

SANDIA REPORT

SAND2016-10669

Unlimited Release

September 2016

Advanced Computational Methods for Thermal Radiative Heat Transfer

John Tencer, Kevin Carlberg, Marvin Larsen, Roy Hogan

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-mission laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: <http://www.osti.gov/scitech>

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: <http://www.ntis.gov/search>



SAND2016-10669
Unlimited Release
Printed October 2016

Advanced Computational Methods for Thermal Radiative Heat Transfer

John Tencer, Kevin Carlberg, Marvin Larsen, Roy Hogan
Thermal Sciences and Engineering
Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185-MS0840

Abstract

Participating media radiation (PMR) in weapon safety calculations for abnormal thermal environments are too costly to do routinely. This cost may be substantially reduced by applying reduced order modeling (ROM) techniques. The application of ROM to PMR is a new and unique approach for this class of problems. This approach was investigated by the authors and shown to provide significant reductions in the computational expense associated with typical PMR simulations. Once this technology is migrated into production heat transfer analysis codes this capability will enable the routine use of PMR heat transfer in higher-fidelity simulations of weapon response in fire environments.

ACKNOWLEDGMENTS

Supported by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. This document has been reviewed and approved for unclassified, unlimited release under 2016-XXXX

CONTENTS

1. Introduction.....	9
1.1. Participating Media Radiation	9
1.1.1. Discrete Ordinates Method	10
2. Model Reduction Strategy	11
2.1. Sampling Methods	13
3. Results.....	17
3.1. 1-D Example	17
3.2. 2-D Examples.....	19
3.3. 3-D Examples.....	22
4. Error Estimation.....	28
4.1. Decoupled	32
4.2. Decoupled with Variable Truncation.....	34
4.3. Concurrent Training Sets	35
4.4. Concurrent Training Sets with Variable Truncation.....	37
4.5. k-fold Cross Validation.....	40
4.6. k-fold Cross Validation with Variable Truncation	42
4.7. Fixed Dimension Greedy Search	46
4.8. Lagged Dimension Greedy Search	47
4.9. Constant Energy Greedy Search	49
5. Conclusions.....	52
6. References.....	54
Distribution	58

FIGURES

Figure 1. Decay of singular values demonstrates the presence of a reduced order basis for the intensity solution for several source distributions. (a) 20 modes are sufficient to capture all behavior in the 1D geometry. (b) 50 modes would capture a very large percentage of the solution behavior in the 2D geometry.....	12
Figure 2. Distribution of sample points generated through greedy search algorithm for a radial temperature distribution $T(r)=300+400r^2$ colored by the ROM residual. Circles correspond to points in the high-order quadrature. Filled circles correspond to points used to train the ROM. 15	
Figure 3. Convergence of angle-integrated intensity distributions for quadratic temperature profile with increasing LOM quadrature order.....	18
Figure 4. Convergence of angle-integrated intensity distributions for quadratic temperature profile with greedy sampling method.	19
Figure 5. Normalized heat flux predictions $q''(x,1)$ for (a) a discontinuous temperature profile and (b) a linear temperature profile for the 14 th order PNTN quadrature, 32 nd order PNTN quadrature, and the ROM derived from the 14 th order quadrature solutions but evaluated at the 32 nd order quadrature points.....	20

Figure 6. (a) Convergence of LOM and corresponding ROM for a temperature distribution of $300+700(1-y)$ for various spatial mesh resolutions (b) Reduction in error provided by the ROM relative to the low-order model for various spatial mesh resolutions. This is the L2 error of the ROM solution divided by the L2 error of the low-order model solution.....	21
Figure 7. Decay of singular values demonstrates that many additional modes are required to capture the behavior of the fully 3D problem.	22
Figure 8. Accuracy as a function of cumulative solution time for a discontinuous temperature distribution $T(x,y,z)=100$ if $y=0$, 0 else.	23
Figure 9. Accuracy as a function of cumulative solution time for a linear temperature distribution, $T(x,y,z)=300+700(1-y)$	24
Figure 10. Accuracy as a function of cumulative solution time for a radially quadratic temperature distribution $T(x,y,z) = 300+400(x^2+y^2+z^2)$	25
Figure 11. Example of construction of ROM error estimate from training data acquired during the greedy search algorithm for a linear temperature distribution, $T(x,y,z)=300+700(1-y)$	28
Figure 12. Convergence of adaptive ROM and associated error estimate of an angularly integrated quantity for a linear temperature distribution, $T(x,y,z)=300+700(1-y)$	29
Figure 13. Comparison of deterministic and Gaussian process estimates of the L2 error in the angular intensity.	30
Figure 14. (a) Observed frequencies from above table (b) Empirical CDFs of validation data normalized by GP mean and standard deviation compared with normal CDF.....	32
Figure 15. Empirical CDF of expanded validation data normalized by GP mean and standard deviation compared with normal CDF for prediction ROM of basis 75 trained generated with first 100 FOM solutions.	33
Figure 16. Variation of Kolmogorov-Smirnov Statistic with number of ROM training points and ROM basis dimension.	34
Figure 17. ROMES training couples generated from training ROMs of varying dimension compared with validation data for prediction ROM of dimension 75.	35
Figure 18. Empirical CDFs of validation data for ROMES generated using concurrent training sets normalized by the GP mean and standard deviation.	36
Figure 19. Variation of Kolmogorov-Smirnov statistic with ROM basis dimension for ROMES generated using concurrent training sets.	37
Figure 20. ROMES generated using concurrent training sets with variable truncation along with the ROMES training couples used.	38
Figure 21. ROMES generated using concurrent training sets with variable truncation along with the ROMES training couples used.	39
Figure 22. (a) Empirical CDFs and (b) Kolmogorov-Smirnov statistics for ROMES generated using concurrent training sets with variable truncation normalized by the GP mean and standard deviation.	40
Figure 23. Empirical CDFs of validation data for ROMES generated using LOOCV normalized by the GP mean and standard deviation.	41
Figure 24. Kolmogorov-Smirnov statistics for ROMES generated using k-fold cross validation normalized by the GP mean and standard deviation.	42
Figure 25. Kolmogorov-Smirnov statistics for ROMES generated using k-fold cross validation with variable truncation with 200 FOM evaluations normalized by the GP mean and standard deviation.	43

Figure 26. Kolmogorov-Smirnov statistics for ROMES generated using k-fold cross validation with variable truncation with 50 FOM evaluations normalized by the GP mean and standard deviation.....	44
Figure 27. Training couples compared with validation data shows good overall agreement for sufficiently large training ROM dimensions.....	45
Figure 28. Training couples compared with validation data shows poor agreement for low dimensional training ROMs.....	46
Figure 29. (a) Empirical CDFs and (b) Kolmogorov-Smirnov statistics for ROMES generated using a fixed dimension greedy search normalized by the GP mean and standard deviation.	47
Figure 30. Kolmogorov-Smirnov statistics for ROMES generated using a lagged dimension greedy search.	48
Figure 31. Kolmogorov-Smirnov statistics for ROMES generated using a constant energy greedy search.	50

NOMENCLATURE

- I radiative intensity
- I_n Angular intensity at quadrature point n
- σ_T Macroscopic total cross section or extinction coefficient
- w_n Quadrature weight
- σ_s Macroscopic scattering cross section or scattering coefficient
- κ Macroscopic absorption cross section or absorption coefficient
- σ Stefan-Boltzmann constant
- T Material temperature
- I_b Black-body intensity, $I_b = \frac{\sigma T^4}{\pi}$
- $\vec{\Omega}_i$ Unit vector pointing in the ordinate direction corresponding to quadrature point i
- ε Surface emissivity
- \vec{n} Surface normal unit vector
- $\vec{\phi}$ Reduced basis

1. INTRODUCTION

Numerical modeling for the assessment of weapon safety in abnormal thermal environments plays an extremely important role in assuring the safety of both the existing stockpile and ongoing LEPs. To this point, the inclusion of participating media radiation in full system simulations has been prohibitively computationally expensive. The methods developed here have potential value to thermal-mechanical analyses associated with SNL satellite programs and thermal-mechanical analyses of interest to LANL, thus creating the potential for new markets for SNL/ASC code capabilities.

In assessment of abnormal thermal environments, radiation heat transfer can be a dominant and important energy transfer mechanism. Radiation heat transfer can often dominate computational resource requirements (by several orders of magnitude) in these studies. Such analyses are impractical, even with current high performance computing (HPC) resources. Consequently, analysts have often ignored dynamic boundaries and participating media or applied approximate models and other simplifications in these analyses to reduce the computational expense to an acceptable level at the expense of severely compromising the underlying physics considered. The goal of this work is to make accurate modeling of radiative heat transfer with participating media tractable in the modern HPC environment. To this end, we investigate a novel approach to improve the computational efficiency of higher fidelity radiative heat transfer formulations which involves applying parametric model reduction to problems which include participating media radiation (PMR) heat transfer.

1.1. Participating Media Radiation

The steady-state gray radiation transport equation (RTE) with isotropic scattering is given by

$$\vec{\Omega} \cdot \vec{\nabla} I(\vec{\Omega}) + \sigma_T I(\vec{\Omega}) = \kappa I_b + \frac{\sigma_s}{4\pi} \int I(\vec{\tilde{\Omega}}) d\vec{\tilde{\Omega}} \quad (1)$$

The boundary conditions for Eq. (1) are Dirichlet; they specify the outgoing intensity at a surface to be equal to the sum of the surface emission and the reflected intensity. This equation defines the radiative intensity in terms of position and direction of travel, $\vec{\Omega}$, which results in a function of 5 independent variables: 3 in space and 2 in angle. The inputs that parameterize this equation are the absorption and scattering coefficients and the emission source (or alternatively the temperature). The intensities in different directions are only coupled through the right-most (in-scattering) term in Eq. (1) and reflective boundary conditions.

Model-reduction techniques are useful for decreasing the computational cost of many-query problems and are increasingly popular in the areas of optimal design, optimal control, uncertainty quantification, and inverse problems. At first glance, the solution of the RTE does not appear to belong to this class of problems. However, if the angular discretization is performed according to the discrete ordinates method, the angular coordinates may be viewed as independent parameters that must be sampled in much the same way as the parameters in the previously mentioned applications. This observation is the central concept that allows for the

application of parametric model reduction to the problem of simulating participating media radiation.

1.1.1. Discrete Ordinates Method

The discrete ordinates method is a common approach for representing the angular dependence of the radiative intensity. In the discrete ordinates method, the RTE is satisfied along a set of discrete directions and a quadrature rule is used to evaluate integrals over angle. An effective choice of quadrature rule is somewhat problem dependent and there are many options available [1-12]. The choice of quadrature rule defines a set of N directions, $\vec{\Omega}_i$ and weights, w_i $i=1,2,\dots,N$. Eq. (1) is then approximated as a set of N first-order PDEs in only the three spatial dimensions.

$$\vec{\Omega}_i \cdot \vec{\nabla} I_i + \sigma_T I_i = S \quad i = 1, 2, \dots, N \quad (2)$$

The source, S is generally the sum of the emission and in-scattering terms. For systems with scattering, this couples the solution in each ordinate direction with that in every other ordinate direction. This added complexity may be avoided by the use of scattering source iteration [7, 13] in which Eq. (2) is solved repeatedly with different source distributions via fixed point iteration.

$$\vec{\Omega}_i \cdot \vec{\nabla} I_i^n + \sigma_T I_i^n = S^n \quad (3)$$

With the source at the n^{th} iteration given by

$$S^n = \kappa I_b + \frac{\sigma_s}{4\pi} \sum_k w_k I_k^{n-1} \quad (4)$$

The solution to Eq. (2) is then approached as n approaches infinity. For highly scattering media, n of several tens may be necessary but for most problems $n < 10$ is sufficient.

For a given source distribution S , Eq. (2) may be solved independently for each ordinate direction. These solutions are all similar, varying only in the direction parameter, $\vec{\Omega}_i$. For 1D slab geometry problems, this parameter space is one-dimensional as only the component of the direction of travel aligned with the spatial dimension enters into the governing equation. For 2D and 3D problems, this parameter space is two-dimensional. Any point on the unit sphere (or hemisphere in the 2D case) may be described by two variables such as the azimuthal and circumferential angles.

2. MODEL REDUCTION STRATEGY

The goal of the parameterized model reduction strategy used is to generate a reduced basis for possible solution vectors (i.e. intensity distributions) across the parameter (angular) space. The first step is to discretize Eq. (2). In the results to follow, the spatial discretization is accomplished through linear finite elements. The discrete problem is given by

$$\overline{\overline{K}}(\vec{\Omega})\vec{I}(\vec{\Omega}) = \vec{S} \quad (5)$$

where $\overline{\overline{K}}$ is an $m \times m$ matrix where m is the number of nodes in the spatial mesh. Eq. (5) is linear with respect to both the unknown intensity as well as the directional parameter. The existence of this reduced basis may be confirmed by sampling the parameter space and computing the dimensionality of the span of the solution vectors generated. For a fixed source distribution, for several values of the directional parameter, $\vec{\Omega}_1, \vec{\Omega}_2, \dots, \vec{\Omega}_K$ the solution vector $\vec{I}_1, \vec{I}_2, \dots, \vec{I}_K$ is computed. These solution vectors are then concatenated to construct the global $m \times K$ basis matrix.

$$\overline{\overline{M}} = [\vec{I}_1, \vec{I}_2, \dots, \vec{I}_K] \quad (6)$$

If a large number of samples K are taken, it is likely that the column rank of $\overline{\overline{M}}$ is less than K . To address this situation, the singular value decomposition (SVD) of $\overline{\overline{M}}$ is generated [15] yielding 3 matrices $\overline{\overline{U}}$, $\overline{\overline{S}}$, and $\overline{\overline{V}}$ where $\overline{\overline{M}} = \overline{\overline{U}}\overline{\overline{S}}\overline{\overline{V}}^T$. $\overline{\overline{S}}$ is a diagonal matrix containing the singular values and $\overline{\overline{U}}$ is a full matrix containing the modes as its columns. The presence of a reduced basis may be confirmed by examining the decay of these singular values. In the 1D example cases, 99.9999% of the energy is captured by the first 5 or 6 modes depending on the source distribution. In the 2D example cases, capturing the same fraction of the total energy requires between 10 and 35 modes depending on the source distribution.

Fig. 1 shows the decay for a number of different cases. The precise behavior of the singular values is problem dependent, however, the qualitative behavior appears to be similar for problems of the same spatial dimension. The rapid decay in singular values suggests that a relatively small number of modes is likely to be sufficient to capture the vast majority of the solution behavior. Hence the necessary condition of the existence of a low-dimensional solution space is satisfied.

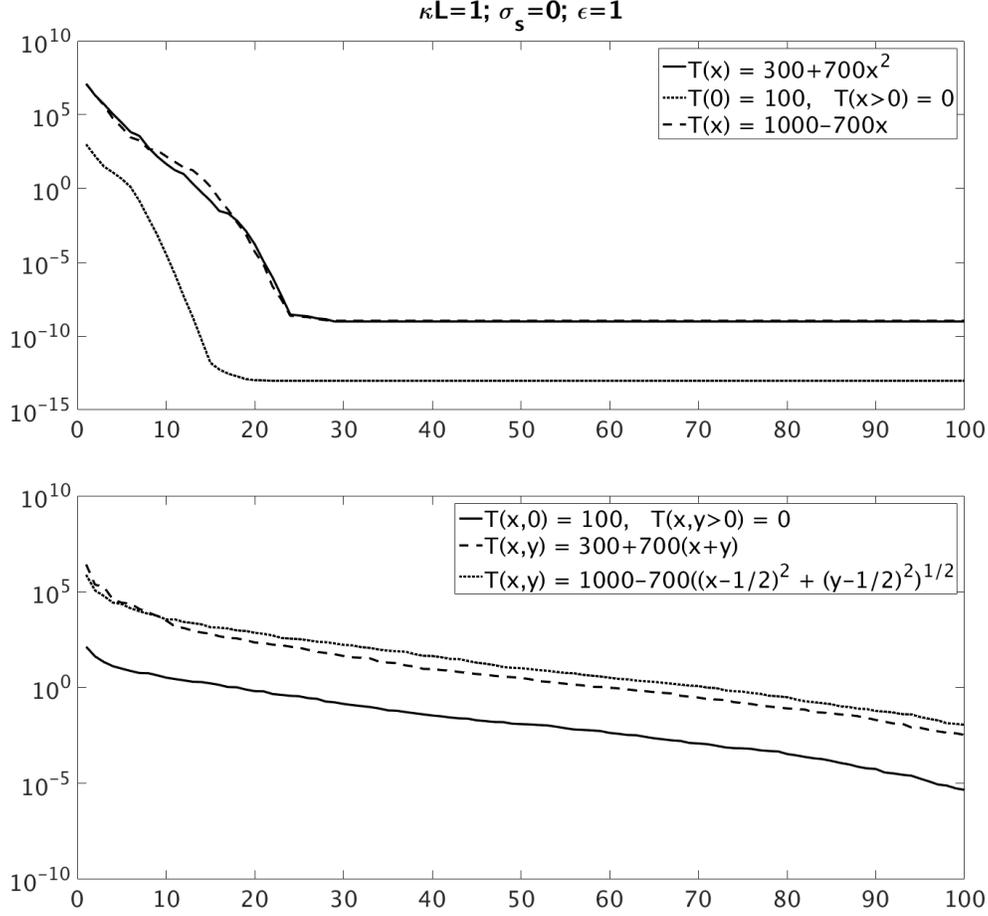


Figure 1. Decay of singular values demonstrates the presence of a reduced order basis for the intensity solution for several source distributions. (a) 20 modes are sufficient to capture all behavior in the 1D geometry. (b) 50 modes would capture a very large percentage of the solution behavior in the 2D geometry.

As a result, only the primary modes of $\overline{\overline{M}}$ corresponding to the first $k < K$ columns of $\overline{\overline{U}}$ are used. These k columns define the reduced basis, $\overline{\overline{\phi}} = [\vec{U}_1, \vec{U}_2, \dots, \vec{U}_k]$ through proper orthogonal decomposition (POD). The discretized intensity in any direction is then approximated as

$$\vec{I}(\vec{\Omega}) \approx \overline{\overline{\phi}} \vec{x} \quad (7)$$

Where \vec{x} is a k -dimensional vector and $\overline{\overline{\phi}}$ is a tall skinny $m \times k$ matrix with $k \ll m$. The reduced order model (ROM) is then given by substituting Eq. (7) into the full-order model (FOM) given by Eq. (5).

$$\left(\overline{\overline{K}}(\vec{\Omega}) \overline{\overline{\phi}} \right) \vec{x} = \vec{S} \quad (8)$$

Eq. (8) is over-determined, i.e., a solution may not exist. Instead, a unique solution can be computed by enforcing the Galerkin orthogonality condition or by applying least-squares Petrov-Galerkin projection [29], which minimizes the residual and solves

$$\left(\overline{\overline{K}}(\overline{\overline{\Omega}})\overline{\overline{\phi}}\right)^T \left(\overline{\overline{K}}(\overline{\overline{\Omega}})\overline{\overline{\phi}}\right)\overline{\overline{x}} = \left(\overline{\overline{K}}(\overline{\overline{\Omega}})\overline{\overline{\phi}}\right)^T \overline{\overline{S}} \quad (9)$$

Eq. (9) is a $k \times k$ dense linear system whereas Eq. (5) is a $m \times m$ sparse linear system. Provided that $k \ll m$, finding the solution of Eq. (9) is much less computationally expensive than finding the solution of Eq. (5). Because the discrete ordinates method requires many solutions to Eq. (5) a significant reduction in computational cost may be realized by replacing a significant fraction of those solutions with solutions to Eq. (9) instead.

2.1. Sampling Methods

The most straight-forward way to integrate over the angular parameter space is using an established quadrature rule. The nodes of this quadrature rule correspond to the values of $\overline{\Omega}$ for which either Eq. (5) or Eq. (9) must be evaluated. Traditionally, Eq. (5) is used at all points and by necessity a low-order quadrature is used (due to computational constraints). Using Eq. (9) instead allows for the use of a significantly higher order quadrature and, as will be shown in subsequent sections, a significantly more accurate solution. The question remains if using Eq. (9) how best to determine the reduced basis, $\overline{\overline{\phi}}$. This is a question of sampling.

The most intuitive sampling method is to simply sample at the nodes of a low-order quadrature rule. This insures that sample points are reasonably well-distributed throughout the angular domain. Additionally, this approach provides a convenient assessment of the accuracy gained by incorporating a ROM. To be more precise, the full-order model is evaluated for a low-order quadrature just as one would typically do when a high-order quadrature is either impractical or unnecessary. This yields a set of angular intensities in the given ordinate directions $\vec{I}_1, \vec{I}_2, \dots, \vec{I}_K$ where K is the number of nodes in the low-order quadrature rule. Important angularly-integrated quantities such as the angle-integrated intensity (scalar flux) may then be computed using the quadrature rule.

$$\vec{G}^{LOM} = \sum_{i=1}^K w_i \vec{I}_i \quad (10)$$

This is the classical approach to the DOM and is referred to as the low-order model (LOM). In order to apply parametric model reduction to this problem, the intensities $\vec{I}_1, \vec{I}_2, \dots, \vec{I}_K$ might instead be viewed as samples for generating a reduced basis, $\overline{\overline{\phi}}$. This basis defines a ROM which may be queried inexpensively. We may define a high-order quadrature with $L \gg K$ nodes. Solving Eq. (9) for each of these nodes results in a large set of approximate intensities $\vec{I}_1^{ROM}, \vec{I}_2^{ROM}, \dots, \vec{I}_L^{ROM}$. These approximate intensities may then be used in combination with the high-order quadrature weights to produce a different estimate of the angle-integrated intensity.

$$\vec{G}^{ROM} = \sum_{i=1}^L w_i \vec{I}_i^{ROM} \quad (11)$$

The accuracy of these two different estimates to the angle-integrated intensity may be compared. It is observed that \vec{G}^{ROM} is significantly more accurate than \vec{G}^{LOM} provided K is sufficiently large. However, it is also observed that this benefit is greatly reduced for large K due to the inefficiencies associated with using many popular quadrature rules as a sampling method. Quadrature rules tend to possess a high degree of symmetry which necessitates the presence of many samples in directions where the angular intensity is not changing rapidly resulting in a high degree of redundancy. Additionally, high-order level-symmetric quadratures tend to place many samples near the coordinate axes causing these regions of the angular space to be over-sampled relative to the rest of the space.

These difficulties are overcome by adopting a more efficient sampling strategy such as a greedy search algorithm [17 - 20]. In the greedy search algorithm, samples are chosen adaptively by placing the new sample point at the location where the estimated error in the ROM prediction is maximum. This process bears some similarity to previously developed adaptive quadrature methods [21-27]. However, it is fundamentally different from those techniques, which seek to enrich a low order quadrature through local refinement. In the proposed method, a ROM is constructed to represent a very high order quadrature and the ROM training points (ordinate directions) are chosen adaptively.

As has been exploited in ROM error modeling methods [30], the ROM residual for a given direction is closely related to the ROM error for that direction and the residual is selected as an effective error indicator. The ROM residual tends to be a highly oscillatory function of the angle with many local minima and maxima making the continuous optimization problem inherent to each step of the greedy algorithm difficult. To avoid this difficulty, a discrete optimization problem is solved instead where the sample points are constrained to belong to a high order quadrature.

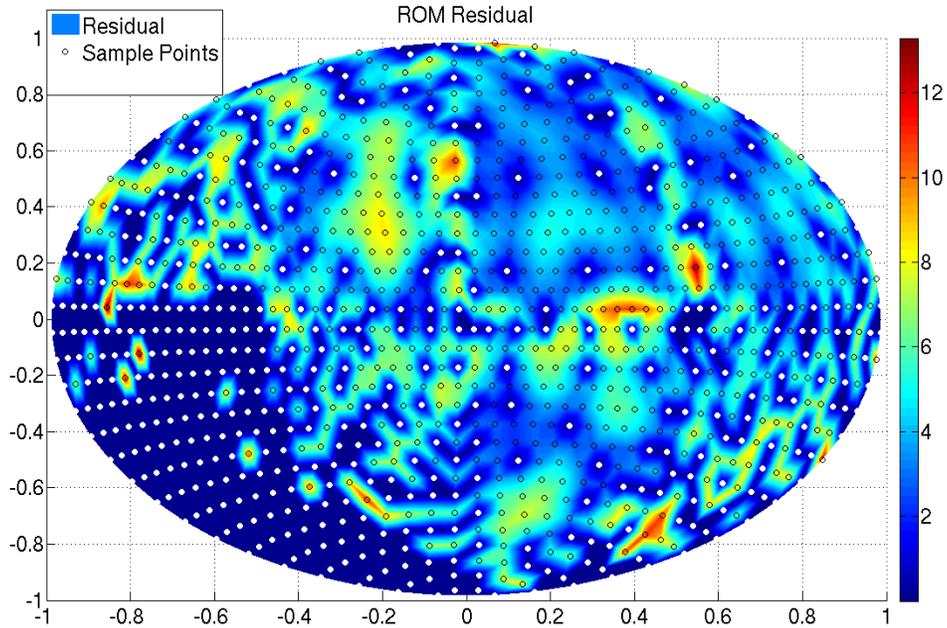


Figure 2. Distribution of sample points generated through greedy search algorithm for a radial temperature distribution $T(r)=300+400r^2$ colored by the ROM residual. Circles correspond to points in the high-order quadrature. White circles correspond to points used to train the ROM.

Fig 2 shows the distribution of sample points resulting from one such greedy search after 440 sample points have been selected. The filled white circles show the locations of the sample points while the open black circles show the locations of nodes of the high-order quadrature not selected by the greedy search algorithm. The greedy search algorithm preferentially places sample points in the directions with large directional intensity values. In this case, those are directions pointing back towards the origin (lower left octant in Fig 2). In practice, acceptable levels of accuracy may be reached with far fewer sample points than are shown in Fig 2.

The greedy search algorithm is inherently suboptimal so there is little value added by finding the exact location of the maximum error at each step. Significant time savings may be achieved by choosing the maximum of as few as 10 randomly selected points and only evaluating the ROM residual at these points and then sampling at the point with the largest residual norm. This greatly reduces the number of ROM evaluations required. Although a ROM evaluation is significantly less expensive than a FOM evaluation, evaluating the ROM for every quadrature point of the HOM prior to generating each sample can become expensive. Additional computational savings may be had by seeding the ROM with a number of predetermined sample points prior to initiating the greedy search algorithm reducing the number of steps required for a given number of samples.

3. RESULTS

The parametric reduced order modeling approach described in section 2 provides significant reductions in the computational effort required to generate solutions to PMR problems using the DOM. The algorithm also provides analysts with enhanced control over solution accuracy and increased confidence in simulation results through the use of robust error estimators to be described in detail in section 4. The computational cost benefits of this algorithm are demonstrated for a set of 1D, 2D, and 3D example problems. The 1D and 2D problems make use of structured grids. The 3D problems are solved using an unstructured tetrahedral mesh.

3.1. 1-D Example

Consider the 1D case of a purely absorbing slab surrounded by black walls. The temperature profile is assumed quadratic, $T(x) = 300 + 700x^2$. For 1D geometries, the Gauss-Legendre quadrature rules are typically employed. Consider the angle-integrated intensity generated by the Gauss-Legendre quadrature with 6 directions (S6). The error in the angle-integrated intensity in this case is approximately 2%. Contrast this with the distribution generated by the Legendre-Gauss quadrature with 200 directions (S200) which has an expected error of about 0.001%. The S6 model will be referred to as the low-order model (LOM). The S200 model will be referred to as the high-order model (HOM). The intensities from the S6, LOM solution may be used to generate a ROM capable of predicting intensities in any arbitrary direction. If this ROM is then used to generate intensities for the S200 quadrature, the resulting error in the angle-integrated intensity is less than 0.3%. The convergence of the LOM and ROM solutions as the LOM quadrature order is increased is plotted in Fig 3a.

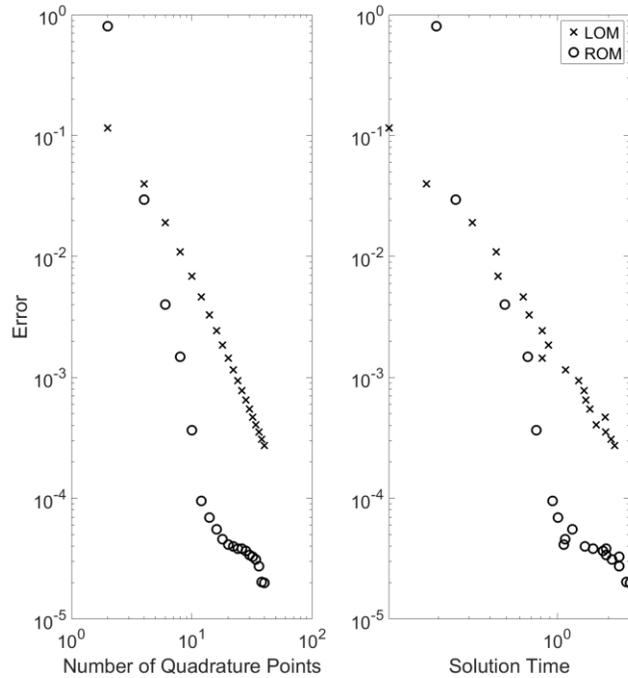


Figure 3. Convergence of angle-integrated intensity distributions for quadratic temperature profile with increasing LOM quadrature order.

The results in Fig 3a do not tell the whole story since the ROM is both more accurate and necessarily more expensive than the LOM with the same number of quadrature points. Fig 3b shows the comparison of execution time and accuracy for each method. The trends in Fig 3b are only valid for sufficiently large problems ($m \gg k$) otherwise model reduction provides little advantage. Often it is possible in 1D problems to get away with an extremely coarse mesh such that this condition is not satisfied. However, this condition is almost always true for any practical 2D or 3D problem.

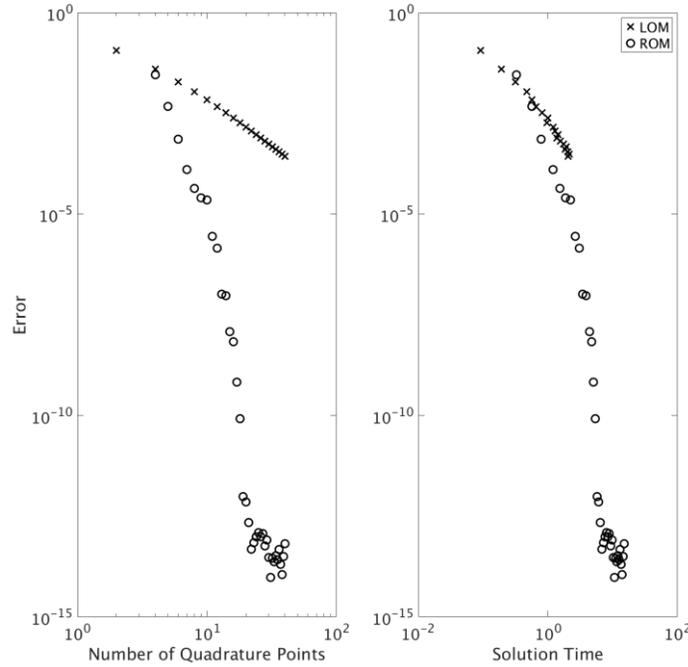


Figure 4. Convergence of angle-integrated intensity distributions for quadratic temperature profile with greedy sampling method.

The inefficiencies associated with using a Gauss-Legendre quadrature to generate sample points is apparent as the number of quadrature points in the LOM, K becomes large (see Fig 3). The rapid convergence of the ROM solution slows to approximately the same rate as that of the LOM solution. Alternatively, the greedy sampling approach may be used. The results of this approach are shown in Fig 4. The ROM solution is seen to converge rapidly to machine precision after approximately 20 samples. This is consistent with the decay of singular values shown in Fig 1a. These results show a great deal of promise for ROMs to reduce the computational expense and increase the accuracy of discrete ordinates simulations in 1D.

3.2. 2-D Examples

To demonstrate this solution methodology in higher dimensional geometries, consider the case of a square surrounded by black walls and filled with a purely absorbing medium with opacity equal to the inverse of the side length of the square. Consideration of a purely absorbing medium is sufficient for demonstrating the validity of this approach as this is mathematically equivalent to considering any individual source iteration step for a scattering problem. The LOM considered here is the 14th order PNTN quadrature [9] which still contains 112 ordinate directions. The HOM is the 32nd order quadrature of the same type which contains 544 ordinate directions. The PNTN quadrature rule is defined by using the Gauss-Legendre quadrature set to define the levels along the z -axis as well as the total weight for each level. The azimuthal angles for each level are set equal to the roots of the Chebyshev polynomials.

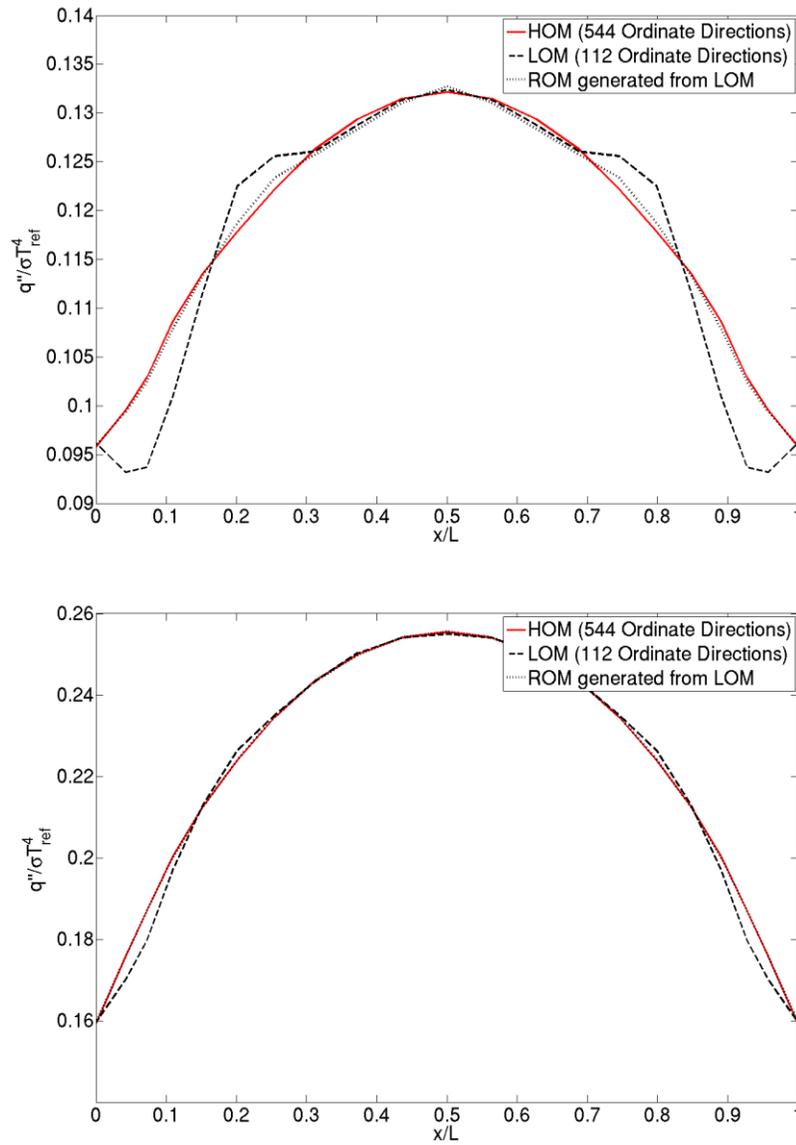


Figure 5. Normalized heat flux predictions $q''(x,1)$ for (a) a discontinuous temperature profile $T(x,y)=1-H(y)$ and (b) a linear temperature profile $T(x,y)=300+700(1-y)$ for the 14th order PNTN quadrature, 32nd order PNTN quadrature, and the ROM derived from the 14th order quadrature solutions but evaluated at the 32nd order quadrature points.

Fig 5 shows the heat flux evaluated at the top ($y=1$) surface for a pair of temperature distributions. Both the accuracy of the solution and the improvement provided by the ROM relative to the LOM are seen to depend strongly on the source distribution, the order of the LOM, and the mesh resolution. However, in all cases a rapid increase in accuracy was observed once the LOM achieved a sufficient order.

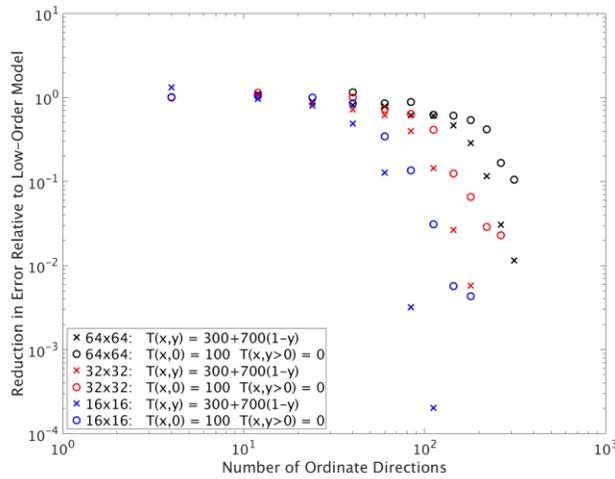
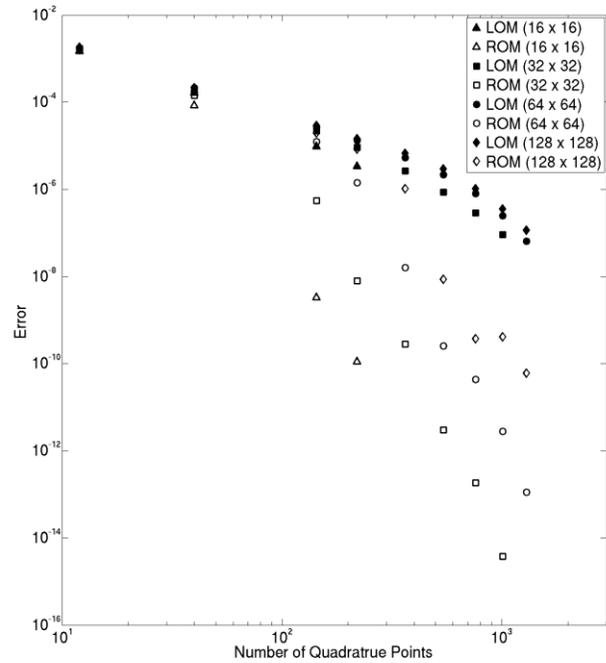


Figure 6. (a) Convergence of LOM and corresponding ROM for a temperature distribution of $300+700(1-y)$ for various spatial mesh resolutions (b) Reduction in error provided by the ROM relative to the low-order model for various spatial mesh resolutions. This is the L2 error of the ROM solution divided by the L2 error of the low-order model solution.

Fig 6a shows the convergence of the LOM and corresponding ROM as the quadrature order of the LOM is increased for a temperature distribution of $300+700(1-y)$. We do not observe the same type of stagnation as was seen in the 1D example. This is due to the more long-tailed decay of the singular values in the 2D case as seen in Fig 1. Fig 6b shows the reduction in error in the angularly integrated intensity distribution relative to the LOM for a variety of source distributions and mesh resolutions. In all of these cases a regular rectangular grid is used. The

relative improvement is seen to be more pronounced for the linear temperature profile than for the discontinuous temperature profile. This is related to the fact that the LOM error tends to be smaller for the linear temperature profile. The number of ordinate directions required in the LOM to sufficiently inform the ROM and achieve the greatly enhanced accuracy is observed to increase with mesh refinement. This is likely due to the reduction in false scattering with increasing mesh resolution.

3.3. 3-D Examples

We now move on to 3D applications. Consider a purely absorbing cube with an optical side-length of one surrounded by black walls. Figure 7 shows the relatively slow decay of the singular values for this geometry for several temperature distributions. Fig 7 implies that the methodology is significantly less efficient for the 3D problem than was demonstrated for the 2D problem. Between 711 and 1284 modes are required to capture 99.9999% of the energy depending on the source distribution while between 42 and 296 are required to capture even 99.9% of the energy. This is partially due to reduced symmetry and partially due to the mesh size. The same quadrature includes twice as many angles in 3D as it does in the 2D case. Also, the number of degrees of freedom in the 3D simulation is significantly greater than the number in the 2D simulation. The end result is that the LOM may be required to be impractically large in order to adequately inform the ROM and achieve significant accuracy gains.

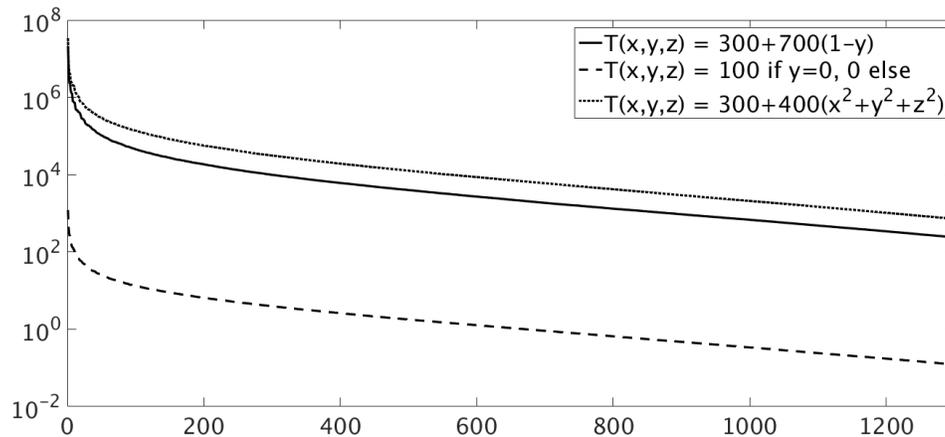


Figure 7. Decay of singular values demonstrates that many additional modes are required to capture the behavior of the fully 3D problem.

However, this obstacle is not insurmountable. The failure of the previously described sampling technique is largely due to inefficient sampling of the angular parameter space. This difficulty may once again be overcome by employing an adaptive sampling method such as a greedy search. In the examples to follow, the greedy sampling method is used and the ROM is seeded with 48 sample points analogous to the P6-T6 quadrature.

The HOM is chosen to be the 40th order PN-TN quadrature (P40-T40) which includes 1680 ordinate directions. This number of ordinate directions has been shown to be sufficient to eliminate ray effects from the solution for the spatial mesh resolutions considered here [28]. The

adaptive ROM performs better as the mesh is refined. This is because the cost of FOM evaluations increases much faster than the cost of ROM evaluations as mesh size increases since the size of the ROM linear system does not increase with the size of the mesh.

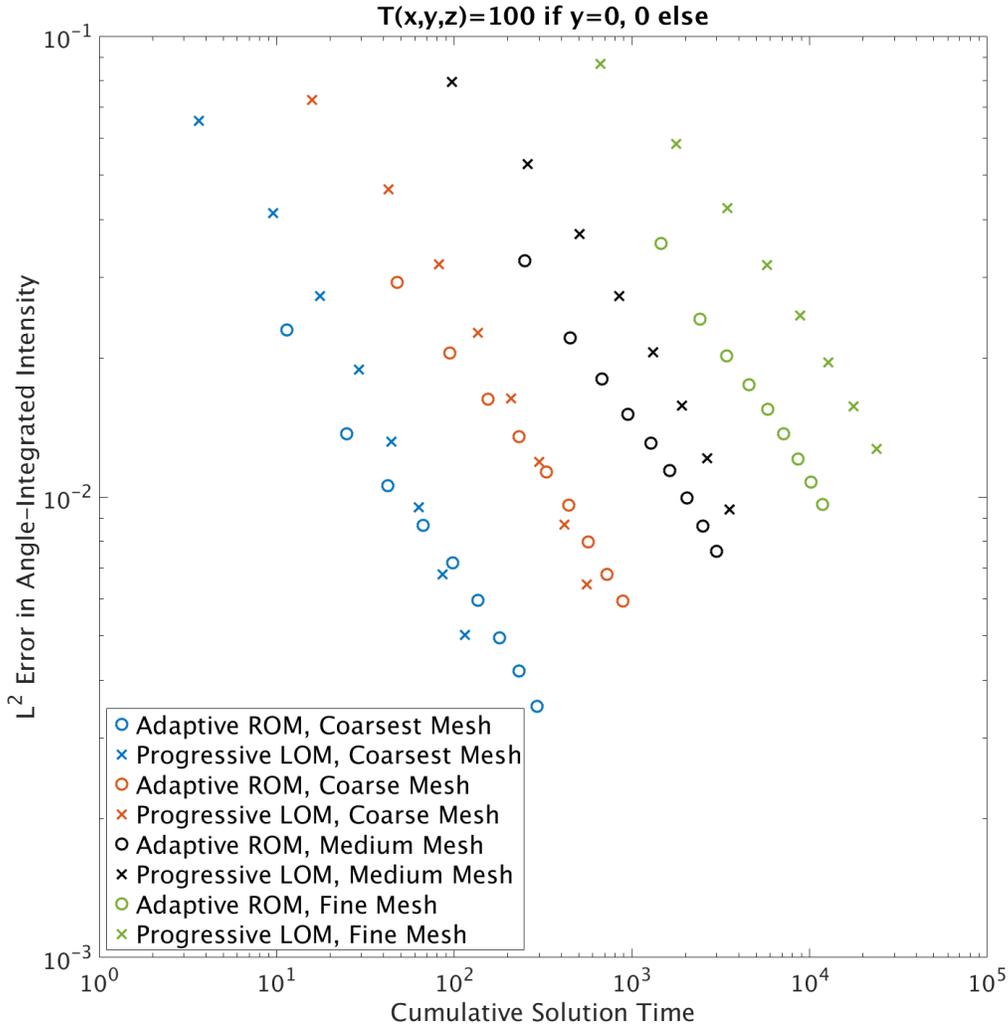


Figure 8. Accuracy as a function of cumulative solution time for a discontinuous temperature distribution $T(x,y,z)=100$ if $y=0$, 0 else.

Fig 8 shows the results of a timing study for a 3D example problem. Because the adaptive ROM yields a much more accurate solution than the LOM when given the same number of FOM evaluations but is slower because the adaptive ROM evaluates the ROM a large number of times in addition to the set number of FOM evaluations, it makes sense to compare the two approaches on the basis of error and solution time.

The error in the adaptive ROM may be (inexpensively) estimated at any time and used as a stopping criterion for the greedy search algorithm. The timings shown in Fig 8 include this error

estimation although a maximum number of FOM evaluations was used as the stopping criteria rather than a predetermined error level. There is presently no analogous technique for estimating the error incurred by the discrete ordinates method if too small a quadrature set is used as the LOM. Instead, an understanding of the magnitude of the error is typically obtained by performing multiple LOM evaluations with increasing quadrature orders. For example, one might evaluate the LOM using the P6-T6, P8-T8, and P10-T10 quadratures which include 48, 80, and 120 ordinate directions respectively and use the differences between the solutions to infer an approximate level of accuracy for the P10-T10 solution. If the inferred level of accuracy is inadequate, successively higher order quadratures may be invoked until the accuracy is acceptable.

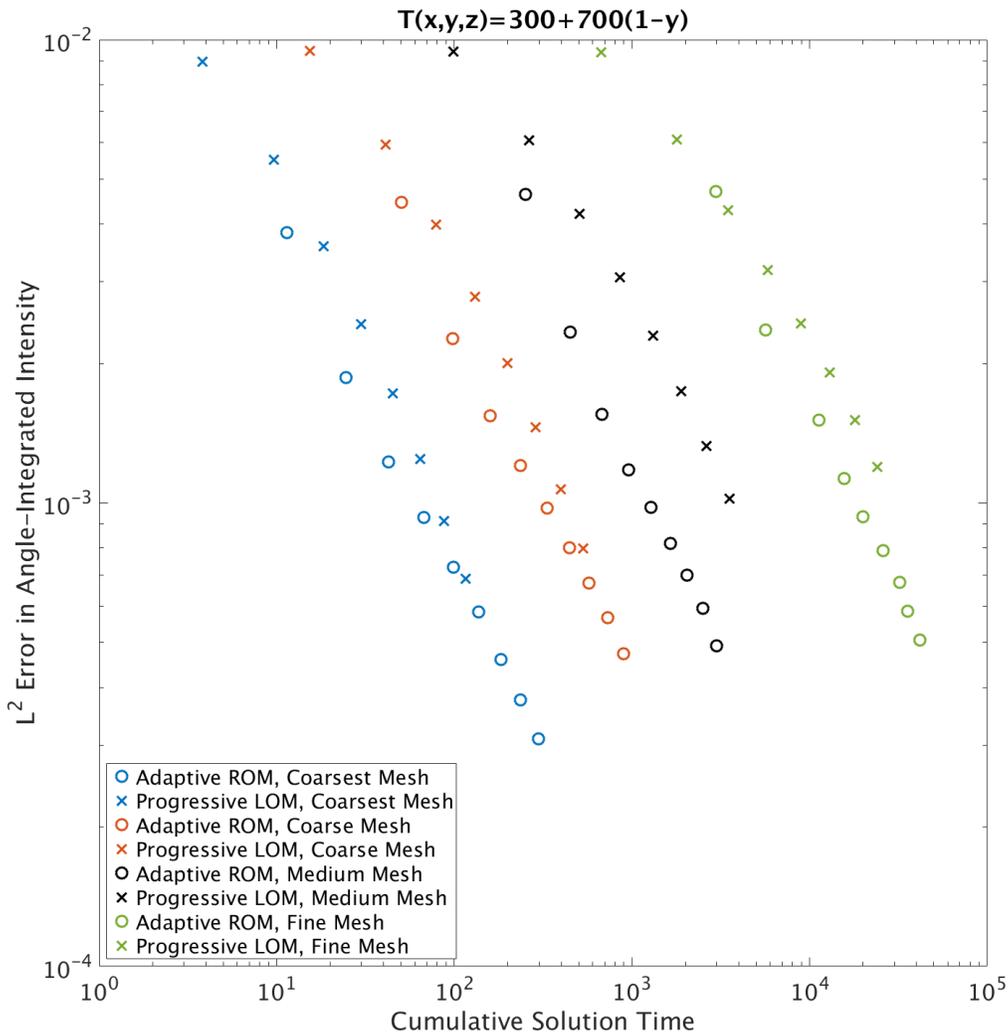


Figure 9. Accuracy as a function of cumulative solution time for a linear temperature distribution, $T(x,y,z)=300+700(1-y)$.

For this reason, the cumulative solution time is used in Figs 8, 9, and 10 which demonstrate the effect of spatial mesh resolution on the convergence behavior for three different temperature profiles. It is shown in all three figures that seeding the adaptive ROM with 48 ordinate directions and applying no refinement (the first data point) is both more accurate than the P10-T10 solution and significantly faster than evaluating the 3 LOMs listed above. This initial speed advantage increases with spatial mesh resolution.

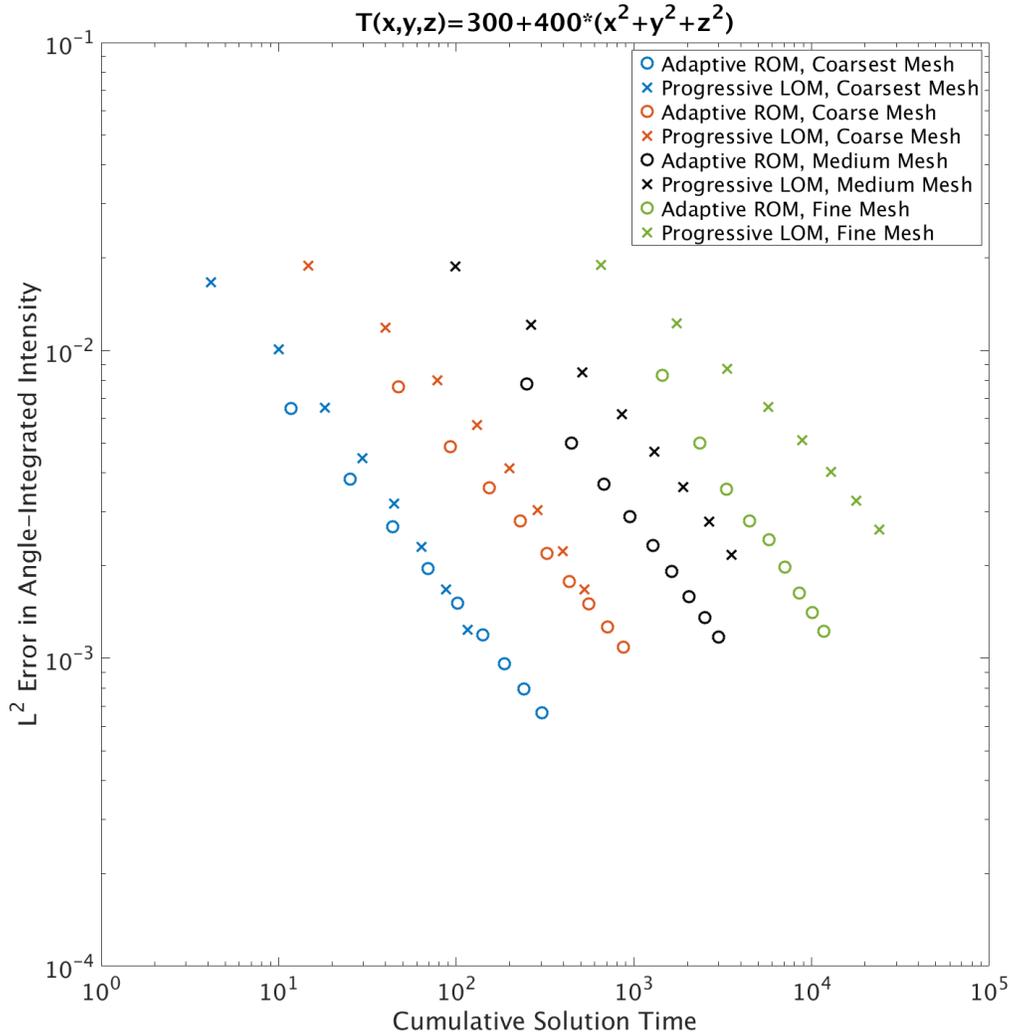


Figure 10. Accuracy as a function of cumulative solution time for a radially quadratic temperature distribution $T(x,y,z) = 300+400(x^2+y^2+z^2)$.

These conclusions hold true for a wide variety of temperature distributions. The series of meshes used in this set of example problems included approximately 3.4k, 12.2k, 36.5k, and 111.8k nodes. In all cases, the medium is purely absorbing and the walls are black. The wall temperature is equivalent to the temperature of the adjacent medium.

4. ERROR ESTIMATION

The adaptive ROM error may be estimated at any time by constructing a relationship between the ROM residual and the ROM error. At each step in the greedy search algorithm, both of these quantities are known for what will become the next sample point. The form of the relationship between the residual norm and the error norm is roughly a power law as shown in Fig 11. The training data in Fig 11 represents the residual norm – error norm pairs that are known from previous iterations of the greedy search algorithm. These errors are the L2 errors in the directional intensity. The prediction data are the errors calculated for all quadrature points in the HOM not included in the adaptive quadrature. The training data are used to generate a curve fit that allows for the approximation of the error at any point for which the ROM residual norm is known; this can be interpreted as a reduced-order model error surrogate (ROMES) model [30].

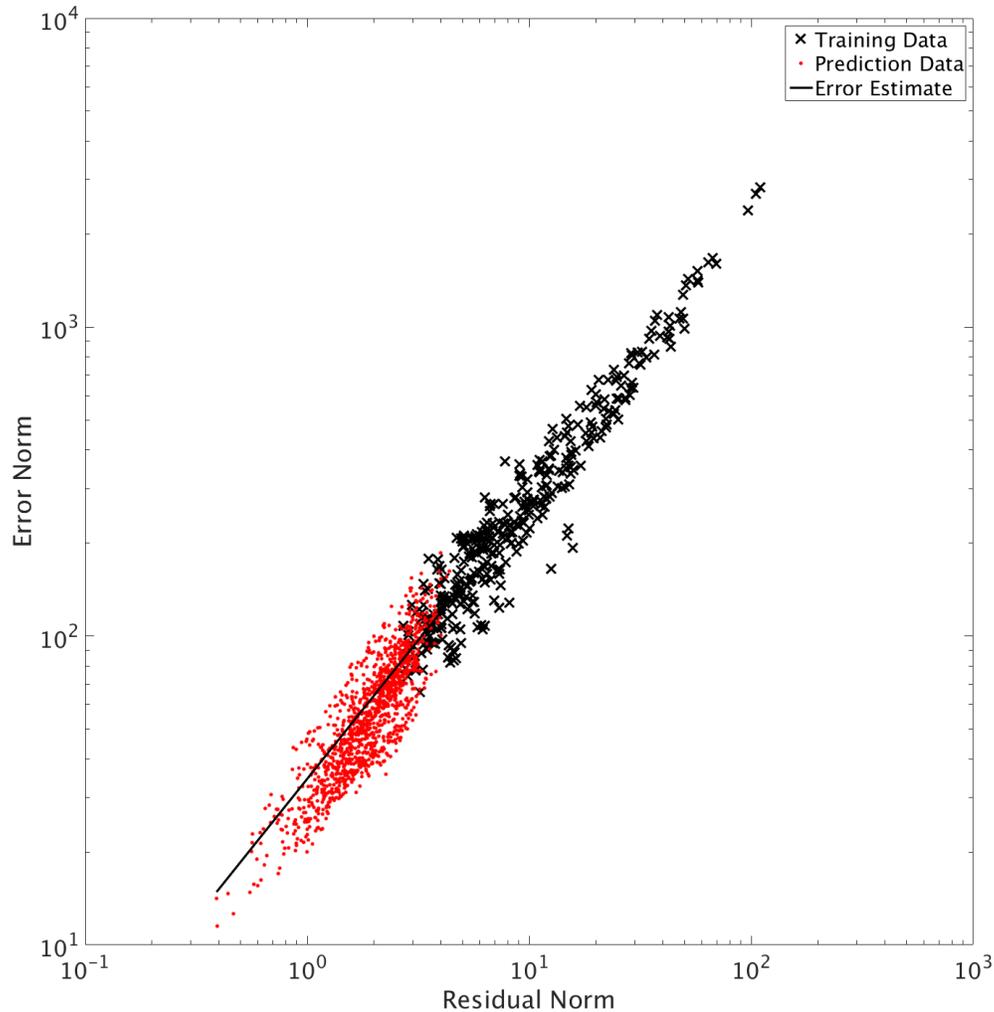


Figure 11. Example of construction of ROM error estimate from training data acquired during the greedy search algorithm for a linear temperature distribution, $T(x,y,z)=300+700(1-y)$.

This process provides an estimated distribution of the ROM error over the angular space. The resulting error in angularly integrated quantities may then be estimated as well. Fig 12 shows the convergence of one such quantity.

$$\text{Error} = \sum_i w_i^2 \int (I_i^{FOM} - I_i^{ROM})^2 dV \quad (12)$$

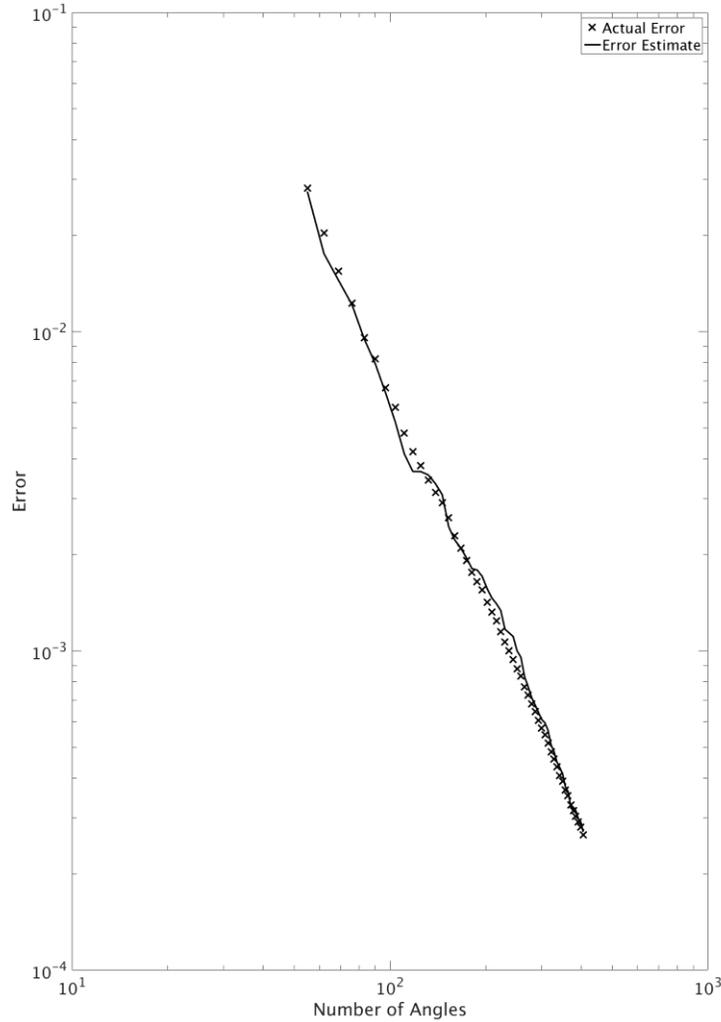


Figure 12. Convergence of adaptive ROM and associated error estimate of an angularly integrated quantity for a linear temperature distribution, $T(x,y,z)=300+700(1-y)$.

As seen in Fig 12, the deterministic estimate of the error derived from this process is quite accurate and is easily sufficient for use as a stopping criteria. However, for the purposes of uncertainty quantification (UQ) it is often desired not only to know the approximate error level

but also the degree of consequence in the prediction of that error level. Toward this end, we can use the available residual/error data shown in Fig 11 to generate a stochastic rather than a deterministic model of the relationship between the residual and the error using a Gaussian process (GP). Fig 13 shows one example of this. The mean of the GP is very similar to the deterministic estimate as one would expect given how well the deterministic estimate predicts the average behavior.

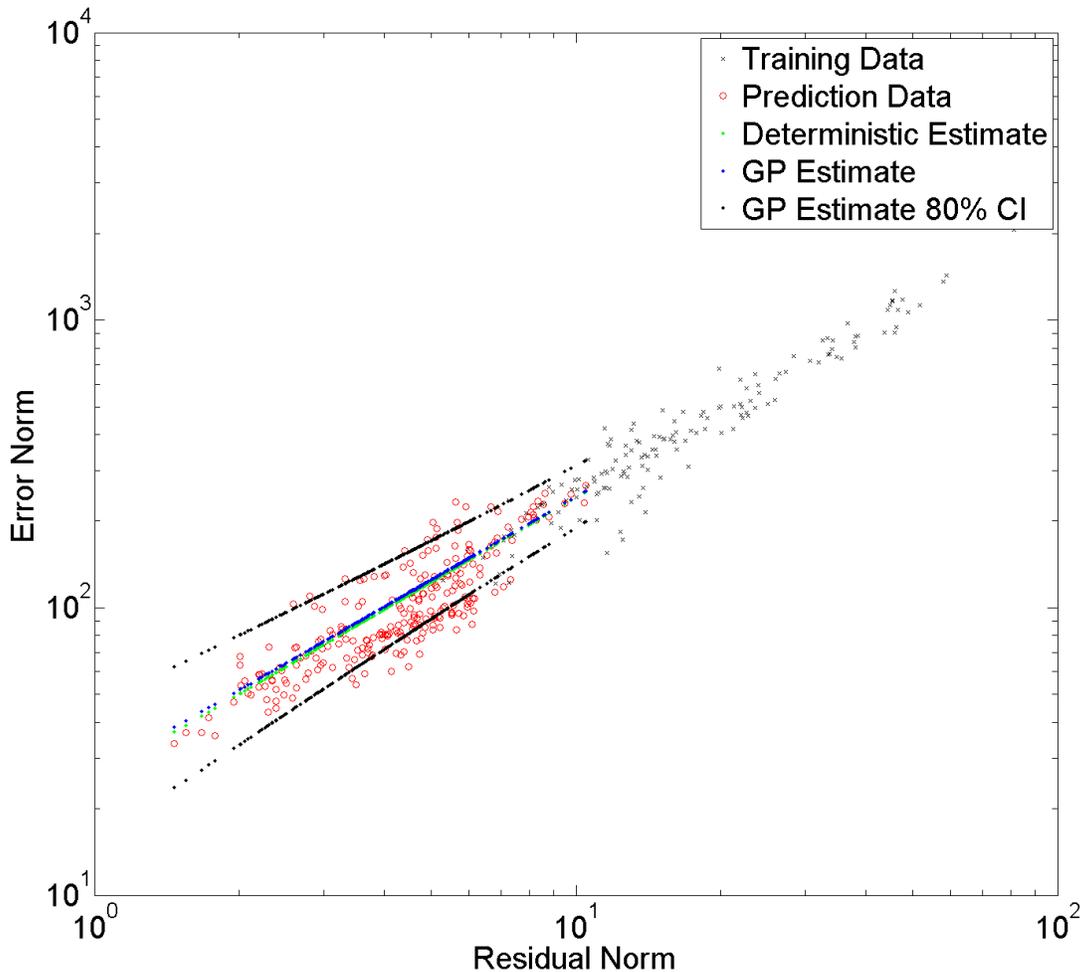


Figure 13. Comparison of deterministic and Gaussian process estimates of the L2 error in the angular intensity.

The difference is in the extra information contained in the GP. Rather than only being able to predict the average error, the GP is also able to predict the variability in the error or rather the uncertainty in the error prediction. In Fig 13, the GP is trained using the residual norm – error norm pairs generated during the greedy search process. This is far from the only option. Assuming that the solutions to all of the FOM evaluations performed over the course of the greedy search process are retained, the GP ROMES may be generated to assess the model-form uncertainty at any point, even after the conclusion of the simulation.

Given a set of FOM solutions, it remains to be determined how best to generate a ROMES. This generally comes down to a question of how to allocate the available FOM solutions to train the ROM, ROMES, or both. Given n FOM solutions to be used as training points, some number $m \leq n$ of solutions are used to inform the ROM. This ROM is referred to as the prediction ROM and has basis dimension d_p where $d_p \leq m$. It is this ROM that is used to make the final prediction and it is the error of this ROM that the ROMES seeks to reproduce. Residual norm – error norm pairs or “ROMES training couples” are required to train the ROMES. These are generated by evaluating one or more training ROMs at points where the FOM solution is known. The training ROMs are generated using $p \leq n$ FOM solutions and are of dimension $d_t \leq p$. There is no requirement that $d_t = d_p$. Different training ROMs may be generated using different FOM solutions for training points and/or different values of d_t .

In the following subsections, 9 potential options are examined.

1. Decouple the training of the ROM and the ROMES. Set $m + p = n$ and $d_t = d_p$
2. Distinct training sets with variable truncation. Set $m + p = n$ and $d_t \in \{d_p, d_p - 1, d_p - 2\}$
3. Concurrent training sets. Set $m = p = n$ and $d_t = d_p$
4. Concurrent training sets with variable truncation. Set $m = p = n$ and $d_t \in (1, d_p]$
5. k-fold cross validation. Set $m = n$, $p = n - k$, and $d_t = d_p \leq n - k$
6. k-fold cross validation with variable truncation. Set $m = n$, $p = n - k$, and $d_t \in \{d_p, d_p - 1, d_p - 2\}$
7. Fixed dimension greedy search. $d_t = d_p$ but ROM training set is gradually expanded in same order as greedy search. One ROMES training couple generated at each step in greedy search process for which $m > d_p$
8. Lagged dimension greedy search. Similar to the fixed dimension greedy search except $d_t \in \{d_p, d_p - 5, d_p - 10\}$
9. Constant energy greedy search. ROM training points are added sequentially through the greedy search process. The prediction ROM dimension, d_p is set so as to capture a certain fraction of the statistical energy in the full training set. At each step d_t is set so as to capture the same fraction of the statistical energy of the subset of the training set thus far explored by the greedy search process.

In the examples to follow a training set of size $n = 200$ is reused. This training set is used to generate both a ROM and a corresponding ROMES using each of the 9 options enumerated above. The performance of the ROMES is then assessed by evaluating both the ROM and the FOM at an additional 4000 points in the parameter space and comparing the actual difference between the ROM and FOM solutions with the ROMES prediction.

4.1. Decoupled

The ROM and ROMES trainings are completely decoupled. The prediction ROM is only trained on a subset of the available data points. The training ROM is generated using the same subset. The unused data points are used to train the ROMES. The training and prediction ROMs are of the same dimension. Since they are generated with the same data points, this means that both ROMs are identical. $\frac{1}{2}$ of the available 200 training points will be used to generate the ROM. The ROM is then evaluated at the remaining 100 data points to generate 100 ROMES training couples. The resulting error surrogate is then evaluated against the validation data. Fig 14 shows how well the ROMES predicts the ROM error.

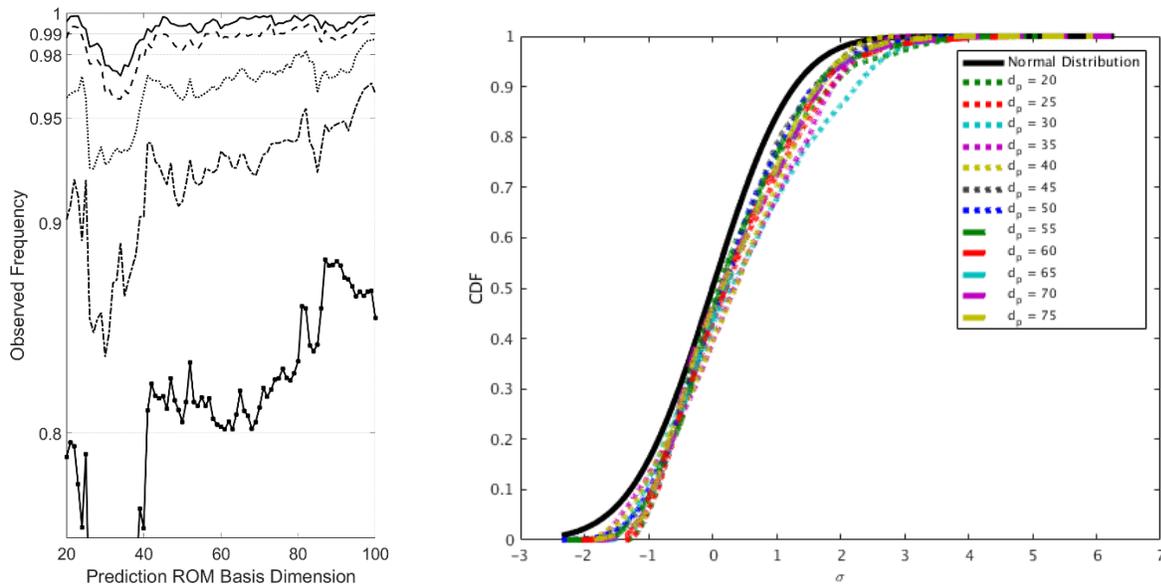


Figure 14. (a) Observed frequencies from above table (b) Empirical CDFs of validation data normalized by GP mean and standard deviation compared with normal CDF.

This approach to building the error surrogate works reasonably well for a wide range of ROM basis dimensions. However, it is limited in two ways. First, the total number of ROMES training couples is limited to a fraction of the number of available data points. Second, the prediction ROM does not utilize all of the available data and is thus less accurate than is possible given all the data. This may result in insufficient training data for either the ROM, the ROMES or both.

The appropriateness of the Gaussian Process estimate may be assessed using the Kolmogorov-Smirnov statistic, D_n which is the maximum difference between the empirical CDFs shown in Fig 16 and the CDF corresponding to the standard normal distribution. If the validation data comes from the distribution described by the GP, then D_n will converge to 0 almost surely as the number of validation points goes to infinity.

The CDFs in Fig 14b match the standard normal distribution reasonably well with KS statistics ranging from 0.08 to 0.18, but it's difficult to tell how valid the assumption of normality is for

this data. With only 100 ROMES training couples, it is difficult to determine if the underlying distribution is significantly non-normal or if the ROMES has been insufficiently trained. To answer this question, the number of ROMES training couples is greatly expanded to 10,000. For the same prediction ROM of basis dimension 100 the experiment is repeated with the larger ROMES training set. This results in the CDF in Fig 15. The KS statistic of 0.04 represents an approximate lower bound for this combination of ROM training points and ROM basis dimension.

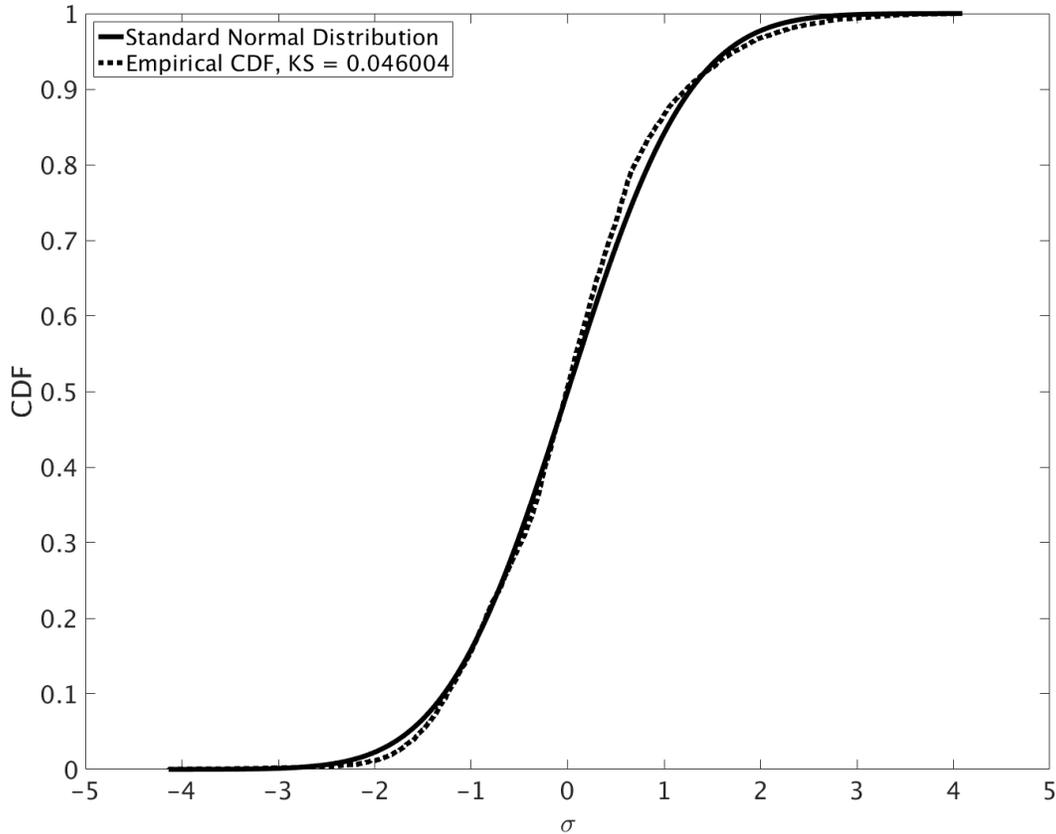


Figure 15. Empirical CDF of expanded validation data normalized by GP mean and standard deviation compared with normal CDF for prediction ROM of basis 75 trained generated with first 100 FOM solutions.

What this lower bound is varies with the number of ROM training points, m , and the prediction ROM basis dimension, d_p . The KS statistic tend to generally decrease with increasing m and d_p as shown in Fig 16.

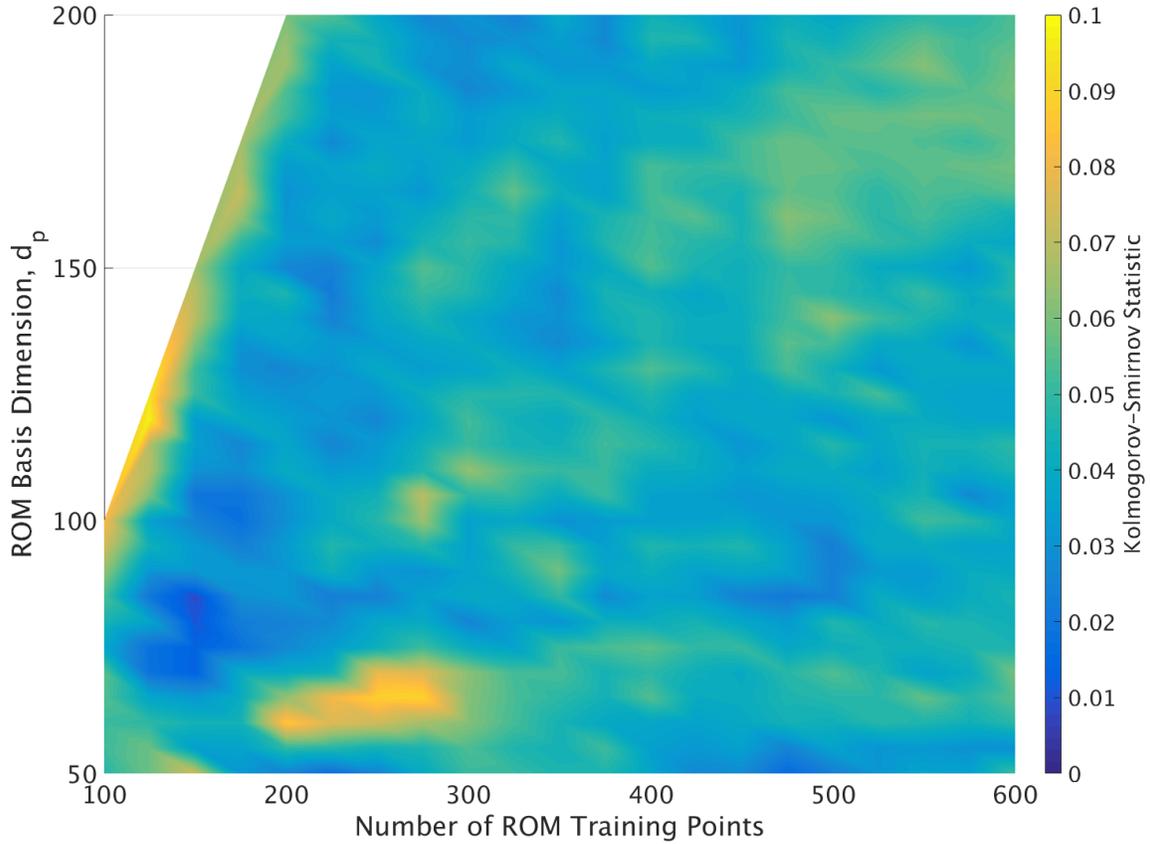


Figure 16. Variation of Kolmogorov-Smirnov Statistic with number of ROM training points and ROM basis dimension.

This is consistent with the error behavior becoming more Gaussian as the quality of the prediction ROM improves. The one exception to this rule is that the KS statistic tends to increase significantly as $d_p \rightarrow m$. This is likely due to the prediction ROM becoming too closely linked to the specific ROM training points used.

4.2. Decoupled with Variable Truncation

In this case, the first limitation of the decoupled approach is addressed. Additional ROMES training couples are generated by varying the training ROM dimension. The prediction ROM dimension will be fixed at 75. The training ROMs will be of dimension 73, 74, and 75. This triples the number of ROMES training couples.

This results in a slight increase in the Kolmogorov-Smirnov statistic from 0.1063 to 0.1197. This is a marginal reduction in accuracy. There is a trade-off here. If the original training set available by the decoupled method examined is too small, the ROMES may be improved by enlarging that set to include data from training ROMs of different dimensions. However, if the

training set is already sufficiently large to inform the ROMES, the additional training couples serve only to pollute the training set and result in a net reduction in accuracy.

In this case, the effect is small but due to the already large size of the training set (100 training couples) the net effect is negative. The effect is primarily due to including the training ROM of dimension 73. Most of the couples from the training ROM of dimension 74 are nearly coincident with couples from either 75 or 73 as shown in Fig 17.

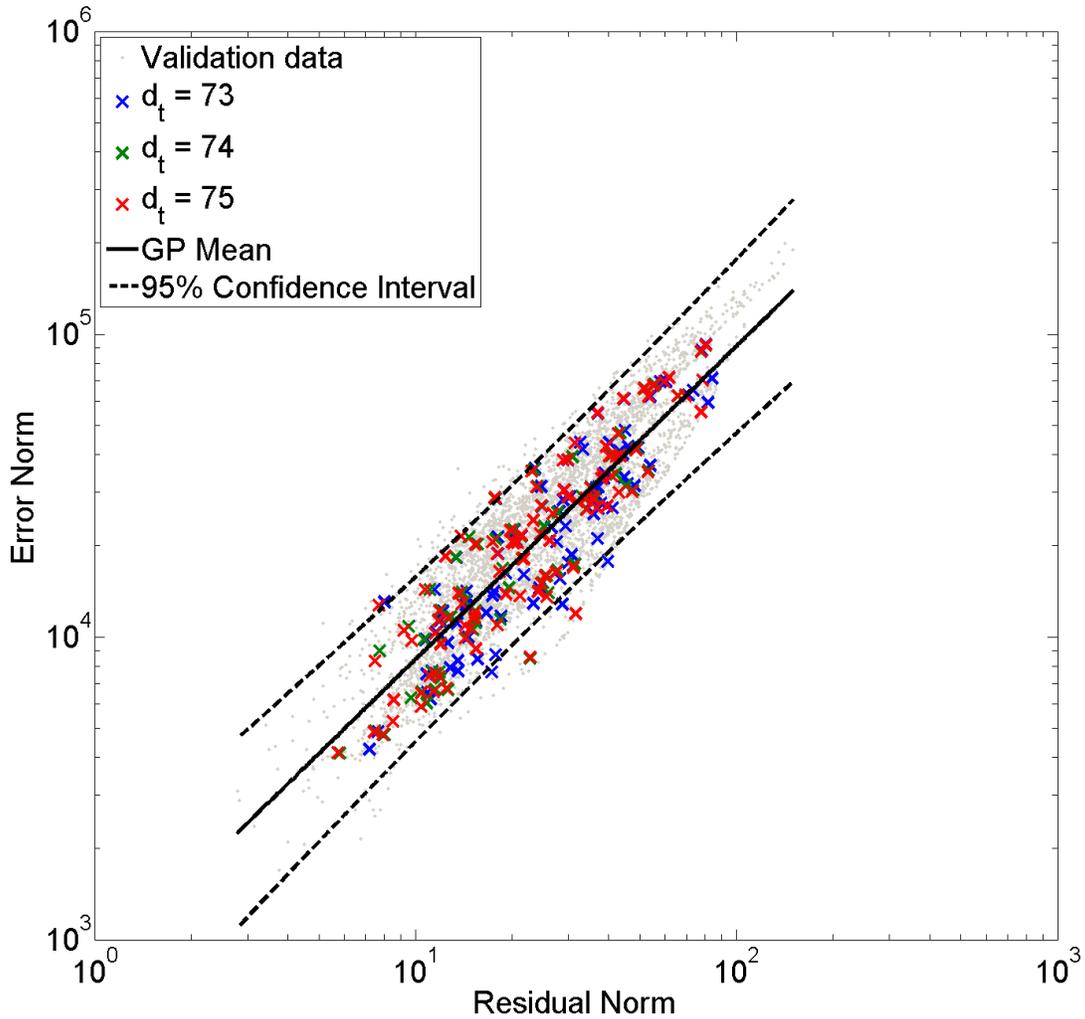


Figure 17. ROMES training couples generated from training ROMs of varying dimension compared with validation data for prediction ROM of dimension 75.

This approach shows promise, but like the decoupled approach, the prediction ROM is less accurate than is desirable given the number of available data points. This issue is addressed by some of the other approaches.

4.3. Concurrent Training Sets

Use all data from the training set to inform both the prediction and training ROM(s). In this case, there is only one training ROM and the prediction ROM and the training ROM are equivalent (since the basis dimension for both is fixed to the same value) and may be referred to simply as “the ROM”. The ROM is evaluated at every point in the training set and the residual norm / error norm pairs are recorded and used to train the ROMES. This provides 1 ROMES training couple for every point in the training set (a total of 200).

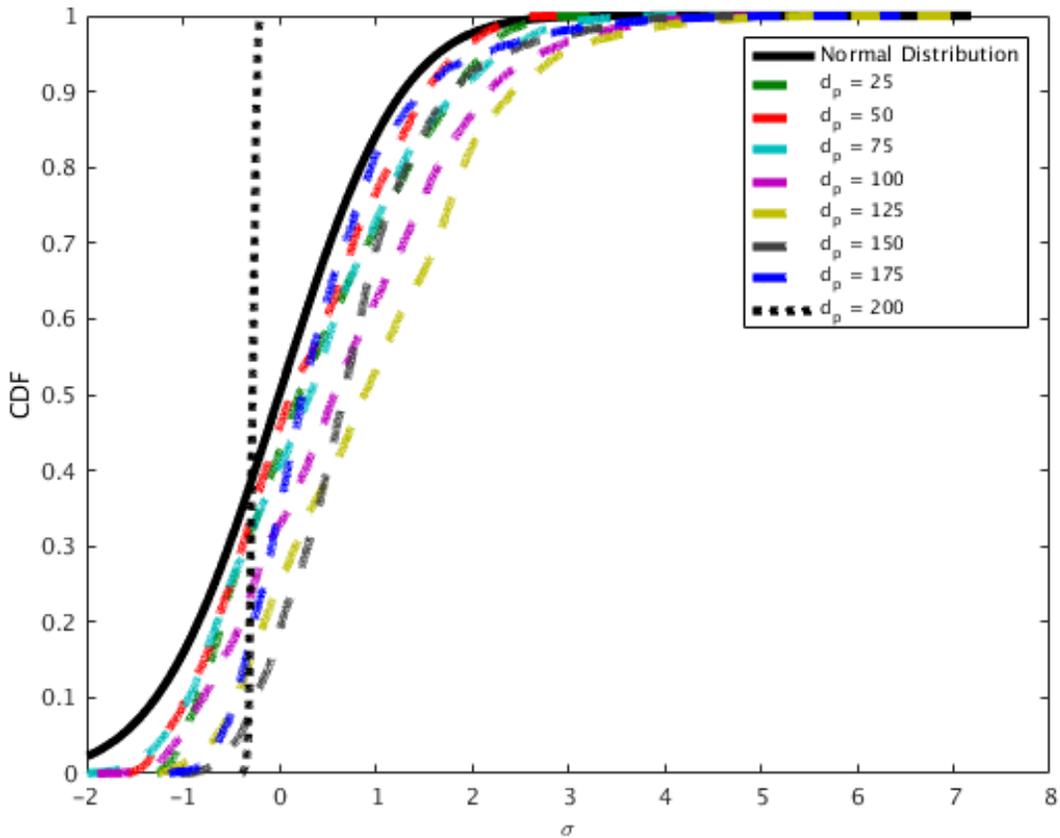


Figure 18. Empirical CDFs of validation data for ROMES generated using concurrent training sets normalized by the GP mean and standard deviation.

This process works reasonably well for highly truncated ROMs but begins to perform poorly for ROMs with basis dimension more than about 100 as seen in Figs 18 and 19. The requirement for this method to work appears to be that the ROM must do a poor enough job reproducing its own training data that the error at the training points is comparable to the error at the validation points.

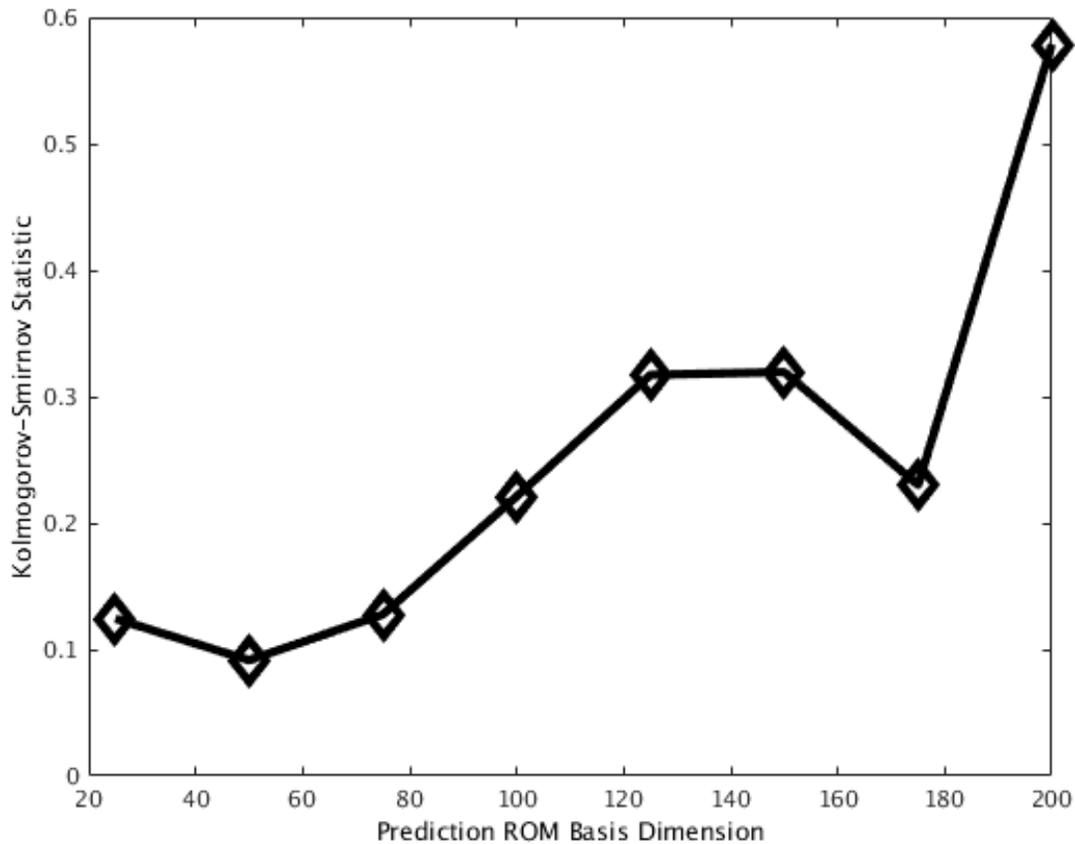


Figure 19. Variation of Kolmogorov-Smirnov statistic with ROM basis dimension for ROMES generated using concurrent training sets.

In these examples, the GP is not expected to perfectly match the true distribution and so the Kolmogorov-Smirnov statistic is expected to converge to some constant value as the size of the validation set approaches infinity. The 4000 points in the validation set are enough that the K-S statistic plotted in Fig 19 accurately represents the difference between the GP distribution and the true distribution.

4.4. Concurrent Training Sets with Variable Truncation

Use all data from the training set to inform both the prediction and training ROMs. The basis dimension of the prediction ROM is fixed. The basis dimension of the training ROMs is varied. This provides 200 ROMES training couples for every basis dimension considered.

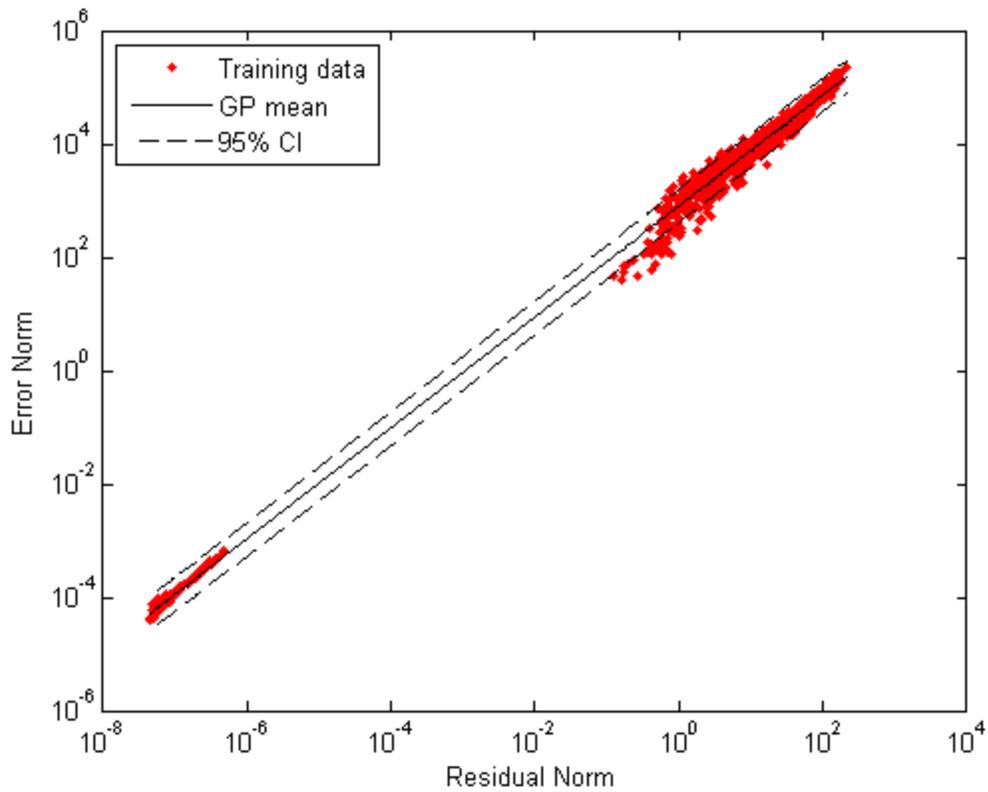


Figure 20. ROMES generated using concurrent training sets with variable truncation along with the ROMES training couples used.

Training ROMs of basis dimension 50, 75, 100, 125, 150, 175, and 200 were considered resulting in the 14,000 ROMES training couples plotted in Fig 20. ROMs of basis dimension less than 50 were not included. The ROMES generated from these 14,000 couples is also shown. This ROMES model may then be applied to the prediction ROM. Fig 21 shows the ROMES model overlaid onto the validation data for various prediction ROM dimensions. Notice the significant bias error introduced as d_p increases.

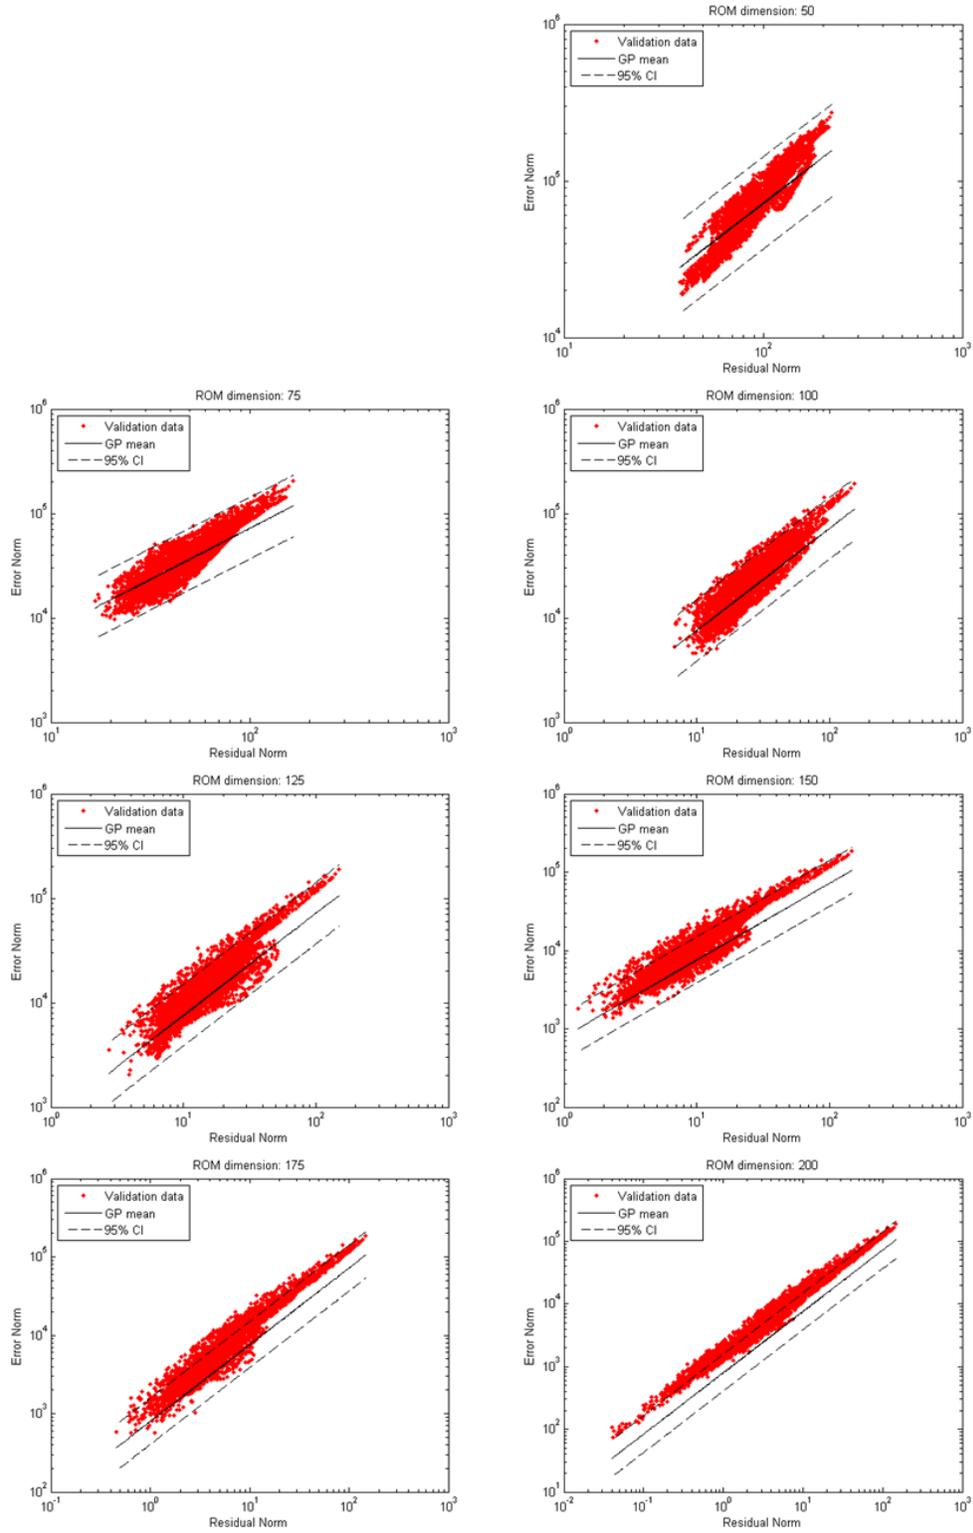


Figure 21. ROMES generated using concurrent training sets with variable truncation along with the ROMES training couples used.

Fig 22 shows the CDFs and corresponding KS statistics for this method. This method works best for highly truncated ROMs but does not work particularly well in any case. Note especially the large bias introduced at larger prediction ROM dimensions. This is born out in the large KS statistic values in Fig 22b.

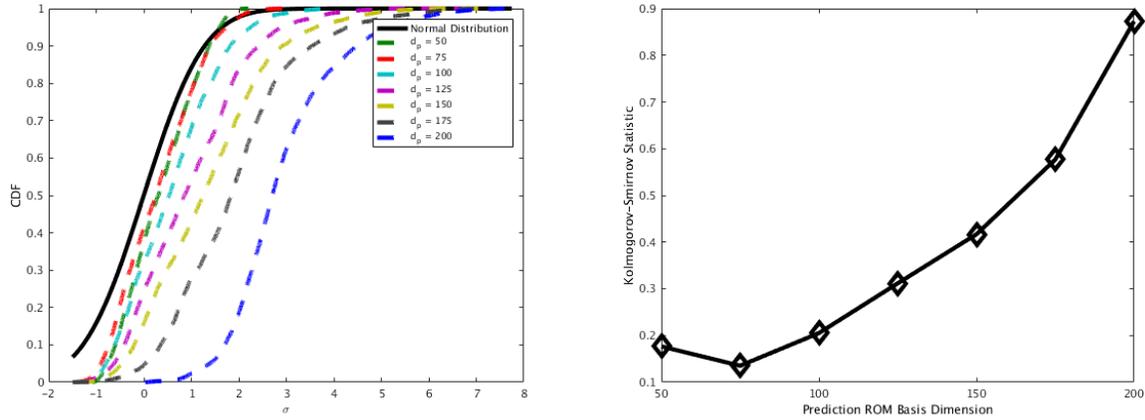


Figure 22. (a) Empirical CDFs normalized by the GP mean and standard deviation and (b) Kolmogorov-Smirnov statistics for ROMES generated using concurrent training sets with variable truncation.

The use of concurrent training sets with variable truncation to generate ROMES training couples is not recommended. The distribution of training couples sampled by this method is not the same as – and for many cases does not even approximate – the distribution of errors generated by the prediction ROM.

4.5. k-fold Cross Validation

In the previous two methods all of the available data was used to inform the training ROM(s). It is expected that the ROMs will perform differently (better) at points within their training set than at points outside their training set. This caused a discrepancy between the ROMES training couples and the validation data resulting in poor behavior of the error surrogate. This is in contrast to the first 2 methods in which some of the available data is used to inform the training ROM(s) while other data is used to generated the ROMES training couples. This latter scenario was observed to result in a good approximation to the distribution of the validation data but often required an unfeasible number of FOM evaluations to provide a sufficient number of ROMES training couples. Furthermore, the decoupled methods are undesirable because a significant fraction of the costly FOM evaluations are used exclusively to train the ROMES. This yields a less accurate ROM than would otherwise be possible.

Both of these types of deficiencies are addressed through k-fold cross validation. The simplest implementation of this is Leave-One-Out Cross-Validation (LOOCV). In LOOCV, the available data ($n=200$ points) is split into two parts. $n-1$ points are used to construct a training ROM. The training ROM is then evaluated at the remaining point providing a ROMES training couple. This process is repeated n times, leaving a different data point out each time. This provides 1

ROMES training couple for every point in the training set (a total of 200). This produces good agreement between the resulting ROMES and the validation data as seen in Fig 23. The KS statistic for the empirical CDFs in Fig 23 range from 0.06 to 0.11. The KS statistic is below 0.08 for all basis dimensions besides $d_p = 25$. Initially, the training and prediction ROMs will be required to have the same basis dimension. In the next section this requirement will be relaxed.

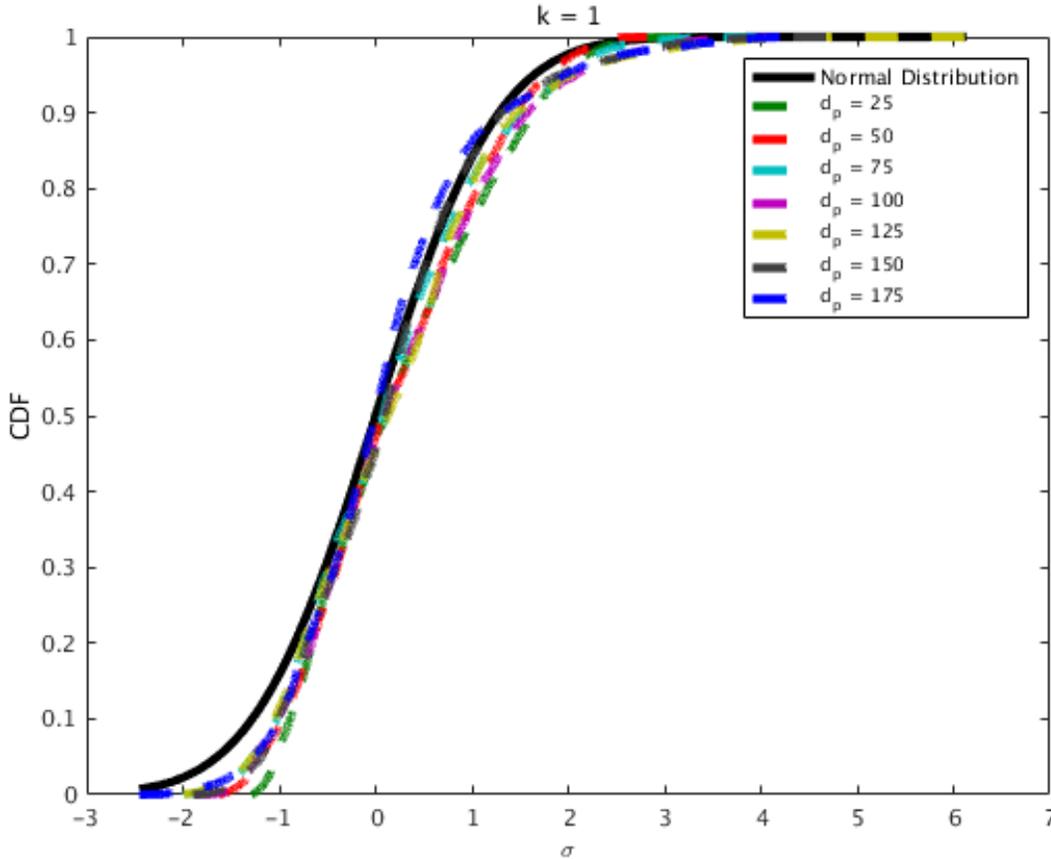


Figure 23. Empirical CDFs of validation data for ROMES generated using LOOCV normalized by the GP mean and standard deviation.

LOOCV is a special case of k -fold cross-validation. In k -fold CV the available data is divided into k groups or folds. All but one of the folds is used to construct the training ROM which is then evaluated on the k points of the other fold. This provides k ROMES training couples per training ROM. The same number of total training couples is generated regardless of k but fewer training ROMs must be constructed for larger values of k .

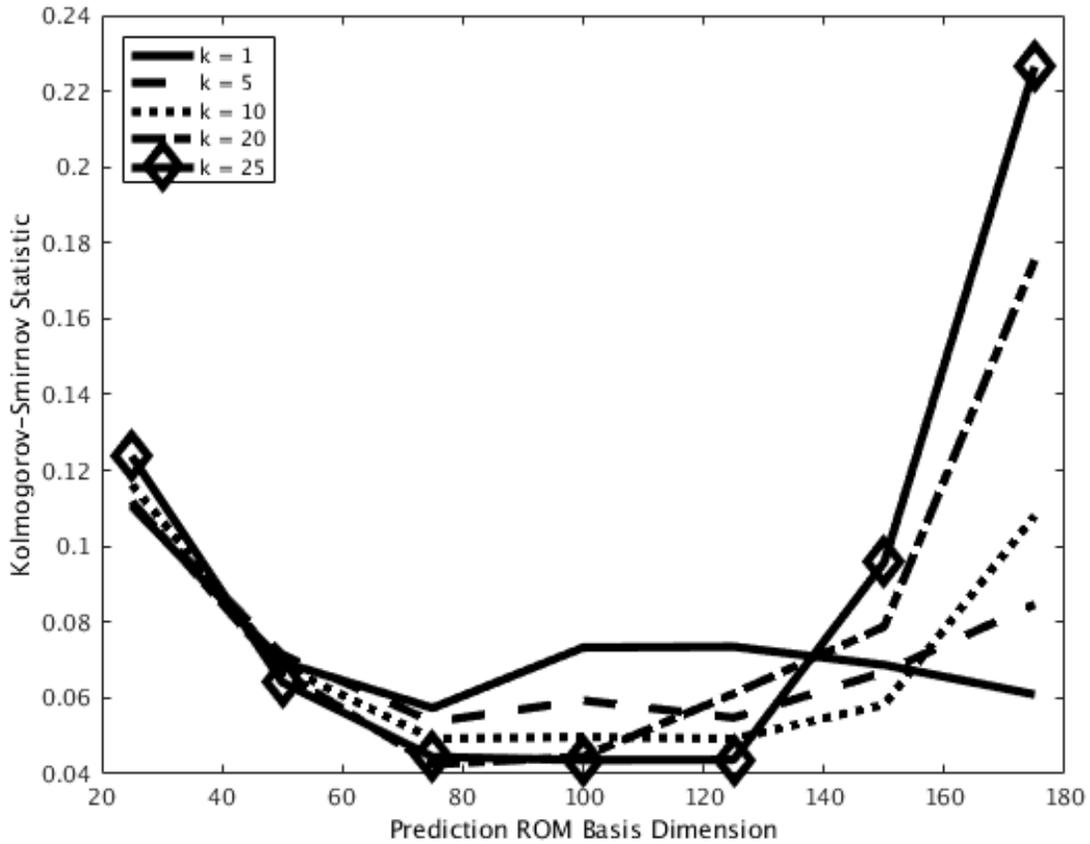


Figure 24. Kolmogorov-Smirnov statistics for ROMES generated using k-fold cross validation.

k-fold CV works well for a wide range of k and ROM dimension as shown in Fig 24. However, there are tradeoffs. As seen in Fig 16, the error behavior is less Gaussian for ROMs below a certain dimension. Outside of this limit, for cases where the ROM dimension is small relative to the number of sample points used to build it, larger values of k result in lower variance as the individual training ROMs are less correlated. However, for higher ROM dimensions, larger values of k yield a faster evaluation time at the cost of accuracy. In these cases, the training ROMs become increasingly sensitive to their individual training sets and less closely resemble the prediction ROM. What sort of accuracy is required and how much computational expense is affordable will be application dependent.

4.6. k-fold Cross Validation with Variable Truncation

Regardless of the value of k chosen, k-fold CV has the drawback of only producing 1 ROMES training couple for each available data point. This limits the total number of ROMES training couples available. This isn't an issue when a sufficiently large number of data points is available as in the case above but may be significant when smaller numbers of data points are available. We can seek to augment the number of ROMES training couples by varying the dimension of the

training ROM. The results in the following two tables were generated by taking the training ROM dimension, d_t to be three different values based on the prediction ROM dimension, d_p . $d_t \in \{d_p, d_p - 1, d_p - 2\}$. Two different values of k are used, 1 and 5.

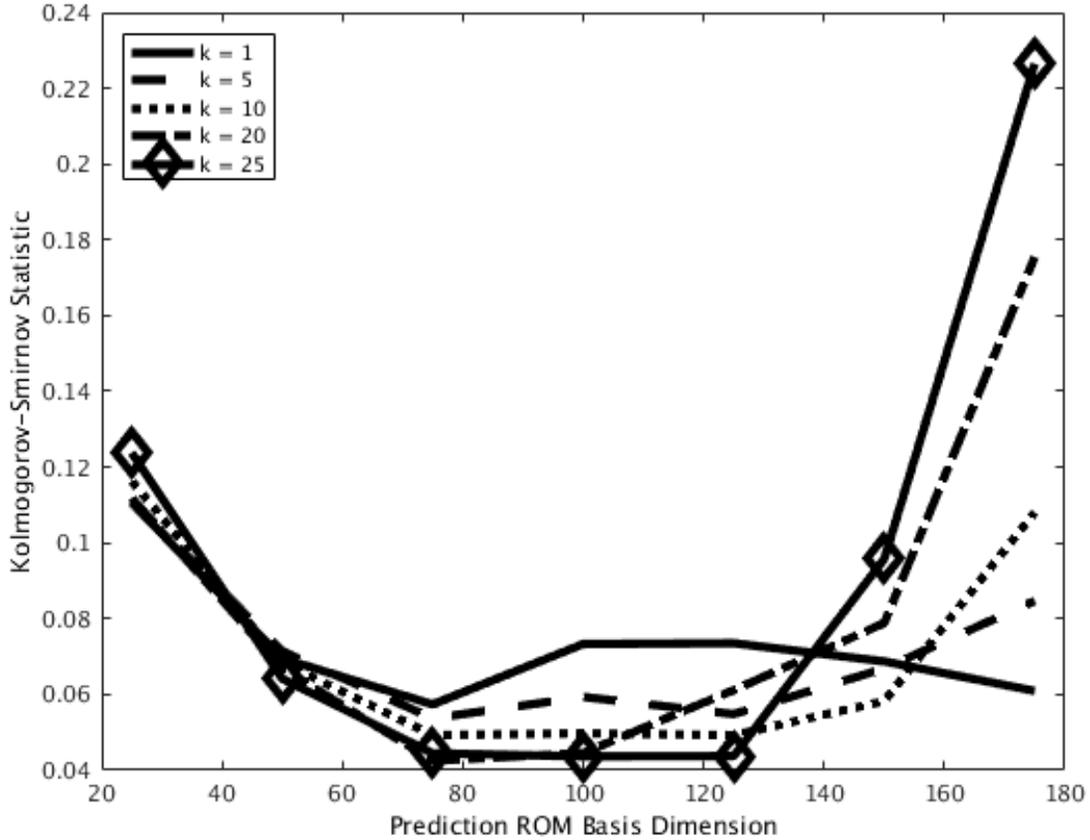


Figure 25. Kolmogorov-Smirnov statistics for ROMES generated using k-fold cross validation with variable truncation with 200 FOM evaluations.

This works reasonably well. There is a little degradation of performance relative to the case where $d_t=d_p$ for some combinations of k and d_t but also some improvement for other combinations. Overall, the results in Fig 24 and 25 are very similar. However, the true test is if this provides a benefit when the number of available data points is small. Next, we limit ourselves to LOOCV ($k=1$) and only use the first 50 available data points. We will examine prediction ROMs of basis dimension $d_p=25, 35,$ and 45 .

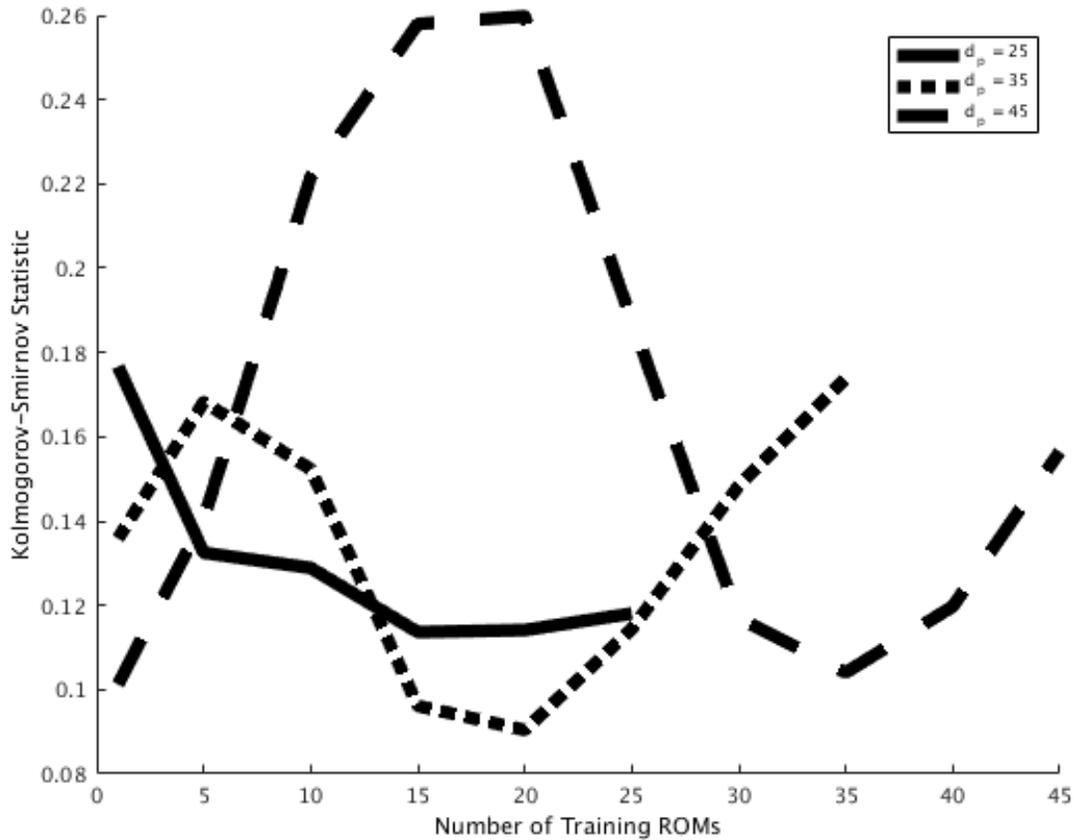


Figure 26. Kolmogorov-Smirnov statistics for ROMES generated using k-fold cross validation with variable truncation with 50 FOM evaluations.

The results in Fig 26 show that for this few data points it is possible to improve the solution by utilizing multiple training ROMs. However, more is not always better. Training ROMs below a certain dimension (about 15 for this data set) appear to behave differently enough from higher dimensional ROMs that including these couples degrades the accuracy of the ROMES. An example of this is in Figs 27 and 28.

Consider the training couples generated using higher order training ROMs. They all agree reasonably well with the validation data as shown in Figure 27 and the ROMES prediction accuracy increases as more ROMs within this range are included. However, the ROMES prediction accuracy is diminished if too many low order training ROMs are included. This is because training ROMs below a certain order do not approximate the behavior of the prediction ROM. The point at which this occurs is dependent upon the prediction ROM dimension as well as the number of data points available.

The training couples generated from training ROMs of dimension 5, 10, and 15 bear no resemblance to the validation data. However, some of the couples from the training ROMs of dimension 20 and 25 (especially those within the range of the residual norm for which the

ROMES will be evaluated) likely contain useful information. This is reflected in the fact that the best ROMES performance for $d_p = 45$ is achieved by including couples from 30 training ROMs of dimension 16-45.

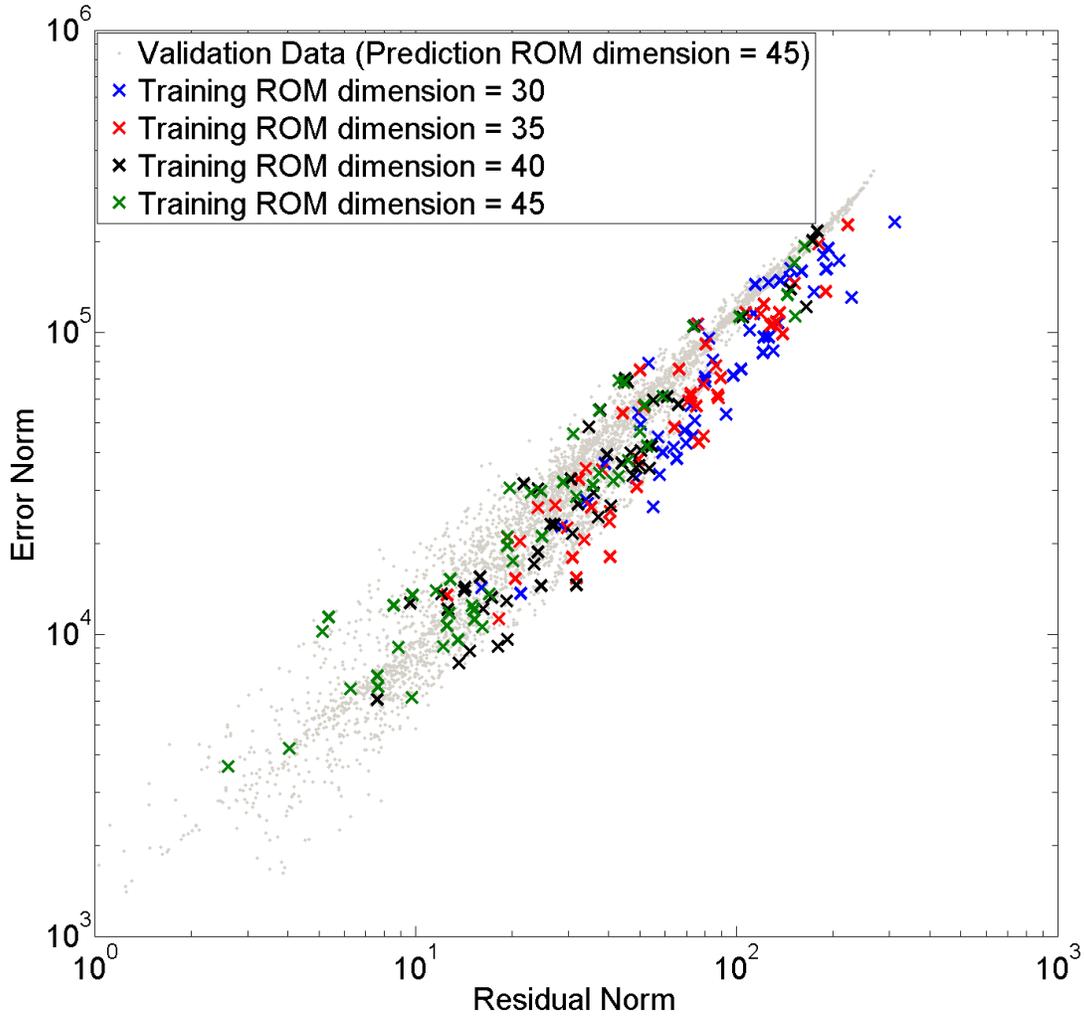


Figure 27. Training couples compared with validation data shows good overall agreement for sufficiently large training ROM dimensions.

Furthermore, it is noted that exploring every possible ROM dimension is expensive and generates a much larger number of training couples than is reasonably required. Many of these training couples are redundant. The information yielded by a training ROM of dimension 45 and a training ROM of dimension 44 are very similar. This redundancy may be reduced by using only a subset of the available training ROMs. Intervals of 5 are observed to work well.

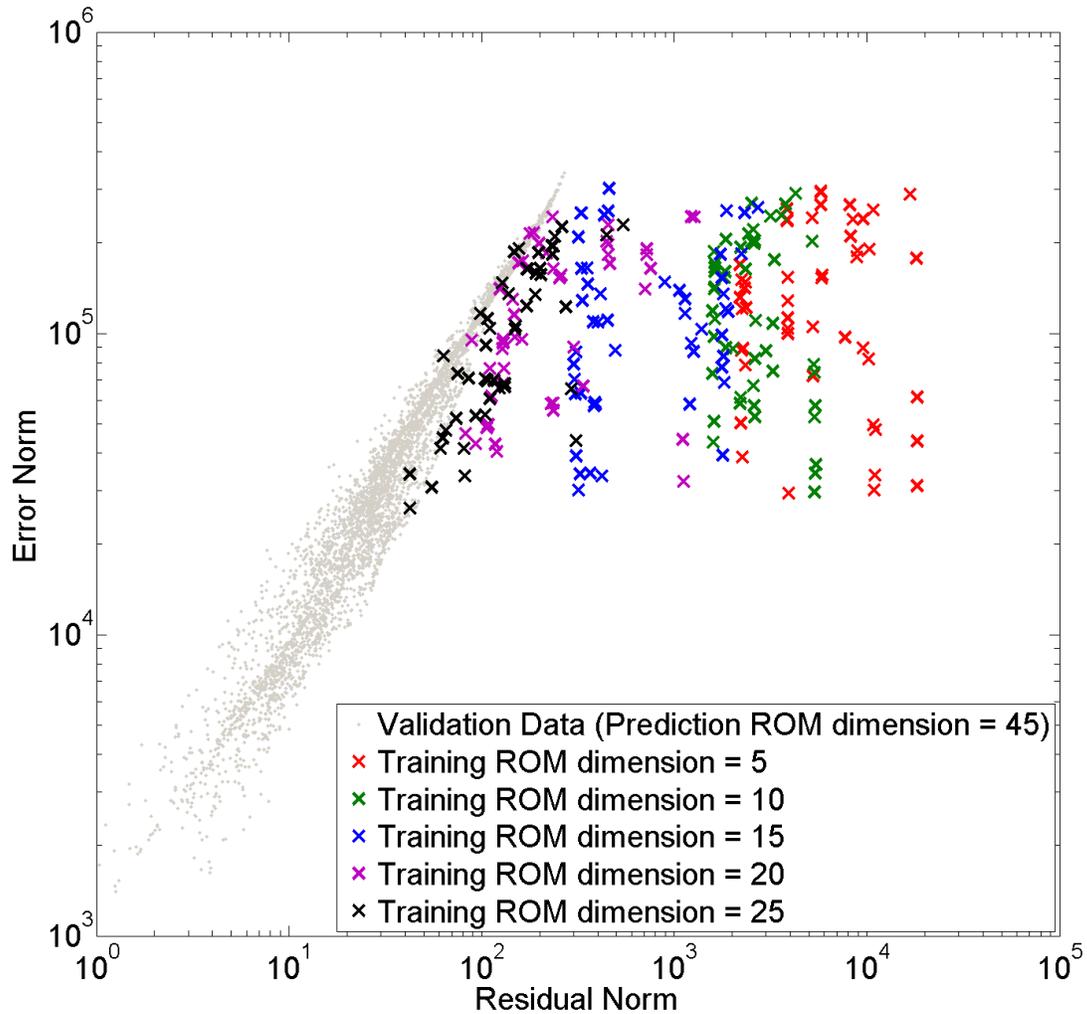


Figure 28. Training couples compared with validation data shows poor agreement for low dimensional training ROMs.

4.7. Fixed Dimension Greedy Search

We’ve found that an efficient way to sample the angular space is through a greedy search. In this process, a small number of samples are initially taken to generate a ROM. Additional samples are then taken to refine the ROM basis one at a time. Each successive sample is chosen such that it maximizes the residual norm of the ROM generated using all previous samples. At each step in the process, the ROM residual norm (for the previous iteration’s ROM) and the corresponding error are known. This generates a series of potential ROMES training couples.

In this case, the ROM dimension (both training and prediction) is fixed. The same 200 sample points used throughout this document are reused here. The first d_p sample points are used for training the initial ROM. The remaining $200 - d_p$ sample points are then added to refine the

ROM basis in the order given by the greedy search algorithm producing $200 - d_p$ ROMES training couples. The ROMES generated from these training couples is then validated against the validation data. The result is shown in Fig 29.

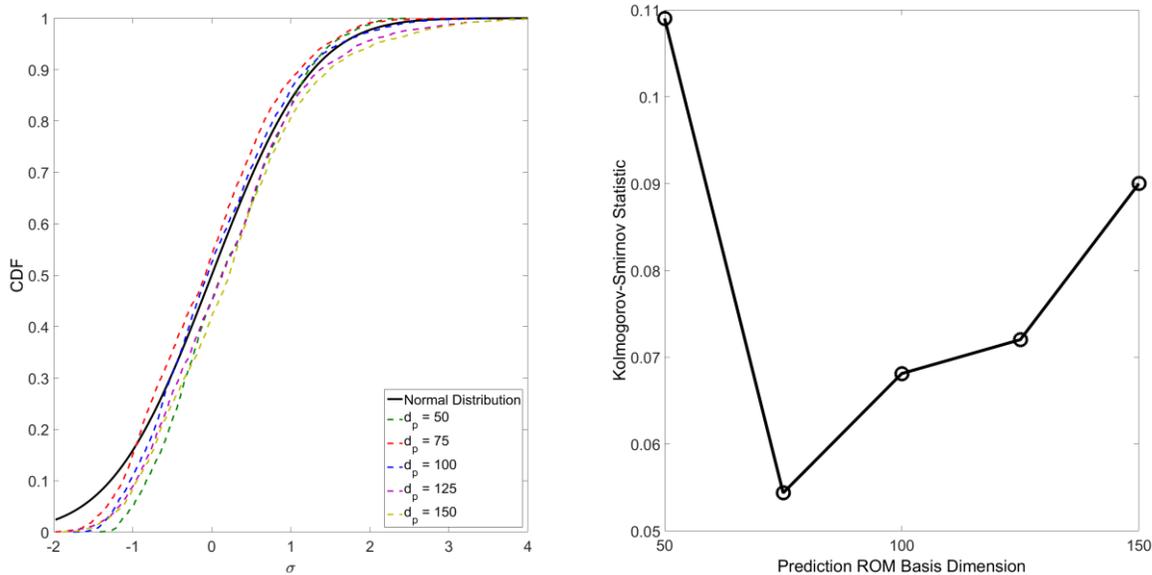


Figure 29. (a) Empirical CDFs normalized by the GP mean and standard deviation and (b) Kolmogorov-Smirnov statistics for ROMES generated using a fixed dimension greedy search.

This seems to work somewhat, especially for moderate ROM basis dimensions, but it has a couple of significant drawbacks. 1) The ROM basis dimension must be specified beforehand without any knowledge of the true dimensionality of the problem and 2) For a fixed number of training points, there is a tradeoff between the ROM dimension and the number of available ROMES training couples. The training couples closely match the validation data for the larger ROM dimensions but the small number of training couples limits the ROMES accuracy.

4.8. Lagged Dimension Greedy Search

We attempt to increase the number of available ROMES training couples by varying the dimension of the training ROM. In this case, we triple the number of available ROMES training couples by using $d_t \in \{d_p, d_p - 5, d_p - 10\}$.

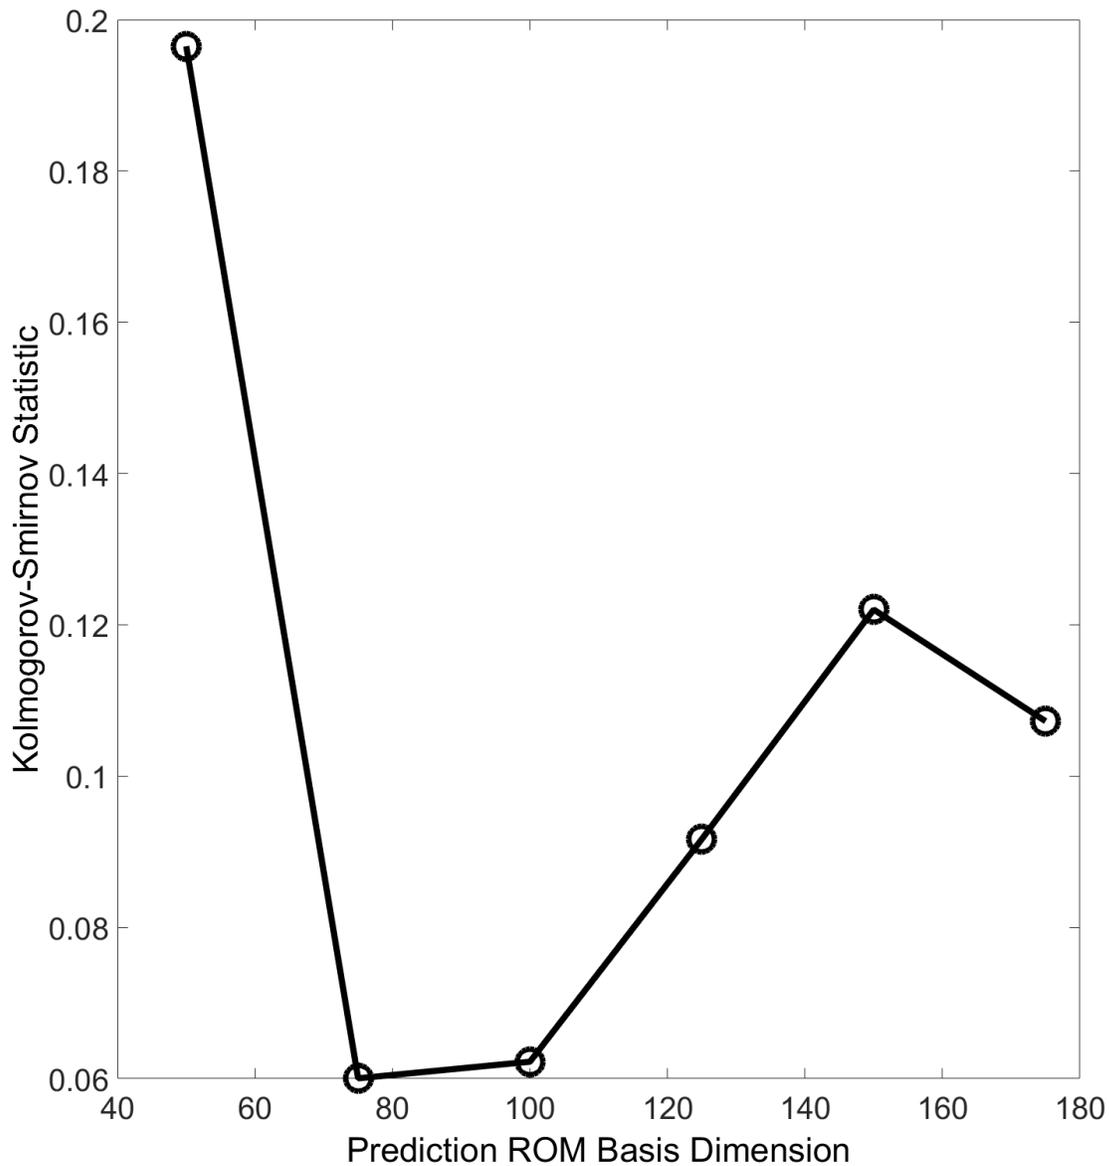


Figure 30. Kolmogorov-Smirnov statistics for ROMES generated using a lagged dimension greedy search.

This improves the performance relative to the results of the fixed dimension greedy search for prediction ROMs of basis dimension 75 and 100 but degrades performance for the cases of either high or low prediction ROM basis dimension. For the prediction ROM of dimension 50 two things are true. 1) The 150 ROMES training couples found using only the training ROM of the same dimension as the prediction ROM is already sufficient to train the ROMES and 2) The ROMES training couples generated with the training ROM of dimension 40 do not necessarily all follow the same distribution as the training couples generated with higher dimensional

training ROMs causing the resulting ROMES to be overly conservative. This results in the large decrease in accuracy for the ROM basis dimension seen in Fig 30. For large ROM basis dimension, the accuracy is again seen to be a trade-off between an increased number of ROMES training couples and a decreased average quality of those training couples.

4.9. Constant Energy Greedy Search

With the constant energy greedy search, we tackle the first problem with the fixed dimension greedy search: the requirement that the prediction ROM dimension be specified before the start of the simulation. In this approach, the ROM dimension is allowed to vary. At each step in the greedy search process, the ROM dimension is set such that the ROM captures 99.999% of the statistical energy. This causes the ROM dimension to increase with the number of samples taken until it reaches a plateau due to the underlying dimensionality of the problem.

The greedy search algorithm is seeded with a number of initial samples, randomly chosen, to generate the first ROM. Figure 31 shows the impact of varying this number of initial samples. Reducing the number of initial samples means more steps in the greedy search algorithm before the 200 available data points are exhausted and more ROMES training couples generated. Unfortunately, these additional ROMES training couples are necessarily generated with a low-dimensional training ROM. This could potentially lead to problems as we have seen. However, the results below are not bad, comparable to or better than the lagged dimension greedy search in most cases.

The explanation for this behavior is that although the low-dimensional ROMs perform differently than the prediction ROM, these differences are most pronounced at smaller values of the residual. The constant energy greedy search approach only includes as a ROMES training couple the point with the maximum residual value.

The constant energy greedy search approach has one additional advantage. It is the cheapest of all the methods to evaluate as the ROMES training couples may be generated online at no additional cost. The ROMES training couples used by the constant energy greedy search are the same as those used by the previously described deterministic approach.

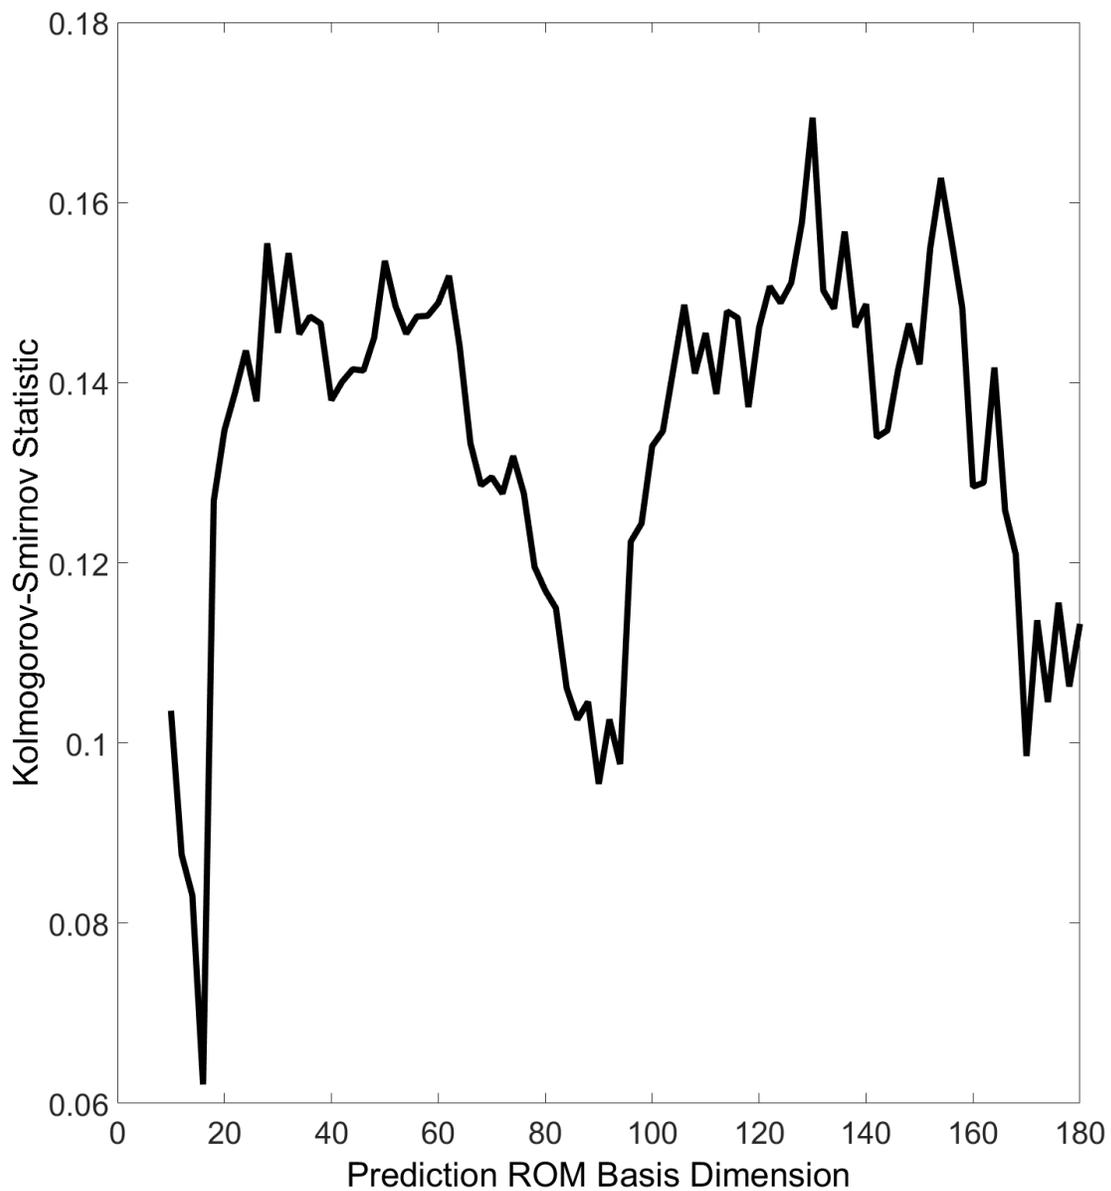


Figure 31. Kolmogorov-Smirnov statistics for ROMES generated using a constant energy greedy search.

5. CONCLUSIONS

The discrete ordinates method is shown to be amenable for reduced order modeling. For 1D and 2D problems, any sufficiently large sampling of ordinate directions results in adequate sampling of the angular domain to construct a highly accurate reduced order model. For 3D problems more care must be taken to sample the angular domain efficiently. Efficient sampling also results in benefits in the 1D and 2D cases. A greedy sampling approach is proposed using the ROM residual as an error indicator. Timing studies show this to be highly effective relative to successively increasing the angular quadrature order. The advantages of this approach increase with increasing mesh resolution since the cost of additional ROM evaluations is only weakly related to the mesh size (unlike the cost of additional FOM evaluations). Additionally, an error estimate is proposed that closely matches the actual error for the problems considered. This error estimate is inexpensive to generate and evaluate and may be easily used as a stopping criteria for the greedy search algorithm.

Although the deterministic error model presented performs well in predicting the mean behavior, additional utility is possible through the use of stochastic error models which provide a distribution of error estimates. Stochastic ROMES models may be effectively trained using the same FOM solutions used to generate the prediction ROM. This may either be done online as part of the simulation or offline at a later date. ROMES training methods vary in accuracy and computational cost. Nine potential options were examined with k-fold cross validation providing the most robustly accurate results at a reasonable computational cost. The constant energy greedy search result which uses does not require the retention of any FOM solutions or any subsequent ROM evaluations is the most computationally expedient. However, this approach tends to produce a lower quality ROMES than would be possible through k-fold cross validation. The use of concurrent training sets to generate ROMES data by varying the training ROM dimension was observed to result in large bias errors and an inaccurate ROMES.

The utility of the parametric model reduction approach has been demonstrated for the discrete ordinate method for accurately simulating systems involving PMR. The next logical step is implementation in the SIERRA suite of codes.

6. REFERENCES

- [1] B. G. Carlson, "Mechanical Quadrature and the Transport Equation," Los Alamos Scientific Laboratory, Los Alamos, NM, 1961.
- [2] B. G. Carlson, "Transport Theory: Discrete Ordinates Quadrature over the Unit Sphere," Los Alamos Scientific Laboratory, Los Alamos, NM, 1970.
- [3] S. H. Kim and K. Y. Huh, "A New Angular Discretization Scheme of the Finite Volume Method for 3-D Radiative Heat Transfer in Absorbing, Emitting and Anisotropically Scattering Media," *International Journal of Heat and Mass Transfer*, vol. 43, pp. 1233-1242, 2000.
- [4] R. Koch and R. Becker, "Evaluation of Quadrature Schemes for the Discrete Ordinates Method," *J. Quant. Spectrosc. & Radiat. Transfer*, vol. 84, pp. 423-435, 2004.
- [5] R. Koch, W. Krebs, S. Wittig and R. Viskanta, "Discrete Ordinates Quadrature Schemes for Multidimensional Radiative Transfer," *J. Quant. Spectrosc. & Radiat. Transfer*, vol. 53, no. 4, pp. 252-372, 1995.
- [6] K. D. Lathrop and B. G. Carlson, "Discrete Ordinates Angular Quadrature of the Neutron Transport Equation," Los Alamos Scientific Laboratory, Los Alamos, NM, 1965.
- [7] E. E. Lewis and W. F. Miller, Jr., *Computational Methods of Neutron Transport*, La Grange Park, Illinois: American Nuclear Society, Inc., 1993.
- [8] B. W. Li, J. H. Zhou, X. Y. Cao, H. G. Chen and K. F. Cen, "The Spherical Surface Symmetrical Equal Dividing Angular Quadrature Scheme for Discrete Ordinates Method," *ASME Journal of Heat Transfer*, vol. 124, no. 3, pp. 482-490, 2002.
- [9] G. Longoni, *Advanced Quadrature Sets, Acceleration and Preconditioning Techniques for the Discrete Ordinates Method in Parallel Computing Environments*, Gainesville, FL: University of Florida, 2004.
- [10] G. Longoni and A. Haghghat, "Development and Application of the Regional Angular Refinement Technique and its Applications to Non-Conventional Problems," in *Proc. PHYSOR*, Seoul, Korea, 2002.
- [11] S. A. Rukolaine and V. S. Yuferev, "Discrete Ordinates Quadrature Schemes Based on the Angular Interpolation of Radiation Intensity," *J. Quant. Spectrosc. & Radiat. Transfer*, vol. 69, pp. 257-275, 2001.
- [12] C. Thurgood, A. Pollard and P. Rubini, "Development of TN Quadrature Sets and HEART Solution Method for Calculating Radiative Heat Transfer," in *Int. Symp. on Steel Reheat*

Furnace Technology, Hamilton, Ontario, Canada, 1990.

- [13] M. L. Adams and E. W. Larsen, "Fast Iterative Methods for Discrete-Ordinates Particle Transport Calculations," *Progress in Nuclear Energy*, vol. 40, no. 1, pp. 3-159, 2002.
- [14] P. Benner, S. Gugercin and K. Willcox, "A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems," *SIAM Review*, 2015.
- [15] D. Ryckelynck, "Hyper-Reduction of Mechanical Models Involving Internal Variables," *Int. J. Numer. Meth. Engng*, vol. 77, pp. 75-89, 2009.
- [16] D. Amsallem, M. Zahr, Y. Choi and C. Farhat, "Design Optimization Using Hyper-Reduced-Order Models," *Structural and Multidisciplinary Optimization*, vol. 51, pp. 919-940, 2015.
- [17] M. Grepl, "Reduced-Basis Approximation and A Posteriori Error Estimation for Parabolic Partial Differential Equations," Ph. D. thesis, MIT, Cambridge, MA, 2005.
- [18] M. Grepl and A. Patera, "A Posteriori Error Bounds for Reduced-Basis Approximations of Parametrized Parabolic Partial Differential Equations," *ESAIM Math. Modelling Numer. Anal.*, vol. 39, pp. 157-181, 2005.
- [19] K. Veroy, C. Prud'homme, D.V. Rovas, and A.T. Patera, "A Posteriori Error Bounds for Reduced-Basis Approximation of Parametrized Noncoercive and Nonlinear Elliptic Partial Differential Equations," in *Proceedings of the 16th AIAA Computational Fluid Dynamics Conference*, Orlando, FL, 2003.
- [20] T. Bui-Thanh, K. Willcox, and O. Ghattas, "Model Reduction for Large-Scale Systems with High-Dimensional Parametric Input Space," *SAIM J. Sci. Comput.*, vol. 30, no. 6, pp. 3270-3288, 2008.
- [21] J. Stone and M. Adams, "Adaptive Discrete-Ordinates Algorithms and Strategies," in *Mathematics and Computation, Supercomputing, Reactor Physics and Nuclear and Biological Applications*, Palais des Papes, Avignon, France, 2005.
- [22] D. Koeze, "Goal-Oriented Angular Adaptive Neutron Transport using a Spherical Discontinuous Galerkin Method," M.S. thesis, Delf University of Technology, Netherlands, 2011.
- [23] J. Stone, "Adaptive Discrete-Ordinates Algorithms and Strategies," Ph.D. thesis, Texas A&M University, 2007.
- [24] J. Stone and M. Adams, "Progress on Adaptive Discrete-Ordinates Algorithms and Strategies," in *Nuclear Mathematical and Computational Sciences: A Century in Review, A Century Anew*, Gatlinburg, Tennessee, 2003.

- [25] G. Longoni and A. Haghghat, "Development and Application of the Regional Angular Refinement Technique and its Applications to Non-Conventional Problems," in *Proc. PHYSOR 2002*, Seoul, Korea, 2002.
- [26] P. N. Brown and B. Change, "Locally Refined Quadrature Rules for SN Transport," Report UCRL-JRNL-220755, Lawrence Livermore National Laboratory, 2006.
- [27] G. Longoni and A. Haghghat, "Development of New Quadrature Sets with the Ordinate Splitting Technique," in *Proceedings of the 2001 ANS International Meeting on Mathematical Methods for Nuclear Applications*, Salt Lake City, UT, 2001.
- [28] J. Tencer, "A Comparison of Angular Discretization Techniques for the Radiative Transport Equation," in *Proceedings of the ASME 2015 International Mechanical Engineering Congress & Exposition*, Houston, TX, 2015.
- [29] K. Carlberg, C. Bou-Mosleh, and C. Farhat. "Efficient non-linear model reduction via a least-squares Petrov–Galerkin projection and compressive tensor approximations." *International Journal for Numerical Methods in Engineering* 86, no. 2 (2011): 155-181.
- [30] M. Drohmann and K. Carlberg. "The ROMES method for statistical modeling of reduced-order-model error." *SIAM/ASA Journal on Uncertainty Quantification* 3, no. 1 (2015): 116-145.

DISTRIBUTION

1	MS0359	D. Chavez, LDRD Office	1911
1	MS0825	Roy Hogan	1514
1	MS0828	Flint Pierce	1541
1	MS0836	Marvin Larsen	1514
1	MS0840	John Tencer	1514
1	MS0899	Technical Library	9536 (electronic copy)
1	MS9159	Kevin Carlberg	8959



Sandia National Laboratories