



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Perspective on machine learning for advancing fluid mechanics

M. P. Brenner, J. D. Eldredge, and J. B. Freund

Phys. Rev. Fluids **4**, 100501 — Published 16 October 2019

DOI: [10.1103/PhysRevFluids.4.100501](https://doi.org/10.1103/PhysRevFluids.4.100501)

A Perspective on Machine Learning for Advancing Fluid Mechanics

M. P. Brenner^{1,2}

¹*School of Engineering and Applied Sciences,
Harvard University, Cambridge, MA, 02138*

²*Google Research, Mountain View CA 94103
mpbrenner@gmail.com*

J. D. Eldredge

*Mechanical and Aerospace Engineering
University of California, Los Angeles
Los Angeles, CA 90095, USA
jdeldre@ucla.edu*

J. B. Freund

*Mechanical Science & Engineering and Aerospace Engineering
University of Illinois at Urbana–Champaign
Urbana, IL 61801, USA
jbfreund@illinois.edu*

Abstract

A perspective is presented on how machine learning (ML), with its burgeoning popularity and the increasing availability of portable implementations, might advance fluid mechanics. As with any numerical or experimental method, ML methods have strengths and limitations, which are acknowledged. Their potential impact is high so long as outcomes are held to the long-standing critical standards that should guide studies of flow physics.

I. INTRODUCTION

In the last five years there has been a leap in the visibility of machine learning, with computational algorithms now carrying out tasks that a decade ago seemed far off or even impossible. Examples range from the vast improvements in speech recognition and machine translation of natural language to medical diagnostics.

The progress in machine learning that led to these advances has been mostly technological, and is well reviewed for fluid mechanics elsewhere [1]. Though any compact expression will not express all the complex ways ML algorithms can use data, simplistically they take an input \mathbf{x} and map it to an output \mathbf{y} : $\mathbf{y} = \mathbf{F}_{\mathbf{w}}(\mathbf{x})$ via a mapping parameterized by \mathbf{w} . The \mathbf{w} values are chosen by attempting to minimize a function $C(\mathbf{x}, \mathbf{y})$, often termed a cost or loss function. Machine learning *per se* is the determination of \mathbf{w} —learning it—so that \mathbf{F} can be used for some task. Historically such algorithms used a relatively small number of parameters with analytically tractable loss functions for ease of computation. For linear regression with a quadratic cost function, the problem can be solved analytically. To facilitate analysis, cost functions have often been chosen to be convex so that there is a single, unique minimum. Although fluid mechanics primarily focuses on physical mechanisms for fluid phenomena, variants of this broader machine learning framework have played an important role in situations where mechanisms are difficult to parameterize. A prominent example is Proper Orthogonal Decomposition [2], which is equivalent to a linear ML autoencoder.

Recent progress derives from the surprising discovery in multiple application domains that choosing $\mathbf{y} = F_{\mathbf{w}}(\mathbf{x})$ with an extraordinarily large numbers of parameters (\mathbf{w}) and with non-convex $C(\mathbf{x}, \mathbf{y})$ can produce models that are both accurate and generalizable, working outside the set of training examples used to set \mathbf{w} . Particularly successful mapping functions have been parameterized with neural networks, invented more than thirty years ago, though progress initially lagged due to the computational complexity for increasing numbers of parameters exceeding the computational power then available. Modern neural networks can easily have hundreds of millions of weights. Although their mathematical formulation is straightforward to describe as layers of simple neuron functions, mathematical explanations are incomplete regarding why or how minimizing non-convex functions, with so many parameters, can lead to generalizable models. To many the idea of working with a model that

is not fully understood is disturbing, and caution is indeed warranted. Still, it is impossible to dismiss the utility of these models in multiple domains.

Two factors made it possible to train these complex models (seek local minima). The first was the expanded use of automatic differentiation and back propagation algorithms, especially the ease of programming them. These algorithms are essentially the neural-network analogue of adjoint-based sensitivity methods, which have been used in fluid mechanics for over fifty years [3]. The adjoint in that case is derived from the governing equation whereas for ML the neural-network parameters are inferred based on data representing flow solutions. In general, deriving an adjoint for such an optimization is tedious; robust packages now exist for carrying out these calculations automatically for neural-network models as well as for discrete representations of flow equations, either automatically [4] or directly derived for greater implementation efficiency. The second factor is computational, with open-source toolkits that enable ML models to leverage modern graphical processing units [5, 6]. To broad surprise, the results were both wondrous and practical. Special purpose hardware is even being developed that is optimized for evaluating neural networks. These same technologies, as they continue to improve with greater power and parallel scaling, have the potential of making a major impact in fluid mechanics, giving a methodology for asking $\mathbf{y} = F_{\mathbf{w}}(\mathbf{x})$ questions and answering them in flow regimes that were once unimaginable.

II. AN EXAMPLE

To be concrete about potential we consider an example where ML has impacted science, in which a ML classifier was trained to detect diseases of the retina [7]. The input \mathbf{x} was the image pixels and the output \mathbf{y} the classification. The authors constructed a large data set of more than 50,000 fundus images of the eye, and then a panel of ophthalmologists labeled the images for whether or not they had eye disease, and if so, the level of severity. A large convolutional neural network, trained on these images, performed better than a panel of ophthalmologists in its first instantiation and, with improvements, better than a panel of retinal specialists [8].

Though of obvious practical importance, more interestingly this study also led to a discovery. The retinal images also came with additional data from the person's medical record. A

retrained version of the model was able to predict blood pressure, age, and sex. Though unprecedented, age and blood pressure were not particularly surprising, except perhaps in the accuracy (the 95% confidence interval for age was 3.26 years), because morphology changes with age and blood vessels are the most obvious features in a fundus image. The fundus image–sex correlation, however, was a scientific discovery, as up until that point there had been no known connection between the structure of the retina and sex. Yet the model predicted this correlation with 97% accuracy.

What is interesting about these results is that the computer with data, set up to construct what might be considered a mundane albeit intricate correlation, enabled discovery of something that had not been previously suspected. This example is particularly interesting because it was possible to then disentangle the decision tree of a neural network to extract the rules for diagnosing macular disease, and provide new guidelines for diagnosis. Computer classification led to diagnosis rules that can be implemented by humans. Such disentanglement is hard, and the sex–retina correlation is still unexplained, though it is now a target of hypothesis generation for potential discovery and generalization of findings more broadly to human physiology.

Although fluid mechanics is not typically presented as classification problems, it can be. For example, images of flows under different conditions (turbulent or not, Reynolds number, stable or not, separated or not) can be classified, so a ML classifier can be trained to categorize. At the outset this might be technically challenging and scientifically mundane, maybe as mundane as implementing a numerical method or setting up a PIV system, but it facilitates questions about how the classifier makes the classification that it does. We would expect flow structure to be classified as higher and higher Reynolds number as eddies of decreasing size show influences of viscosity. An opportunity exists in dissecting the logic of how networks carry out classifications, ideally mathematically though more likely first conceptually, to advance understanding. This is easy to envision in a currently understood simple context, such as near a stability threshold; extending this to increasingly complex nonlinear situations is less clear, though such situations are where opportunities lie. This should be done with care, and some concerns are addressed in the following section, followed by examples from fluid mechanics selected to illustrate how progress might follow.

III. SKEPTICISM AND RESPONSES

Recent advances in ML methods are in the computational science underlying function approximation, which is common in fluid mechanics. Indeed, the apparent truth of many correlations is almost codified (e.g., the law of the wall), with debates about their coefficients. A difference is that these established correlations are typically supported by a combination of dimensional analysis and physical insights, and their mathematical form is explicit, with few inputs and outputs. An aspect of ML that fosters discomfort is the complexity of its mapping from inputs to outputs. Correspondingly, the desiderata for a function can be specified with much more freedom than ever and tailored to particular problems. This is contrary to the traditional approach for function approximation, such as Galerkin expansion, Fourier transforms, or Proper Orthogonal Decomposition. Even boundary layer approximations are general, applying whenever the underlying problem has the appropriate separation of length scales. Of course, however an approximation is constructed, if it is done well, then it can be usefully applied. And, if it is done badly, then it is undeniably not useful. Few can say we should not deploy every useful mathematical technique we can find to advance fluid mechanics, and for that reason the community should take advantage of the recent advances if they are useful.

There are, however, significant pitfalls that must be avoided. First and foremost, the core goal of the science of fluid mechanics is to discover mechanisms. Reducing a flow process to its essence can, for example, require recognizing dominant mechanisms within the Navier–Stokes equations and understanding why these terms govern the process. In contrast, many advances in machine learning proceed without regard to mechanism. For example, it is striking that state-of-the-art approaches to machine translation (from one language to another) forego grammar and are entirely data driven [9]. Similarly, speech recognition models are not based on acoustic models of sound production, but instead on large datasets correlating acoustic input to words [10]. To advance fluid mechanics as a discipline, correlations are clearly not enough, though they might have immediate impact in engineering or other applications of fluid mechanics.

Yet, there is also an opportunity for addressing age old problems in a completely new way. Old methods also have their share of historical accident. Not all stages of science discovery

demand the same level of rigor, and certainly are not restricted to optimizing convex functions that are analytically tractable. Linear algebra should be used when there is an L_2 norm, and Sturm–Liouville theory should be applied when problems are well-described by their bases. We have historically used these in fluid mechanics for computational or calculational convenience, though they are not fundamental representations of flow.

It is also important to recognize that the components of modern neural networks are not far removed from our familiar linear modeling toolbox. For example, a convolutional neural network (CNN) [11] is, at its heart, a local filtering operation. The difference is that, rather than relying on filter kernels that are pre-ordained by some underlying physical law or intuition, kernels are learned from the training process. Sometimes the trained weights that emerge resemble a discrete form of differential operation on the underlying data. Furthermore, the linear filtering operation in a CNN is usually augmented with crucial nonlinear operations, such as locally pooling the output field and applying activation (generally, rectifying or saturating the output), that enrich the network’s ability to approximate functions. Suitable deep networks provide a range of operations that can be useful for extracting correlations. Similarly, a recurrent neural network (and particularly, Long Short-Term Memory, or LSTM [12]) correlates time-varying inputs and output signals by providing a richer version of the already-familiar convolution with an impulse or indicial response function. Some event in the past, experienced in the input data, lingers in a memory propagated from one LSTM cell to the next. But, importantly, the LSTM provides nonlinear enhancements: the response kernel (and its time constant) can vary as time proceeds, and it can forget past events based on a trigger. Thus, we do not completely lose interpretability with such machine learning frameworks; what interpretability is lost is the price we pay for increased potency for function approximation.

IV. APPLICATIONS TO FLOW: RECENT AND HYPOTHETICAL

Whether or not these methods are deemed to impact fluid mechanics depends, in some sense, on the expectations, and maybe where one draws the boundaries of fluid mechanics *per se* versus application of fluid mechanics. We select but a few examples; a much more complete discussion, including a historical perspective, is available in a forthcoming review

article [1].

There are already many examples showing how the correlation/data-compression properties of data-trained mathematical models can aid predictions. Closure of Reynolds-averaged Navier–Stokes (RANS) models against DNS is an important example [13–16], including cases where the trained system provides an ‘alarm’ for anticipated errors [17]. In a similar vein, they can correct approximate models, such as for finite-sized particle drag [18]. Both the RANS governing equations and asymptotic drag models in these examples embody physics, which is then augmented with ML. There are sure to be many such applications. Many of these will unfortunately not have provable properties, because of the complexity of the learned representation, though neither do the models they augment or replace because of the complexity of the physical system. Still, they are potentially useful.

ML models are also starting to have influence closer to the principles of fluid mechanics, when they are used in conjunction with human reasoning to organize complex flow field information in a manner that facilitates the asking and answering of questions. This is similar to the classification discussed in context of the retinal fundus example. Recent examples in these cases involve identifying the nominally important features in two-dimensional turbulence [19], deducing energy-saving principles underlying the schooling dynamics of fish [20], and deducing key features leading to successful ignition of combustibles in turbulence [21]. Each of these leveraged the trained ML model in a different fashion. Jiménez [19] reasoned what an important structure should do, and then used ML analysis of a large database to explore the consequences of this idea for its definition. Verma et al. [20] used reinforcement learning, to model in a sense what a fish’s biological neural network might do given the opportunities afforded by the complex wake structures surrounding fish. Popov et al. [21] used a relatively simple network to facilitate disentanglement of its sub-features (the convolution kernels in the CNN) and a corresponding sensitivity analysis to define important features. More complex scenarios might challenge these methods, though we do not seem to be near a fundamental complexity limit in expanding such applications.

There are also examples of using ML to discover facile definitions within complex systems that are consistent with the intuitive expectations of the users. This is useful when a concept is clear and enabling for subsequent analysis but lacks a specific definition, such as classical boundary layer thickness. An unsupervised self-organizing map was used recently

to cluster regions of a transitioning boundary layer into a consistent and robust data-driven definition of the turbulent boundary layer interface [22], avoiding some concerns with typical threshold definitions. There are other likely opportunities for ML methods to provide means of developing definitions automatically based on simpler notions.

Recent results also show that a trained network can, in a sense, replace challenging dynamics. Examples are the deduction of drag and lift given limited field information [23], inference of flow characteristics from aerodynamics surface measurements [24] super-resolution of unsteady flow details from limited information [25], and anticipation of extreme events in a Kolmogorov flow [26]. These examples are interesting, though only two dimensional, without the stronger chaos of the true turbulence that would exist in corresponding applications. It will be interesting to understand the limits of such an approach in this regard and when a more complete physical description is needed.

It is hoped that this survey is illustrative regarding opportunities for ML in fluid mechanics applications and studies. Many other examples are not discussed. It is noteworthy that there does not yet seem to be an example of ML impacting what might be considered foundational fluid mechanics. All of the examples above were constructed around fluid mechanics principles that were already understood. Still, we can speculate what a fundamental contribution might look like, if only to recognize that it is potentially a long way off. Though it might be unlikely, let us hypothesize that a hitherto undiscovered invariant constrains incompressible turbulent flow, perhaps in the same vein as impulse integral constraints apply in two dimensions. Since it surely could aid the input–output description of the training setup, a deep neural network, trained on trusted turbulence data, might indeed make implicit use of this hypothetical invariant. However, at present there is no sure route for extracting it from the trained network weights as a physical or mathematical law. Its existence in the parameterized ML network might be equally opaque as it is within the flow phenomena. Means of seeking such properties might constitute an interesting research direction in ML applied to scientific applications. Such efforts might build on methods such as Layer-wise Relevance Propagation, which is a network analysis to deduce how it labels images [27]. Speculating still further, it is even harder to imagine how ML would discover a new concept. For this, we might craft another thought experiment: by what route could training data lead to, for example, the distilled form of boundary layer theory at the foun-

dation of high-Reynolds-number fluid mechanics? When or if human understanding will be advanced in such a way will remain a mystery until a point, if ever, such a feat is achieved. Still, the first step might be the more mundane present-day applications oriented around classification and definition.

V. PERSPECTIVE SYNOPSIS

The primary perspective we offer is based on two main observations. Although they are often constructed based on heuristics, ML algorithms are themselves well-defined mathematically, and indeed the training procedures of the presently pervasive deep learning network models leverage this characteristic. They also have a massive infrastructure supporting their broad use. As such, ML will certainly play some role for the correlations they can find, offering ways to enrich fluid mechanics. In addition, we also offer a perspective based on some speculation, which we do with less confidence but in the hope it provides some framework for how ML might develop in fluid mechanics. As with its much touted successes with images and language, success will depend in part upon the availability of training data, which can be hard to come by in the form of high-fidelity simulations or experiments; many ML algorithms need a lot of data to reduce the error to suitably low values, and ways to pre-train in the context of fluid mechanics will be an interesting research direction. The examples provided illustrate its potential in melding observational data into precise fluid mechanics. How this links with the biological neural networks in the head of the fluid mechanician remains unclear, though there is little risk so long as discoveries are held to the same high standards for any generalizable advance in understanding. It would be bold to predict just how these methods will make the greatest inroads in fluid mechanics, and no two people will share a common view—still less, these three authors—but such a flexible method would seem to be well-positioned to help, and we are excited to see attempts.

ACKNOWLEDGEMENTS

The authors are grateful for critical input from Petros Koumoutsakos.

- [1] S. Brunton, B. Noack, and P. Koumoutsakos. Machine learning for fluid mechanics. *arXiv preprint arXiv:1905.11075; to appear Ann. Rev. Fluid Mech*, 2020.
- [2] C. W. Rowley and S. T. M. Dawson. Model reduction for flow analysis and control. *Annual Review of Fluid Mechanics*, 49:387–417, 2017.
- [3] A. E Bryson and Y.-C. Ho. Applied optimal control, revised printing. *Hemisphere, New York*, 1975.
- [4] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind. Automatic differentiation in machine learning: a survey. *Journal of Machine Learning Research*, 18:1–43, 2018.
- [5] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 265–283, 2016. URL <https://www.usenix.org/system/files/conference/osdi16/osdi16-abadi.pdf>.
- [6] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*, 2017.
- [7] V. Gulshan, L. Peng, M.c Coram, M. C. Stumpe, D. Wu, A. Narayanaswamy, S. Venugopalan, K. Widner, T. Madams, J. Cuadros, R. Kim, R. Raman, P. C. Nelson, J. L. Mega, and D. R. Webster. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA*, 316(22):2402–2410, 2016.
- [8] R. Poplin, A. V. Varadarajan, K. Blumer, Y. Liu, M. V. McConnell, G. S. Corrado, L. Peng, and D. R. Webster. Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. *Nature Biomedical Engineering*, 2(3):158, 2018.

- [9] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey, J. Klingner, A. Shah, M. Johnson, X. Liu, L. Kaiser, S. Gouws, Y. Kato, T. Kudo, H. Kazawa, K. Stevens, G. Kurian, N. Patil, W. Wang, C. Young, J. Smith, J. Riesa, A. Rudnick, O. Vinyals, G. Corrado, M. Hughes, and J. Dean. Google’s neural machine translation system: Bridging the gap between human and machine translation. *CoRR*, abs/1609.08144, 2016. URL <http://arxiv.org/abs/1609.08144>.
- [10] Y. He, T. N. Sainath, R. Prabhavalkar, I. McGraw, R. Alvarez, D. Zhao, D. Rybach, A. Kannan, Y. Wu, R. Pang, Q. Liang, D. Bhatia, Y. Shangguan, B. Li, G. Pundak, K. C. Sim, T. Bagby, S.-Y. Chang, K. Rao, and A. Gruenstein. Streaming end-to-end speech recognition for mobile devices. *CoRR*, abs/1811.06621, 2018. URL <http://arxiv.org/abs/1811.06621>.
- [11] Y. LeCun and Y. Bengio. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995, 1995.
- [12] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [13] J. Ling, R. Jones, and J. Templeton. Machine learning strategies for systems with invariance properties. *Journal of Computational Physics*, 318:22–53, 2016.
- [14] J.-X. Wang, J.-L. Wu, and H. Xiao. Physics-informed machine learning approach for reconstructing Reynolds stress modeling discrepancies based on DNS data. *Physical Review Fluids*, 2(3):034603, 2017.
- [15] J.-L. Wu, H. Xiao, and E. Paterson. Physics-informed machine learning approach for augmenting turbulence models: A comprehensive framework. *Physical Review Fluids*, 3(7):074602, 2018.
- [16] K. Duraisamy, G. Iaccarino, and H. Xiao. Turbulence modeling in the age of data. *Annual Review of Fluid Mechanics*, 51:357–377, 2019.
- [17] J. Ling and J. Templeton. Evaluation of machine learning algorithms for prediction of regions of high reynolds averaged navier stokes uncertainty. *Physics of Fluids*, 27(8):085103, 2015.
- [18] Z. Y. Wan and T. P. Sapsis. Machine learning the kinematics of spherical particles in fluid flows. *Journal of Fluid Mechanics*, 857, 2018.
- [19] J. Jiménez. Machine-aided turbulence theory. *Journal of Fluid Mechanics*, 854, 2018.

- [20] S. Verma, G. Novati, and P. Koumoutsakos. Efficient collective swimming by harnessing vortices through deep reinforcement learning. *Proceedings of the National Academy of Sciences*, 115(23):5849–5854, 2018.
- [21] P. P. Popov, D. A. Buchta, M. J. Anderson, L. Massa, J. Capecelatro, D. J. Bodony, and J. B. Freund. Machine learning-assisted early ignition prediction in a complex flow. *Combustion and Flame*, 206:451–466, 2019.
- [22] Z. Wu, J. Lee, C. Meneveau, and T. Zaki. Application of a self-organizing map to identify the turbulent-boundary-layer interface in a transitional flow. *Physical Review Fluids*, 4(2):023902, 2019.
- [23] M. Raissi, Z. Wang, M. S. Triantafyllou, and G. Em. Karniadakis. Deep learning of vortex-induced vibrations. *Journal of Fluid Mechanics*, 861:119–137, 2019.
- [24] W. Hou, D. Darakananda, and J. D. Eldredge. Machine learning-based detection of aerodynamic disturbances using surface pressure measurements. *to appear AIAA J.*, 2019.
- [25] K. Fukami, K. Fukagata, and K. Taira. Super-resolution reconstruction of turbulent flows with machine learning. *Journal of Fluid Mechanics*, 870:106–120, 2019.
- [26] Z. Y. Wan, P. Vlachas, P. Koumoutsakos, and T. Sapsis. Data-assisted reduced-order modeling of extreme events in complex dynamical systems. *PloS One*, 13(5):e0197704, 2018.
- [27] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS One*, 10(7):e0130140, 2015.