footprint of the primal solution (in 1D, i.e. 8x in 3D space) while also allowing for the computational cost to be greatly reduced for the solution of the adjoint solution.

5 CREDITS

Thomas Hunter is responsible for the work related to the adjoint super-resolution from (sub-sampled) primal solutions, the implementation of the truncated POD, and this document. He supervised, alongside Steven Hulshoff, Arnish Sitaram during his Master Thesis at the TU Delft working on the storage footprint reduction and reconstruction of primal solutions using POD, applied to a different framework.

REFERENCES

- [1] Steven L Brunton, Bernd R Noack, and Petros Koumoutsakos. Machine Learning for Fluid Mechanics. *Annual Review of Fluid Mechanics*, 52(1):477–508, jan 2020.
- [2] David A. Venditti and David L. Darmofal. Anisotropic grid adaptation for functional outputs: application to two-dimensional viscous flows. *Journal of Computational Physics*, 187(1):22–46, may 2003.
- [3] Jeffrey Slotnick, Abdollah Khodadoust, Juan Alonso, William Gropp, and Dimitri Mavriplis. CFD Vision 2030 Study: A Path to Revolutionary Computational Aero-sciences. Technical report, 2014.
- [4] Steven M. Kast. An Introduction to Adjoint and Output Error Estimation in Computational Fluid Dynamics. dec 2017.
- [5] Xiaodong Li, Steven Hulshoff, and Stefan Hickel. Towards adjoint-based mesh refinement for Large Eddy Simulation using reduced-order primal solutions: Preliminary 1D Burgers study. *Computer Methods in Applied Mechanics and Engineering*, 379:113733, jun 2021.
- [6] Xiaodong Li, Steven Hulshoff, and Stefan Hickel. Mesh Adaptation Using Adjoint Methods and Reduced-Order Models for Large Eddy Simulation. In *14th WCCM-ECCOMAS Congress*. CIMNE, 2021.
- [7] Yuanyuan Liu, Xiaohui Yuan, Xi Gong, Zhong Xie, Fang Fang, and Zhongwen Luo. Conditional convolution neural network enhanced random forest for facial expression recognition. *Pattern Recognition*, 84:251–261, dec 2018.