

# Fast Neural Network Predictions from Constrained Aerodynamics Datasets

Cristina White\*

*Stanford University, Stanford, California, 94305*

Daniela Ushizima†

*Lawrence Berkeley National Laboratory, Berkeley, California, 94720*

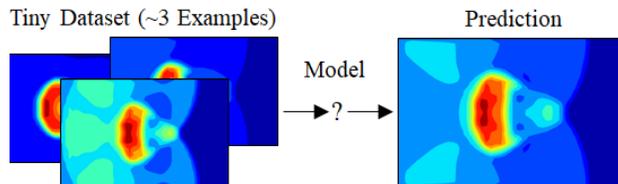
Charbel Farhat‡

*Stanford University, Stanford, California, 94305*

**Incorporating computational fluid dynamics in the design process of jets, spacecraft, or gas turbine engines is often challenged by the required computational resources and simulation time, which depend on the chosen physics-based computational models and grid resolutions. An ongoing problem in the field is how to simulate these systems faster but with sufficient accuracy. While many approaches involve simplified models of the underlying physics, others are model-free and make predictions based only on existing simulation data. We present a novel model-free approach in which we reformulate the simulation problem to effectively increase the size of constrained pre-computed datasets and introduce a novel neural network architecture (called a cluster network) with an inductive bias well-suited to highly nonlinear computational fluid dynamics solutions. Compared to the state-of-the-art in model-based approximations, we show that our approach is nearly as accurate, an order of magnitude faster, and easier to apply. Furthermore, we show that our method outperforms other model-free approaches.**

## I. Introduction

In machine learning, models are data-driven, whereas in computational fluid dynamics (CFD), models are physics-based and well-defined by a set of partial differential equations. However, the solution—the variation in space and time of fluid state variables such as density, velocity, momentum, and energy—is unknown until a simulation of the model has converged. Automatically iterating through new designs for optimization is often infeasible and datasets of previous simulations are highly constrained in terms of the number of samples because a single simulation can take days or weeks for a single design. We explore whether our model-free method can enable engineers to make use of these constrained datasets to generate accurate approximations to CFD solutions in seconds, given the information they already have from the simulations they have already run.



**Fig. 1** The goal is to automatically build data-driven models that use limited experience to make predictions.

Reduced-order models (ROMs) are currently the state-of-the-art model-based methods that address this problem; however, they tend to be highly intrusive. The problem has also been addressed in industry using approximations for

\*Research Assistant, Department of Mechanical Engineering, William F. Durand Building, Room 257, 496 Lomita Mall, Stanford University, Stanford, CA 94305.

†Staff Scientist, Computational Research Division, Lawrence Berkeley National Lab, 1 Cyclotron Rd, Mail Stop 59, Berkeley, CA 94720.

‡Chair, Department of Aeronautics and Astronautics, William F. Durand Building, Room 257, 496 Lomita Mall, Stanford University, Stanford, CA 94305, AIAA Member.

scalar quantities of interest based on surface fitting techniques: these are fast, but can be inaccurate. Such approaches have also been investigated in combination with neural networks with the aim of improving speed or accuracy, but usually have the same drawbacks as the original approach with which they were hybridized.

Here we address the problem by training stand-alone neural networks to learn the functions that approximate complete solutions instead of quantities of interest. For two different CFD test cases, we reformulate the problems to make more efficient use of the data, and we show that fully connected feed-forward networks perform reasonably well, with some room for improvement. We introduce a new neural network architecture, the cluster network, with paired function and context networks, which addresses the limitations of the fully connected network. The cluster network has a stronger inductive bias—that the solutions we are approximating are made up of a small number of local, simple functions. For a one-dimensional but shock-dominated problem, we show that the cluster network delivers an accuracy comparable to that of a ROM approach based on a global reduced-order basis, but runs orders of magnitude faster online.

## II. Background and Related Work

Fluid dynamics introduces a variety of its own unique datasets and problems where machine learning can potentially apply. In the following sections, we map out the areas of intersection between CFD and machine learning and we delineate where our novel proposed method fits.

### A. Model-Based Methods

#### 1. Full Order Models

**CFD Solvers:** Industry CFD solvers simulate the Navier-Stokes equations—a set of partial differential equations governing the motion of fluids (both liquids and gases). High fidelity CFD solvers allow for deeper analysis and better aerodynamic design performance. Methods for improving the approximations and performance of the physical models have been investigated for decades with the goal of accelerating simulations. However, due to the high dimensionality required to resolve small scale physics, even the most highly optimized solver may take too long on a supercomputer, which hinders investigations in design and engineering applications. Therefore, machine learning methods for dramatically reducing the computational cost of obtaining solutions while preserving sufficient accuracy are an active and high impact area of research.

**NN Solvers:** Neural network (NN) architectures can be constructed to perform steps that are similar to a CFD solver after training on simulation data. The effort to use neural networks to solve partial differential equations (PDEs) with a focus on fluid dynamics has grown over time [1–4]. Eventually, these types of solvers matured into physics informed surrogate deep learning models [5, 6]. These iterative neural network solvers do not gain much in terms of speed, with simulation times on the same order of magnitude as the original solver. But they are a proof of concept that neural networks can approximate CFD solver computations, sacrificing very little in terms of accuracy. These methods struggle to compete with the more than 50 years of optimization behind standard PDE solvers. Instead, methods in PDE solving may be more useful *to* the machine learning community for faster neural network training [7].

**CFD Solver / NN Solver Hybrids:** In CFD and NN hybrid technology, the neural network learns terms that are hard for the CFD solver or unknown, like constitutive material models [8] or turbulence closure terms [9]. These model additions can improve the accuracy of simulation, while incurring some computational cost.

**CNN Solvers:** Convolutional neural network (CNN) based surrogate models are a clearer win for neural networks—methods in [10, 11] were found to dramatically improve solver CPU time, while sacrificing very little accuracy. The CNN surrogate models tend to be a couple orders of magnitude faster than a standard solver. However, these methods can only be applied to spatially uniform grids—a very small subset of aerospace industry problems. One field where these CNN based methods are successfully applied is for the acceleration of computer graphics simulations for Eulerian fluids [12].

#### 2. Reduced Order Models

**ROMs:** ROM and machine learning approaches use similar techniques to create a strategy for reducing the computational cost of CFD while retaining high accuracy [13]. ROMs simulate the original governing equations after projecting

them onto a low-dimensional solution manifold using a reduced-order basis, and by introducing other approximations into the governing equations such as hyperreduction [14, 15], reducing computation cost often by several orders of magnitude [16, 17]. ROMs use knowledge gained from previous simulations to construct the reduced-order basis, and can begin to approximate solutions well starting from a handful of examples—hence, they are often useful for industry problems where getting extremely large CFD datasets for training is infeasible. [18, 19]. For training, ROMs utilize an adaptive sampling method known as a “greedy” procedure for sampling the parameter space and generating high-dimensional training solution snapshots, which allows them to minimize the training time as much as possible. For highly nonlinear problems, they have been equipped with the concept of local reduced-order bases [20], which has allowed them to accelerate the solution of three-dimensional, turbulent, viscous CFD problems associated with complete aircraft configurations by orders of magnitude, while delivering unprecedented accuracy, including at shock locations [15, 21]. However, they can be intrusive for compressible flows, for example, which can be a significant drawback.

**ROM / NN Hybrids:** In hybrid models, combining ROMs and neural networks, the neural network is trained to model the terms that are hard for the ROM or unknown, resulting in more accurate ROMs with some computational overhead. Some methods may learn the projection errors [22] or employ neural networks to accurately approximate the coefficients of the reduced model [23].

## B. Model-Free Methods

### 1. Surface Fitting

Surface fitting on a quantity of interest (QoI) simply maps design parameters directly to QoIs using linear or nonlinear fitting. These methods are simple and not intrusive. Their deterministic versions can be inaccurate [24] as they suffer from uncertainty with respect to the choice of the reconstruction data. However, their nondeterministic variants, known as kriging or Gaussian processes, successfully address this issue. Given large datasets, neural networks have been investigated as alternative methods for mapping, and where large enough datasets are available, have been shown improve upon deterministic curve fitting methods [25]. Along with running as many full order models as computationally feasible, various forms of surface fitting constitute the most common methods employed for design optimization in industry. They are employed whenever the analysis of a system can be limited to a static analysis and a few QoIs.

### 2. Solution Fitting

Solution fitting is our novel proposed approach using neural networks to learn to approximate entire solutions from few examples. The method uses a simple network to learn the function that approximates the CFD solutions in a way that generalizes well from only a few examples. Using this method, solutions for very large number of design parameters can be computed nearly instantaneously with one forward pass of a small network, thereby competing with ROMs. This method is unlike surface fitting in that the network does not only learn a QoI for a given design parameter, but the entire solution, from which any QoI can be computed.

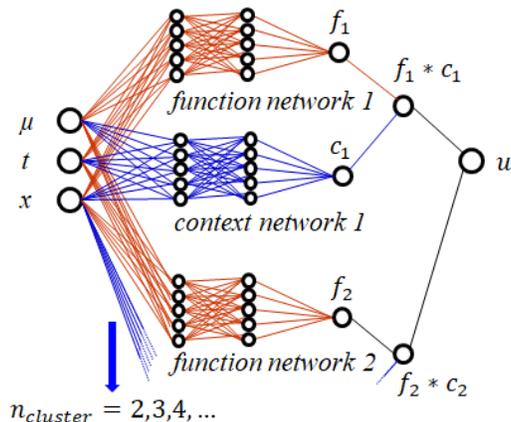
## III. Method

The proposed architectures for full order solution fitting manage the limitation that few examples of design parameters are available by learning the function that approximates the training set solutions. Variables are mapped from  $(x, t, \mu)$  values, which are known, to a fluid state variable of interest,  $u$ , which is unknown. Once either network architecture is constructed, it is trained on the dataset. For evaluation, predictions for unseen design parameters in the test set are computed, and percent error is tabulated and compared to baselines.

### A. Cluster Network

The novel proposed architecture is shown in Figure 2. The network is unique in that different clusters automatically identify regions in the data that behave according to different functions (the function networks), a concept that is similar to that of local reduced-order bases, which are known to greatly improve model order reduction techniques for highly nonlinear dynamical systems [20]. A separate part of the architecture determines when and by how much to turn the other parts on or off (the context networks). Different loss functions are used for each part of the network to train them

as either function or context networks. The separation in intent between different networks is similar to a mixture of experts network [26] and similar in construction to an AI Physicist network [27]. Because each cluster applies to a different function, and because each type of network has a separate role, the network is designed to be easier to interpret than a fully connected network. This network design makes an additional inductive bias over the smooth functions assumed by a fully connected network—that solutions can be broken down into a few regions, defined locally by very simple, smooth functions.

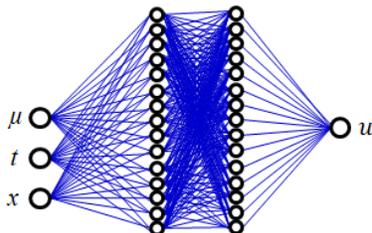


**Fig. 2 Cluster network.**

This network requires a specialized training procedure with the function networks trained by a classification loss function, and the context networks trained successively by a regression loss function. Using the labels in Figure 2, the training procedure for a network with two clusters is (1) train the function networks using a classification loss function, defined as the minimum of the absolute values  $abs(f - f_1)$  and  $abs(f - f_2)$  (2) train the cluster networks while holding the function network weights constant using a regression loss function, defined as the absolute value  $abs(f - y)$ .

### B. Fully Connected Network

The second proposed architecture, for comparison, is the simple feed-forward network in Figure 3. The network performs a mapping directly from a batch of  $(x, t, \mu)$  values to a batch of  $u$  values where  $x$  represents space,  $t$  represents time, and  $u$  represents the output of the network. For the Burgers' equation test case considered in this work,  $u$  represents the velocity. For the shock-induced bubble test case also considered here,  $u$  may represent the density, x-momentum, y-momentum, or energy. Also,  $\mu$  can represent any design variable or parameter in the governing equation than can be easily modified from one simulation to the next. For the Burgers' equation test case,  $\mu$  represents viscosity. For the shock-induced bubble test case, the free-stream Mach number  $M$  is varied instead of  $\mu$ . Because  $x$  and  $t$  can be any scalar value, like the cluster network, this method is easy to apply to any time stepping procedure, grid type, or mesh resolution.



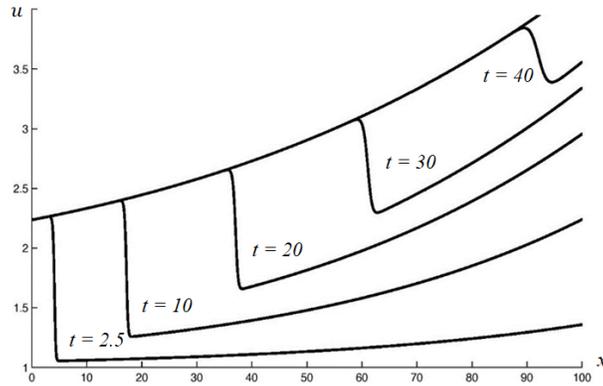
**Fig. 3 Fully connected network.**

## IV. Datasets

Because industry CFD simulations are expensive to compute, our datasets are designed to evaluate the network on its ability to learn well from very few examples. A dataset consisting of 40 simulations, but with variations in 5-10 design parameters would be a typical industry dataset. We test each method on a slightly harder version of the problem, scaled down in order to facilitate faster network evaluation. The networks are given the challenge of learning from only three example simulations with variation in a single design parameter for the two datasets. These datasets we have chosen to constrain and evaluate are the following Burgers' equation and shock-induced bubble test cases.

### A. Burgers' Equation Test Case

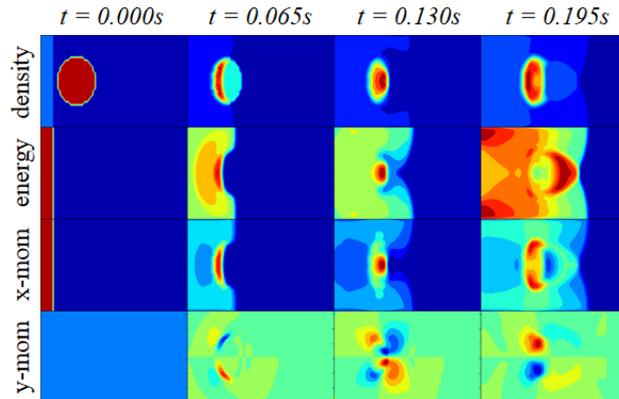
The Burgers' equation test case shown in Figure 4 is commonly used for evaluating ROMs. It is a one-dimensional application of an initial-boundary-value problem, which models the movement of a shockwave across a tube.



**Fig. 4** Burgers' equation evolution in time from 2.5 seconds to 40 seconds showing one full solution in time for a given set of parameters.

### B. Shock-Induced Bubble Test Case

The shock-induced bubble test case is a 2D application of the compressible Navier-Stokes equations, modeling a shockwave moving across a circular high density region representing a 2D bubble. In order to generate the dataset, parameters in the governing equations were varied within a CFD solver, including the free-stream Mach number ( $M$ ), ratio of specific heats, and viscosity. Solutions were generated resolving 0.2 seconds in real time.



**Fig. 5** Shock-induced bubble contour plots showing the progression of the solution for its four state variables—density, energy, x-momentum, and y-momentum—at four points in time. Each solution resolves the progression of the shockwave over the bubble between  $t = 0.000s$  and  $0.200s$  in 201 time steps with step size  $0.001s$ .

At the beginning of the simulation—0.000 seconds—the problem is defined by a high density “bubble” with 5 times the density of the ambient fluid, as visualized in the left-most frames of Figure 5. A moving shockwave is initialized near the left side of the simulation, represented by a small region with high x-momentum, which will move to the right and cross over the bubble. A Mach number greater than 1.0 indicates that the shockwave itself moves at higher than supersonic speed. The four state variables—density, x-momentum, y-momentum, and energy are plotted at  $t = 0.000s$  through  $t = 0.195s$  in Figure 5 for a shockwave Mach number of 1.8.

## V. Numerical Results

### A. Burgers’ Equation Results

#### 1. Results Compared to Baselines

Results for all baseline and proposed methods are tabulated in Table 1 for the Burgers’ equation test case. ROMs, fully connected networks and cluster networks all predict solutions at unseen conditions well, and predict the QoIs calculated from the full solutions well. The fully connected network and cluster network also obtain predictions about two orders of magnitude faster than the full order model.

**Table 1** Burger’s equation test case: performance comparisons with baselines. Online CPU time, percent error on the test set, and percent error on selected quantities of interest are shown for different methods compared to the full order model. Abbreviations: Full order model (FOM), ROM with 100 basis vectors (ROM 100), Fully Connected Network with 4 layers and 25 nodes per layer (FCN 4 20), Cluster Network with 2 clusters 3 layers 5 nodes per layer (ClusterNet 2 3 5), Surface fitting using 1st order polynomial on a QoI (QoI Poly Fit 1).

Method	CPU (s)	% Error	QoI % Error
FOM	16.273		
ROM 100	12.981	0.07	0.05
ROM 50	5.690	0.20	0.16
ROM 20	3.575	0.86	0.60
ROM 5	2.415	4.04	3.00
ROM 3	2.119	5.47	5.28
FCN 4 20	0.245	2.43	3.04
ClusterNet 2 3 5	0.227	0.71	0.42
QoI Poly Fit 3	<0.001		7.28
QoI Poly Fit 2	<0.001		7.15
QoI Poly Fit 1	<0.001		8.42
QoI Poly Fit 0	<0.001		15.03

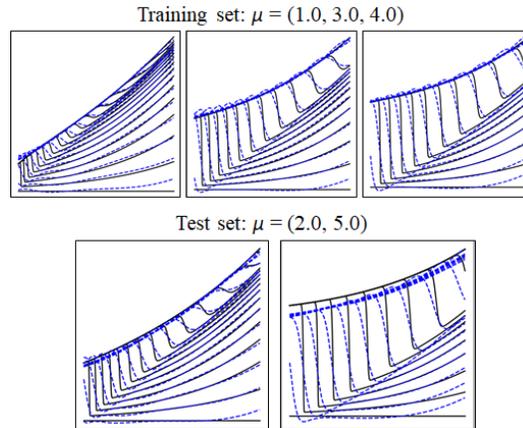
Some QoIs were constructed for more meaningful comparisons because in practice there are usually QoIs that can be computed from a simulation that are important to a design engineer in addition to the full solution. Therefore, four psuedo-QoI’s representative of the kinds of quantities engineers might investigate were calculated from the test set (1) a single random solution point in space and time (2) the average value of the solution at a random fixed point in time (3) the average value of the solution at a random fixed point in space and (4) the average value of the solution over space and time. Similar QoIs were constructed for both Burgers’ equation and the shock-induced bubble test case. The final tabulated QoI is the average of the four psuedo-QoIs.

The percent error is calculated over the entire computational domain by (1) scaling the values for all solutions in the dataset to between 0 and 1, (2) taking the absolute value of the difference between the predicted solutions and solutions computed by the full order model, and (3) multiplying by 100. This percent error metric provides a measure of the difference between the prediction and the ground truth, represented as a percent of the maximum value after the variables have been scaled between 0 and 1. The percent errors are the same as the scaled L1 errors typically investigated in machine learning applications, except multiplied by 100. The QoI percent error is calculated with the

same procedure, except in step (2) we take the absolute value of the difference between the predicted QoI and the QoI computed by the full order model.

### 2. Visualization of Fully Connected Network Results

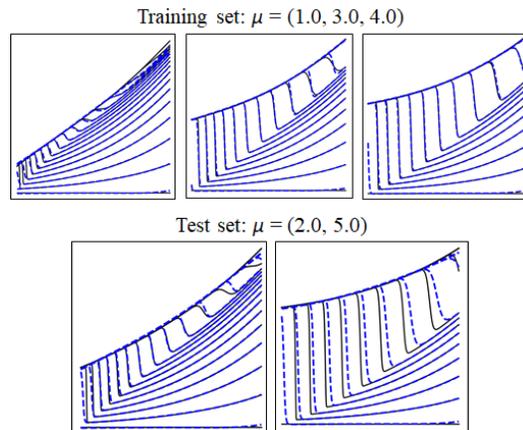
With the right hyper-parameter tuning and the implementation of exponential learning rate decay, the single fully connected network learns to fit the data and predict solution behavior at unseen parameters as shown in Figure 6. Note that the method not only interpolates, but extrapolates to some extent beyond parameters it has seen—this is a surprising result as we only expected neural networks to interpolate well, and the method is only required to interpolate well to be part of a design optimization procedure.



**Fig. 6 Fully connected network results on Burgers' equation.**

### 3. Visualization of Cluster Network Results

The cluster network model was also trained on only three examples, and tested on two examples. Overlays of CFD solutions with neural network predictions in Figure 7 show that the neural network predictions closely approximate the full order model solutions, interpolating and extrapolating better than the fully connected network.



**Fig. 7 Cluster network results on Burgers' equation.**

## B. Shock-Induced Bubble Results

### 1. Results Compared to Baselines

We compare our shock-induced bubble results to the full order model and the QoI baselines in Table 2. Again, the fully connected network and the cluster network outperform the model-free surface fitting methods. The neural

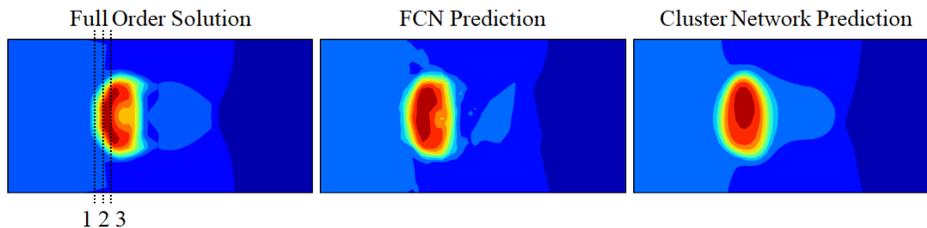
network-based approach has an even greater speedup over the full order model on higher-dimensional problems.

**Table 2 Shock-induced bubble test case: performance comparisons with baselines. Online CPU time and percent error on selected QoIs from the test set shown for different methods compared to the full order model. Abbreviations: Full order model (FOM), Fully Connected Network with 4 layers and 40 nodes per layer (FCN 4 40), Cluster Network with 4 clusters 3 layers 15 nodes per layer (ClusterNet 4 3 15), Surface fitting using 2nd order polynomial fit on a quantity of interest (QoI Poly Fit 2).**

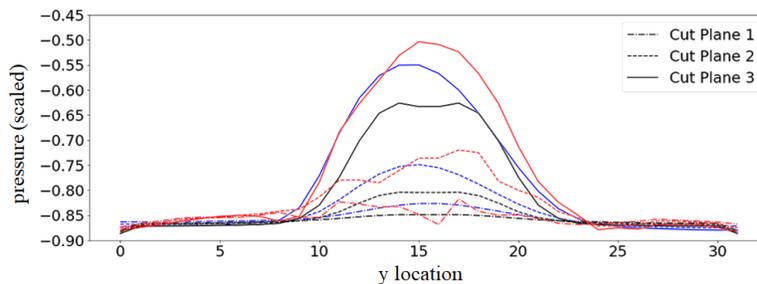
Method	CPU (s)	QoI % Error
FOM	98.374	
FCN 4 40	0.687	0.37
ClusterNet 4 3 15	0.524	0.88
QoI Poly Fit 3	<0.001	5.73
QoI Poly Fit 2	<0.001	5.29
QoI Poly Fit 1	<0.001	1.78
QoI Poly Fit 0	<0.001	1.79

## 2. Visualization of Results

We visualize the accuracy of the network for the shock-induced bubble at a time point near the end of the simulation at 0.180 seconds. The density is shown on a contour plot in  $x$  and  $y$  for each of the shockwave Mach numbers in the training and test sets for the full order CFD solutions and the neural network predictions. The contour plots in Figure 8 show that this relatively small network predicts the solutions at free-stream Mach numbers that it never saw with some significant errors at the locations with more strongly nonlinear density profiles. However, Table 2 reveals that its overall accuracy remains competitive with other model-free methods. Figure 9 shows the shape of the density contours near the shockwave along the three dashed cut planes in Figure 8, where the prediction errors are the highest.



**Fig. 8 2D density contours for the fully connected network and cluster network compared to the full order solution for the shock-induced bubble for one test set example at  $t = 0.18s$  and  $M = 1.8$ .**



**Fig. 9 Density contours for the fully connected network predictions (red), cluster network predictions (blue), and the full order solutions (black) for the shock-induced bubble for one example in the test set at  $t = 0.18s$  and  $M = 1.8$  along three cut planes taken near the shockwave at the dashed lines in Figure 8.**

## VI. Conclusions

Both neural network architectures tested are more accurate than conventional surface fitting. For the one-dimensional Burgers' equation, they are found to deliver the same accuracy as ROMs, but are an order of magnitude faster. Between the two networks, the fully connected network is a good baseline to use for comparison with other neural network architectures for a few reasons: (1) the architecture is simple, (2) the training procedure is simple and (3) the inductive bias assumes smooth functions. However, in practice the network requires extensive hyper-parameter search and tuning. The cluster network is more accurate than the fully connected network for the Burgers' equation test case, and it provides additional benefits: (1) it uses fewer weights, (2) the forward pass is faster than the fully connected network and (3) its inductive bias is more applicable to highly nonlinear CFD problems—a smooth local bias. Ultimately, both network architectures perform well, and the percent error is under 1% on the two test cases for the most successful model—the cluster network.

## VII. Future Work

Future neural network architectures with inductive biases even better suited to fluid dynamics are likely to achieve more accurate results, and compete with model-based methods like ROMs across a wider range of problems. Additionally, demonstrations on high-dimensional industry test cases are necessary for our method to be adopted within industry for design and optimization. Finally, the solution-predicting is only part of the optimization procedure. In order for these tools to ultimately be useful for design and optimization, they will need to be paired with a process for mapping out how quantities of interest vary over the design parameter space. Future work would require a method for choosing the next points for running the expensive simulations, similar to the “greedy” selection procedure used in ROMs, which balances exploration and exploitation. With the addition of a point selection method, we would have a complete procedure for engineers to use for aerodynamic design optimization.

## VIII. Acknowledgments

This first author acknowledges the support by the Department of Energy Computational Science Graduate Fellowship (DOE CSGF) program through US DOE Grant DE-FG02-97ER25308. The third author acknowledges partial support by the Air Force Office of Scientific Research under grant FA9550-17-1-0182, and partial support by a research grant from the King Abdulaziz City for Science and Technology (KACST). At Lawrence Berkeley National Laboratory, Ann Almgren provided the AMReX full order model code used to generate the shock-induced bubble test case and Matt Zahr provided the pyMORTestbed full order model and reduced order model code used to generate and evaluate the Burgers' equation test case. We are grateful to Ann Almgren and Nathaniel Thomas for text revisions and insights.

## References

- [1] Lagaris, I. E., Likas, A., and Fotiadis, D. I., “Artificial neural networks for solving ordinary and partial differential equations,” *IEEE transactions on neural networks*, Vol. 9, No. 5, 1998, pp. 987–1000.
- [2] Parisi, D. R., Mariani, M. C., and Laborde, M. A., “Solving differential equations with unsupervised neural networks,” *Chemical Engineering and Processing: Process Intensification*, Vol. 42, No. 8-9, 2003, pp. 715–721.
- [3] Tsoulos, I. G., Gavriliis, D., and Glavas, E., “Solving differential equations with constructed neural networks,” *Neurocomputing*, Vol. 72, No. 10-12, 2009, pp. 2385–2391.
- [4] Mall, S., and Chakraverty, S., “Comparison of artificial neural network architecture in solving ordinary differential equations,” *Advances in Artificial Neural Systems*, Vol. 2013, 2013, p. 12.
- [5] Raissi, M., and Karniadakis, G. E., “Hidden physics models: Machine learning of nonlinear partial differential equations,” *Journal of Computational Physics*, Vol. 357, 2018, pp. 125–141.
- [6] Raissi, M., Perdikaris, P., and Karniadakis, G., “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *Journal of Computational Physics*, Vol. 378, 2019, pp. 686–707.
- [7] Chen, T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D., “Neural Ordinary Differential Equations,” *arXiv preprint arXiv:1806.07366*, 2018.

- [8] Javadi, A., Tan, T., and Zhang, M., “Neural network for constitutive modelling in finite element analysis,” *Computer Assisted Mechanics and Engineering Sciences*, Vol. 10, No. 4, 2003, pp. 523–530.
- [9] Ling, J., Kurzwski, A., and Templeton, J., “Reynolds averaged turbulence modelling using deep neural networks with embedded invariance,” *Journal of Fluid Mechanics*, Vol. 807, 2016, pp. 155–166.
- [10] Guo, X., Li, W., and Iorio, F., “Convolutional neural networks for steady flow approximation,” *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, 2016, pp. 481–490.
- [11] Stoecklein, D., Lore, K. G., Davies, M., Sarkar, S., and Ganapathysubramanian, B., “Deep learning for flow sculpting: Insights into efficient learning using scientific simulation data,” *Scientific reports*, Vol. 7, 2017, p. 46368.
- [12] Tompson, J., Schlachter, K., Sprechmann, P., and Perlin, K., “Accelerating eulerian fluid simulation with convolutional networks,” *arXiv preprint arXiv:1607.03597*, 2016.
- [13] Amsallem, D., and Farhat, C., “Interpolation method for adapting reduced-order models and application to aeroelasticity,” *AIAA journal*, Vol. 46, No. 7, 2008, pp. 1803–1813.
- [14] Farhat, C., Avery, P., Chapman, T., and Cortial, J., “Dimensional reduction of nonlinear finite element dynamic models with finite rotations and energy-based mesh sampling and weighting for computational efficiency,” *International Journal for Numerical Methods in Engineering*, Vol. 98, No. 9, 2014, pp. 625–662.
- [15] Carlberg, K., Farhat, C., Cortial, J., and Amsallem, D., “The GNAT method for nonlinear model reduction: effective implementation and application to computational fluid dynamics and turbulent flows,” *Journal of Computational Physics*, Vol. 242, 2013, pp. 623–647.
- [16] Farhat, C., Chapman, T., and Avery, P., “Structure-preserving, stability, and accuracy properties of the energy-conserving sampling and weighting method for the hyper reduction of nonlinear finite element dynamic models,” *International Journal for Numerical Methods in Engineering*, Vol. 102, No. 5, 2015, pp. 1077–1110.
- [17] Zahr, M. J., Avery, P., and Farhat, C., “A multilevel projection-based model order reduction framework for nonlinear dynamic multiscale problems in structural and solid mechanics,” *International Journal for Numerical Methods in Engineering*, Vol. 112, No. 8, 2017, pp. 855–881.
- [18] Carlberg, K., Bou-Mosleh, C., and Farhat, C., “Efficient non-linear model reduction via a least-squares Petrov–Galerkin projection and compressive tensor approximations,” *International Journal for Numerical Methods in Engineering*, Vol. 86, No. 2, 2011, pp. 155–181.
- [19] Zahr, M. J., Carlberg, K. T., and Kouri, D. P., “An efficient, globally convergent method for optimization under uncertainty using adaptive model reduction and sparse grids,” *arXiv preprint arXiv:1811.00177*, 2018.
- [20] Amsallem, D., Zahr, M. J., and Farhat, C., “Nonlinear model order reduction based on local reduced-order bases,” *International Journal for Numerical Methods in Engineering*, Vol. 92, No. 10, 2012, pp. 891–916.
- [21] Washabaugh, K. M., Zahr, M. J., and Farhat, C., “On the use of discrete nonlinear reduced-order models for the prediction of steady-state flows past parametrically deformed complex geometries,” *54th AIAA Aerospace Sciences Meeting*, 2016, p. 1814.
- [22] San, O., and Maulik, R., “Neural network closures for nonlinear model order reduction,” *Advances in Computational Mathematics*, 2018, pp. 1–34.
- [23] Hesthaven, J. S., and Ubbiali, S., “Non-intrusive reduced order modeling of nonlinear problems using neural networks,” *Journal of Computational Physics*, Vol. 363, 2018, pp. 55–78.
- [24] Chilenski, M., Greenwald, M., Marzouk, Y., Howard, N., White, A., Rice, J., and Walk, J., “Improved profile fitting and quantification of uncertainty in experimental measurements of impurity transport coefficients using Gaussian process regression,” *Nuclear Fusion*, Vol. 55, No. 2, 2015, p. 023012.
- [25] Trizila, P., Kang, C.-K., Aono, H., Shyy, W., and Visbal, M., “Low-Reynolds-number aerodynamics of a flapping rigid flat plate,” *AIAA journal*, Vol. 49, No. 4, 2011, pp. 806–823.
- [26] Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q., Hinton, G., and Dean, J., “Outrageously large neural networks: The sparsely-gated mixture-of-experts layer,” *arXiv preprint arXiv:1701.06538*, 2017.
- [27] Wu, T., and Tegmark, M., “Toward an AI Physicist for Unsupervised Learning,” *arXiv preprint arXiv:1810.10525*, 2018.