Proceedings of the Combustion Institute 40 (2024) 105730

Contents lists available at ScienceDirect

Proceedings of the Combustion Institute

journal homepage: www.elsevier.com/locate/proci

Artificial intelligence as a catalyst for combustion science and engineering

Matthias Ihme [a,b,c,*], Wai Tong Chung [a]

[a] Department of Mechanical Engineering, Stanford University, Stanford, CA 94305, USA
[b] Department of Photon Science, SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA
[c] Department of Energy Science and Engineering, Stanford University, Stanford, CA 94305, USA

ARTICLE INFO

Keywords:
Artificial intelligence
Machine learning
Combustion
Data-driven methods

ABSTRACT

Combustion and energy conversion play critical roles in all facets of environmental and technological applications, including the utilization of sustainable energy sources for power generation and propulsion, the reduction of pollutant emissions from combustion, or the mitigation of harm from wildfire hazards. Computational and mathematical tools have long been crucial in combustion research in the form of high-fidelity simulations, dynamical-system modeling, and data analytics. With the advent of data-driven methods, machine learning (ML) and artificial intelligence (AI) offer numerous opportunities for predictive modeling, improving existing research methods, and extracting new knowledge from data. In this article, we discuss recent progress on how ML and AI can impact the field of combustion and energy conversion, and discuss the need for domain knowledge for successful ML applications in combustion. Specifically, combustion ML learns from data extracted from large-scale simulations, high-resolution experiments, and sensors, which can introduce challenges tied to dimensionality, interpretability, sparsity, modality, and scarcity. The collective knowledge from these advancements equip combustion researchers and engineers with the ability to adapt to emerging developments in ML foundation models and AI agents, which have begun to offer greater automation across different combustion domains. To this end, we assess opportunities and challenges provided by state-of-the-art ML foundation models, and discuss emerging areas for adapting these new technologies towards solving pressing challenges within sustainable combustion.

1. Introduction

1.1. Motivation, objective, and outline

Combustion science and engineering has facilitated the development of a wide range of technologies, including propulsion and energy conversion systems, thereby significantly enhancing global standards of living [1,2]. However, there are significant challenges ahead that require the attention of combustion researchers and engineers, especially in overcoming increasingly pressing climate-related concerns [3,4]. Promising pathways for addressing these issues include the discovery of methods for (i) synthesizing and characterizing new sustainable fuels [5,6], (ii) designing combustion systems that can adapt to novel fuels [7,8], (iii) implementing advanced combustion strategies for ameliorating greenhouse gas emissions from fossil fuel utilization [9,10], and (iv) developing mitigation strategies for climate-related hazards [11,12].

Each of these pathways requires expertise and effort from every field within combustion. For example, the safe deployment of novel fuels requires developing experimental methods for measuring fundamental fuel properties and realistic combustion behavior [13,14]. These collected measurements can also be complemented by computational chemistry models and high-fidelity calculations, which will require the discovery of new algorithms and the development of advanced software engineering practices within a rapidly changing computing landscape [15,16]. Prior to performing experiments and calculations, significant planning is often required to design optimal target conditions, especially if collecting measurements or performing simulations is costly [17]. In addition, data processing techniques will need improvements to manage the increased complexity and volume of datasets as diagnostics and numerical solvers grow in capabilities [18].

Given the laborious nature of the aforementioned scientific tasks, combustion science and engineering can benefit immensely from increased automation. Cavallotti [19] expressed a similar point for the automation of chemical kinetics, where advances in computational techniques, digital resources, and high-performance computing provide a path towards increased automation across different chemical kinetics-related tasks, such as species and reaction selection, ab initio

* Corresponding author at: Department of Mechanical Engineering, Stanford University, Stanford, CA 94305, USA.
E-mail address: mihme@stanford.edu (M. Ihme).

https://doi.org/10.1016/j.proci.2024.105730
Received 2 February 2024; Accepted 10 July 2024
Available online 14 August 2024

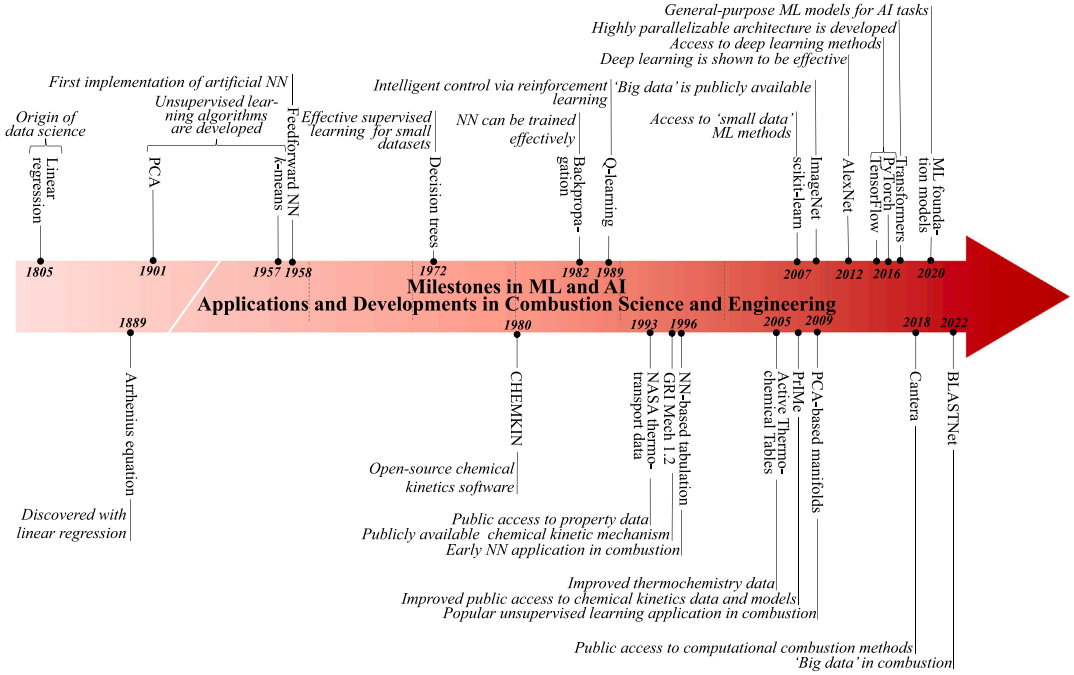

Fig. 1. Timeline of key milestones in ML and AI, along with related developments in combustion.

calculations, thermo-chemical/transport parameterization, rate parameter estimation, and validation, which now enable a greater availability of chemical kinetics models [20–25] for adapting to shifting research trends.

Examples of these advances include stiff numerical methods (which were initially built for atmospheric chemistry [26]) that now automatically solve differential equations through packaged software libraries [27]. Progress in modeling increasingly realistic fuels is partially catalyzed by establishing community-accepted methods and resources, as seen with the NIST webbook [28], the JANAF thermo-chemical tables [29], and general-purpose combustion software packages, such as the NASA equilibrium solver [30], Chemkin [31], and more recently Cantera [32]. The availability of detailed chemical mechanisms has also led to automated chemical optimization and reduction schemes [33] that have been influenced by concepts within linear algebra [34,35], graph theory [36,37], and optimization [38,39]. The integration of these reduced chemical mechanisms, chemistry packages, and computational fluid dynamics (CFD) solvers enable simulations of fundamental and realistic combustion systems [16,40], which provide knowledge for adapting to sustainability challenges.

Increasing computational power has also spurred the exponential growth of highly parallelizable data-driven methods, which are often categorized under an umbrella term: machine learning (ML). Breakthroughs of ML methods in performing human tasks have led to a new class of general-purpose language-based foundation models [41] that have been shown to be able to match and even outperform human performance in an increasingly wide range of tasks including verbal reasoning, programming, and writing [42–44]. These technological developments demonstrate promising paths towards versatile artificial intelligence (AI) agents for automating an even wider range of tasks within different scientific and engineering domains, as seen with emerging studies in biology [45], material science [46], chemical synthesis [47], and mathematics [48]. Note that within this text, we refer to ML as the subset of data-driven methods that fall within the typical supervised, unsupervised, and semi-supervised learning categorizations [49], while AI is the use of any computational method for performing human-like tasks [50].

As these ML and AI techniques become increasingly promising for aiding combustion researchers and engineers in facing pressing challenges, this article aims to:

- Provide readers with a holistic view on the progress of ML and AI.
- Identify how these developments in AI have and can be used to catalyze and automate various aspects of combustion processes.
- Examine challenges and prospects of AI-based autonomous science in combustion.
- Discuss opportunities for combustion researchers and engineers to further integrate AI-based tools for advancing the field of combustion in the presence of emerging needs for sustainable energy utilization.

To this end, we provide an overview of the general progress of AI and ML methods in Section 2. Section 3 is concerned with decomposing the scientific process into sub-processes and tasks, and reviewing recent developments in combustion ML that perform well in these tasks. Since previous articles [18,51–53] have provided in-depth surveys of ML developments within combustion fundamentals and applications, we narrow our review in this section to discuss selected ML studies over the past two years that have focused on aiding combustion researchers in scientific tasks. Section 4 introduces readers to ML foundation models, the leading paradigm within AI, discusses potential directions for integrating foundation models in combustion science and engineering, and demonstrates capabilities and limitations of state-of-the-art ML foundation models in automating various tasks within combustion. We suggest potential directions for employing and furthering AI-based tools within combustion science in Section 5, before summarizing this article in Section 6.

2. Overview of ML and AI progress

In this section, we will review the progress in the development of ML and AI methods in the broader scope of combustion science and engineering. An overview of ML and AI developments is presented in Fig. 1. Supervised learning involves problems (typically in classification and regression settings) that rely on datasets that have been structured into inputs and target outputs through labeling. Linear regression [54] is an early supervised learning approach, with more expressive methods that embrace non-linear operations, such as neural networks (NNs) [55], typically used to solve more challenging

data-driven problems that require higher predictive accuracy. Today, linear regression is ubiquitous in science and engineering applications, and has contributed to several important discoveries, including the Arrhenius equation [56]. Linear regression is also essential in obtaining polynomial fits used for obtaining thermodynamics and transport data [57].

In contrast, principal component analysis (PCA) [58] and k-means [59] are often considered unsupervised learning techniques (typically applied to clustering and dimensionality reduction problems), due to their ability to discover patterns without any form of labeling. In more recent times, dimensionality-reduction techniques have demonstrated promise for automatic identification of low-order manifolds for optimizing combustion chemistry [60]. Data-driven methods involving reinforcement learning (RL), such as Q-learning [61], autonomously create structure without data labels by employing exploitation and exploration schemes that are supervised by a reward function. As highlighted in Fig. 1, the foundations for these techniques were established decades ago, and the interested reader is encouraged to access previous texts [49,62] and software packages [63] that focus on fundamental ML for further details.

Despite its long history, ML methods have seen a sudden growth in popularity due to (i) growing accessibility of large datasets and open-source software tools, as well as (ii) increasingly sophisticated deep learning techniques, which typically involve multiple-layered NN architectures that scale well with large datasets, in solving a wide range of predictive modeling tasks. The seminal moment for this surge in interest began within computer vision, with the publication of the ImageNet dataset [64], which demonstrated the possibility of curating large and diverse structured ML datasets from internet sources. Similar efforts in big data for combustion have begun to gain traction, as seen with the development of the BLASTNet dataset [65], which aims to provide public access to high-fidelity combustion simulation data. This work is a more recent example of the long history of data curation efforts within the combustion community, as seen with NASA thermo-chemical data [57], active thermo-chemical tables [66], and process information systems for chemical kinetics [67–70]. The ImageNet dataset led to the development of AlexNet [71], which re-ignited interest in deep learning methods by demonstrating the superiority of GPU-accelerated NNs over other supervised learning techniques, especially when training with large datasets. We also note that NN-based approaches have been explored within combustion prior to their modern popularity, as shown with early work in chemistry tabulation [72].

Similar to how Chemkin [31] and Cantera [32] spurred the development of increasingly sophisticated representations of combustion processes with detailed thermo-chemical models, such as GRI Mech [73], growing accessibility to community-established ML frameworks, such as TensorFlow [74], PyTorch [75] and JAX [76], has spawned increasingly sophisticated deep learning techniques, which culminated in the transformer architecture [77]. This easily parallelizable architecture, along with well-established empirical scaling relationships [78] between computing power, data size, and ML capabilities, has spurred the development of language-based ML foundation models with increasing model complexity that can match and outperform human performance in text-related AI tasks (such as verbal reasoning, programming, and writing) [41]. The popularity and versatility of ML foundation models have recently inspired researchers across different scientific domains to employ language-based foundation models as autonomous scientist, such as in biology [45], material science [46], chemical synthesis [47], and mathematical science [48]. The current frontier involves extending foundation models beyond language through the introduction of multi-modal capabilities, *i.e.*, the ability of a single model to accept and combine different modalities of inputs (such as text, image, time-series signals, and flowfields). This summarizes the progress of ML and AI up until the present.

To project future developments, we express ML and AI progress in terms of world scopes (WS) [79], as shown in Fig. 2. WS1 represents

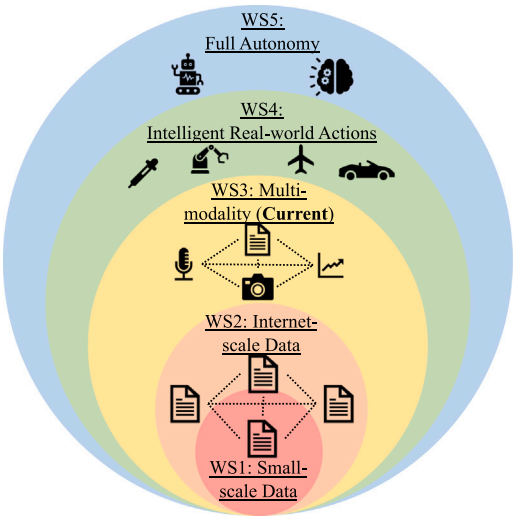

Fig. 2. Progress in ML and AI measured as world scopes (WS) [79].

the early stages of AI development, where researchers develop new ML methods [55,80], and discover new applications for these methods on small and privately accessible datasets. Within combustion and related fields, many recent developments (such as physics-informed NNs [81], as well as ML-based turbulence [82] and chemistry [83] modeling) fall within this scope.

Since ML models have been shown, across different fields, to scale in predictive capabilities with increasing amounts of high-quality training data [84], internet-scale data represents the next stage of progress, as seen with the ImageNet dataset [64]. With the public availability of thermo-chemical data [28,29], chemical kinetic models [21–25,73,85–89], software packages [31,32], as well as 2D [90] and 3D [91] combustion flowfields, the combustion ML research community is showing signs of progressing towards WS2.

In more mature AI fields, different domains are starting to overlap in research efforts in multi-modal AI (WS3), as seen with the increasing intersection of language and computer vision research in recent multi-modal foundation models [41]. This type of cross-disciplinary research represents the next frontier in many scientific domains, where efforts in integrating AI language models into scientific applications are becoming increasingly popular [45–48].

WS4 covers future directions in intelligent real-world actions, where AI systems are expected to reach human-level reasoning when interacting with the physical world. Within this scope, some exploratory work in employing RL within virtual environments [92] and specific real-world tasks [93] have shown success. While similar control techniques have been proposed within both virtual [94] and real-world [95] combustion systems, many of these approaches lack the versatility, generality, and reasoning capabilities of multi-modal foundation models [42]. As such, there are still many open questions within WS4, which have begun to be partially addressed by integrating AI agents (*i.e.*, autonomous intelligent systems performing human-like tasks) with external tools and systems [96] for interacting with the physical world to create embodied AI agents [97].

Finally, WS5 represents fully autonomous and trustworthy AI systems [79] that can be deployed without human monitoring or intervention, which has still not been demonstrated in any capacity. If ML and AI progresses to this stage, the existence of a fully autonomous AI agent could result in the automation of various experimental and computational procedures required for adapting to potential challenges being addressed by combustion researchers.

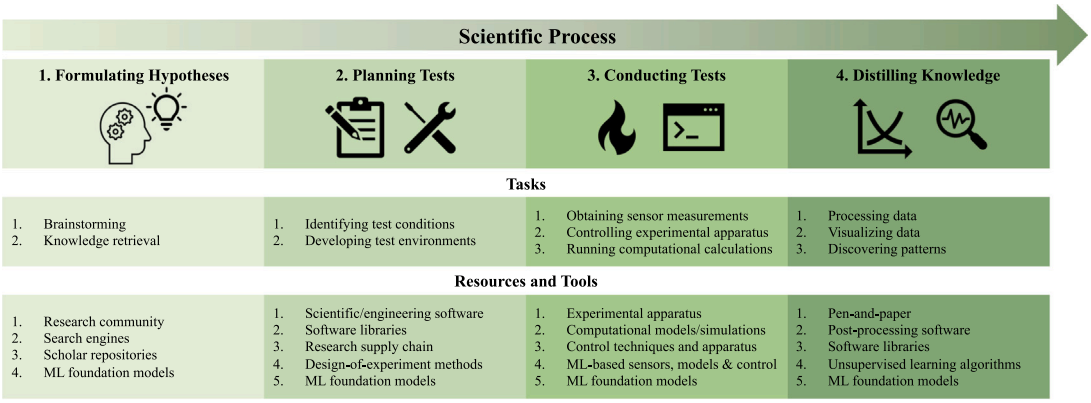

Fig. 3. Stages in scientific process of knowledge discovery, along with tasks, resources and tools (including ML methods and foundation models).

3. ML in combustion science and engineering

In the previous section, we provided a holistic overview of ML and AI techniques, along with factors that have led to their increasing maturity. Here, we provide a review of recent efforts in integrating ML methods towards catalyzing combustion science and engineering.

The systematic discovery of scientific and engineering knowledge can be broken down to the following steps [98], see Fig. 3: (i) formulating hypotheses, (ii) planning tests, (iii) performing tests, and (iv) distilling knowledge. So far, developments in combustion ML have largely been restricted towards aiding scientists in tasks related to distilling knowledge and performing tests. Combustion ML methods have largely been unsuitable for planning tests and generating scientific hypotheses, which are typically been performed exclusively by human scientists within a healthy research and development environment, occasionally with specialized tools such as design-of-experiment methods [99]. As such, we discuss different combustion ML techniques for distilling knowledge in Section 3.1, with methods for aiding scientific tests in Section 3.2.

3.1. Distilling knowledge

Early data processing techniques, such as linear regression [54], resulted in valuable contributions including the Arrhenius equation [56] and transport properties [57], which was used to further understanding of chemical kinetics. By extending linear regression with sparse regularization approaches (that introduce constraints during the optimization process in order to obtain desirable, and often simpler, model behavior), sparse regression can be employed for model discovery by identifying optimal linear combinations of analytical terms within equations [100]. Sparse symbolic regression has been popularized within physical domains through recent efforts in developing automated equation discovery suites for dynamical systems and astrophysics [101,102]. Within combustion, this approach has been used for constructing closure models for turbulent-chemistry interactions [103]. Recent work [104] has shown that models derived via symbolic regression can provide a similar functional form to a model derived via analytical arguments, which demonstrates a promising path towards an automated alternative to mathematical reasoning for discovering combustion-relevant models.

Similarly, dimensionality reduction approaches, such as PCA [58], have a long history of distilling high-dimensional data into more tangible forms for visualization, interpretation, and computation [106]. These approaches are especially useful in modern computational combustion problems for reducing the cost of modeling high-dimensional thermo-chemical state spaces, especially with stiff chemistry associated with fast chemical time-scales. For example, Kumar et al. [107] performed detailed numerical simulation (DNS) of a 3D methane/air slot-burner configuration, with the thermo-chemical state reduced by PCA. This approach resulted in approximately 70-fold and six-fold reduction in computational cost and memory, respectively, while maintaining reasonable accuracy in reconstructed instantaneous and statistical quantities (related to turbulence, flame-structure, and turbulence-chemistry interaction). The same approach was then extended towards more complex fuels, including *n*-heptane and primary reference fuel, in a 2D DNS of an engine-relevant configuration [108]. Thus, progress has been made towards automating computational acceleration for simulations of increasingly realistic combustion systems.

Beyond extensions to challenging configurations, researchers have focused on improving these PCA-based approaches for combustion applications. Jonnalagadda et al. [109] suggested the consideration of fourth-order statistical moments, instead of conventional second-order statistical moments within the PCA formulation in order to better capture anomalous data points within the thermo-chemical state space. Zdybał et al. [105] developed an automated approach for selecting PCA input variables that maximized a proxy quantity for the low-dimensional manifold quality, shown in Fig. 4. These improvements highlight the adeptness of combustion researchers in modifying dimensionality reduction tools towards solving important combustion problems.

When improving ML approaches, it is crucial to consider trade-offs. One common trade-off involves interpretability and expressiveness. While linear reduction operations within PCA can result in an interpretable manifold that can be transported with conservation equations, the linear reconstruction operation can restrict the expressiveness of this data-driven approach. In relation to this, Abdelwahid et al. [110] replaced the linear reconstruction component of PCA with a multilayer perceptron (MLP) to demonstrate that the addition of a non-linear mapping could outperform conventional tabulated chemistry approaches in large-eddy simulations (LES) of an NH_3/H_2/air non-premixed flame configuration. As shown in Fig. 5, Perry et al. [111] unified linear decomposition operations found within PCA with non-linear reconstruction and closure modeling operations into a unified NN architecture, *i.e.*, a co-optimized ML manifold (CMLM), that could be optimized together in a supervised learning manner for an *a priori* study involving a 3D premixed flame DNS. They showed that this co-optimization approach resulted in more accurate predictions of filtered quantities, when compared to a naïvely integrated PCA+NN approach.

Non-linear dimensionality reduction techniques may be more suitable in applications where linear relationships between input data and reduced manifold are not necessary. For example, autoencoder NN architectures can be viewed as non-linear extensions of PCA [112]. Han et al. [113] demonstrated that non-linear manifolds generated by convolutional autoencoders could be employed towards computing the flame stability index directly from raw flame images captured by high-speed cameras. Iemura et al. [114] demonstrated that manifolds extracted from variational autoencoders could be employed towards

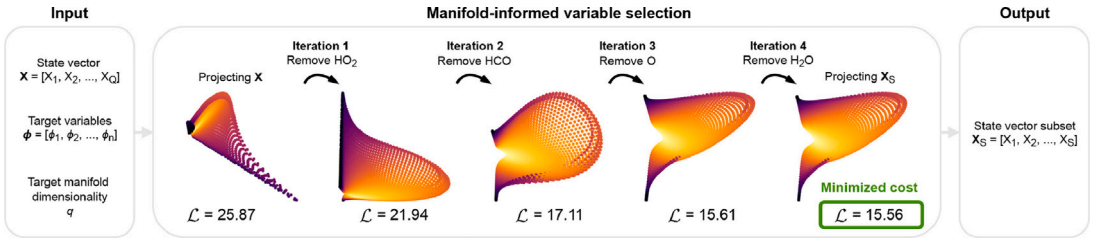

Fig. 4. Iterative improvement of PCA manifolds through optimization of manifold quality criterion $\mathcal{L}$. Reprinted from [105], Copyright 2023, with permission from Elsevier.

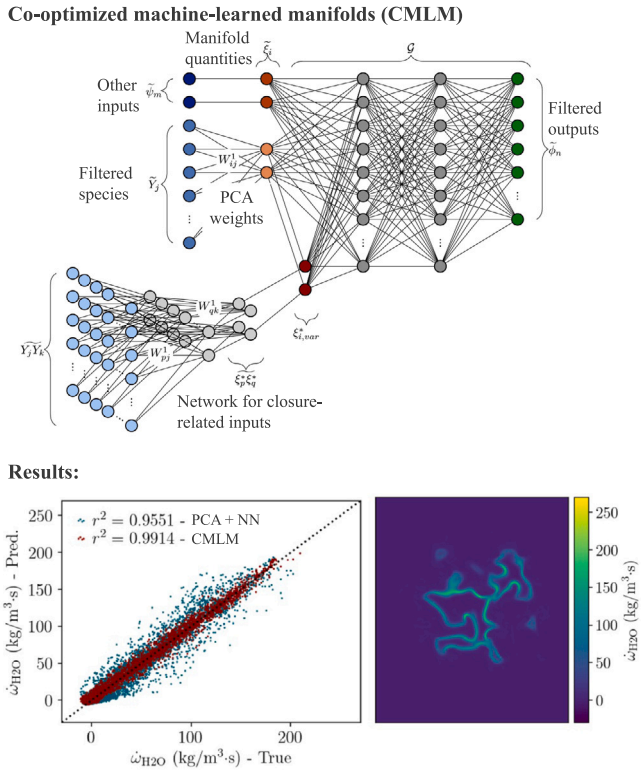

Fig. 5. PCA, non-linear mapping, and closure modeling represented by a unified NN architecture that is shown to outperform a naïvely integrated PCA+NN approach in a closure modeling problem. Adapted from [111], Copyright 2022, with permission from Elsevier.

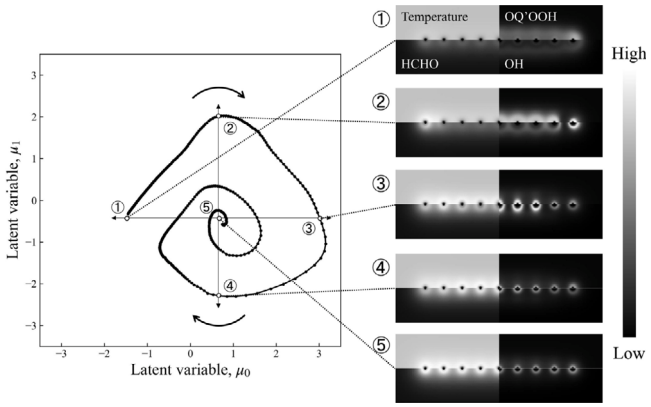

Fig. 6. Variational autoencoder latent space, *i.e.*, reduced manifold representation of five different states of cool flame oscillation dynamics. Reprinted from [114], Copyright 2023, with permission from Elsevier.

analyzing the trajectory of cool flame oscillation phenomena. This latent space, shown in Fig. 6, demonstrates that a coherent trajectory for representing five different states of cool flame oscillation dynamics could be extracted from this deep learning approach. However, we note that deep learning manifolds require careful treatment due to potential over-fitting, which result in non-generalizable approaches, and are much less interpretable than PCA since the relationship between the inputs and the latent space is no longer represented through a linear combination [115]. These trade-offs could be important for developing generalizable and inspectable manifolds for analyzing novel combustion conditions.

Beyond dimensionality reduction, clustering algorithms have also been employed and developed for automating knowledge distillation across a wide range of problems within combustion science and engineering, see Fig. 1. For example, k-means [59], one of the earliest clustering algorithms, is still being actively developed towards solving combustion problems. Ullman et al. [116] extended k-means by specifically segmenting highly reactive regions within combustion flowfields through the inclusion of a cluster-dependent scaling matrix constructed from the Jacobian of source terms. Dave et al. [117] demonstrated the use of Vector Quantization PCA, which is a PCA-based clustering technique, in automatically detecting important features within DNS datasets of MILD combustion. Wang et al. [118] demonstrated that clustering algorithms such as affinity propagation [119] could be used for extracting representative data samples that can be used in place of the full dataset for costly post-processing methods. This cluster-based sampling strategy has also been employed, together with the aforementioned dimensionality techniques, for lowering computational costs of high-fidelity simulations [120,121]. Clustering techniques can also be employed to improve other ML methods, as seen with local approaches [122] that employ PCA on localized clusters identified via data-driven means.

We can expect to see continued discoveries and improvements for combustion applications of both clustering and dimensionality reduction techniques, given the proficiency demonstrated by combustion researchers in this field. This can help automate pattern discovery in increasingly complex datasets as diagnostic and simulation data grow in capabilities, and accelerate high-fidelity calculations towards more realistic configurations.

3.2. Performing tests

As mentioned in Section 2, the proliferation of supervised learning approaches is catalyzed by the increasing availability of software packages for ML. Packages such as scikit-learn [63] provide an easy-to-use interface to bug-free implementations of ML models suited for small datasets [55,80]. In contrast, deep learning packages such as TensorFlow [74], PyTorch [75], and JAX [76] provide access to easily-parallelizable and GPU-compatible methods that provide higher predictive accuracy with datasets in more complex modalities involving representations of data (such as time-series measurements, spatio-temporal flowfields, chemical graphs, and flame images), but with higher computational cost and lower interpretability [123]. Similar to the discussion presented in Section 3.1, choosing between the two

approaches is typically based on trade-offs involving computational resources, inference-speed, and interpretability requirements.

Both categories of supervised learning approaches are useful for performing combustion experimental and computational tests as virtual sensors, control systems, or surrogate models for chemistry, combustion closures, and manifold representations [18]. However, many studies focusing on combustion chemistry tend to employ small-data ML approaches since their resulting datasets are typically represented with tabulated modalities. For example, linear regression with sparse regularization has been demonstrated as a surrogate for molecular dynamics simulations in the aggregation of poly-cyclic aromatic compounds [124]. Gradient-boosted decision trees [125], arguably the most suitable method for tabular data [126], has been recently employed as a surrogate for experimental measurements in determining the rate constants of alkyl ester and atomic hydrogen in a study focused on bio-diesel components [127]. These supervised learning algorithms provide accurate and cost-effective alternatives to costly simulations and experiments, as long as sufficiently representative data is available.

Despite the effectiveness of tree-based ML approaches, NN-based approaches involving MLPs can sometimes be preferred, due to the ability to adapt to problems via cost-effective transfer learning techniques [128], which involve additional training of pre-trained model weights with newer datasets. For example, Zhang et al. [129] demonstrated a 25-fold improvement in data efficiency when employing transfer learning on MLPs in a reaction optimization problem. With data from experimental measurements, Kildare et al. [130] found that transfer learning resulted in approximately 90% improvement in training data efficiency, when predicting temperature fields in a jet-in-hot-coflow flame configuration with deep learning models. The improved data efficiency offered by transfer learning reduces the labor required for collecting combustion measurements. Since backpropagation [55] used in training NNs are a form of automatic differentiation (AD), benefits in computation can also be obtained when using NNs with other AD-based techniques. This was demonstrated when Su et al. [131] combined a modified chemical reaction NN [132] with adjoint sensitivity methods for kinetics parameter optimization of jet propulsion and *n*-heptane fuels.

Deep learning models can also be suitable tools for exploring more complex data modalities. Kim et al. [133] developed a graph NN architecture, which accepted molecular graphs (with atomic and bond energy information) as inputs, for solving a cetane number prediction problem. As illustrated in Fig. 7, they showed that the hidden layers could be manipulated towards identifying functional groups responsible for different cetane numbers. Given the range of conditions involving data availability, problem scope, complementary techniques, and chemical representations within combustion chemistry, we can expect to see continued applications and improvements of a wide range of low-cost and accurate supervised learning techniques for aiding with predictive tasks within this domain.

Combustion experimentalists have also demonstrated a preference for applying a variety of deep learning models towards aiding a wide range of measurement tasks, due to their convenient ability to directly accept raw images with minimal feature engineering. For example, Yoon et al. [134] trained a convolutional NN (CNN) [135], with a loss function based on proper orthogonal decomposition (POD), for denoising time-resolved flame emission spectroscopy data from a methane/air flat flame burner. The inclusion of the POD loss term was shown to reduce the quantified uncertainty of gas property predictions. Gharib et al. [136] proposed to combine a CNN with an MLP for developing a flame-lift-off detector for real-time flame stability diagnostics with high-speed cameras. In another real-time application, Dai et al. [137] explored the use of three different deep learning architectures (based on CNNs and long-short-term-memory (LSTM) networks [138]) for reconstructing 3D soot temperature and volume fraction from flame images. These deep learning approaches show similar predictive capabilities when compared to a traditional volumetric reconstruction

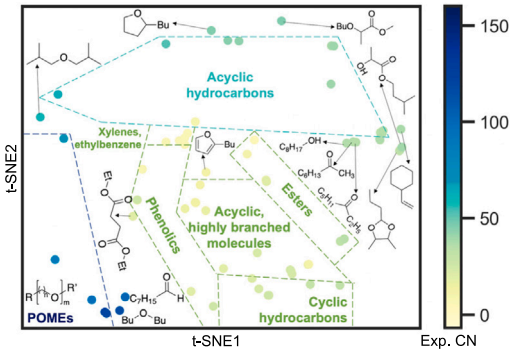

Fig. 7. Embeddings extracted from graph NN hidden layers categorized to different functional groups with different cetane numbers. Reprinted from [133], Copyright 2023, with permission from Elsevier.

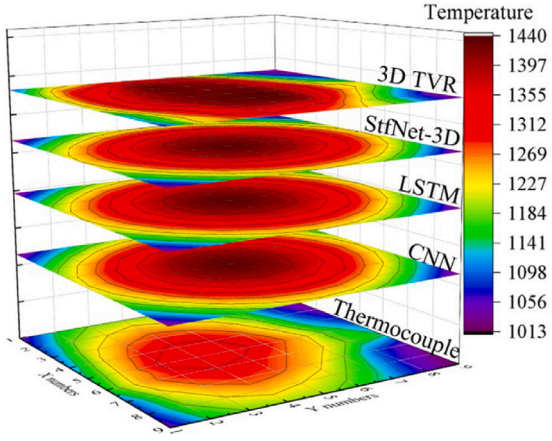

Fig. 8. Comparison of different deep learning approaches in reconstructing soot temperature. Reprinted from [137], Copyright 2023, with permission from Elsevier.

method (as shown in Fig. 8) with up to over 40-fold improvements in computational speed-up. Cheng et al. [139] proposed combining traditional reconstruction techniques with a generative adversarial network (GAN) for reconstructing 3D soot fields from luminosity signals at reasonable computational costs. These real-time ML-based systems are promising for developing improved diagnostics systems that could be relevant for automating tasks related to monitoring emissions and important phenomena in emerging combustion systems.

Deep learning models also enable researchers and engineers to fuse heterogeneous data sources. For example, Chung et al. [140] employed a CNN model that could combine sparse experimental Schlieren measurements of a laser-ignited rocket combustor [141] with corresponding time-averaged flowfields, in order to predict the spatio-temporal dynamics of ignition kernels. The resulting deep learning model was then combined with a stochastic differential equation for generating spatially-resolved ignition probability maps with affordable ensemble calculations, see Fig. 9.

The preference for deep learning techniques is also observed in closure modeling, which help in providing researchers and engineers with more accurate predictions and analyses. In an *a priori* turbulent super-resolution study, Nista et al. [142] demonstrated that a GAN, regularized with physics-based loss terms, for LES closure modeling of unresolved momentum and scalar fluxes could generalize well across different out-of-distribution conditions involving different Reynolds and Karlovitz numbers in a methane/air premixed flame configuration. Given the popularity of physics-informed approaches for turbulent closure modeling [143], Chung et al. [91] benchmarked metrics (related to

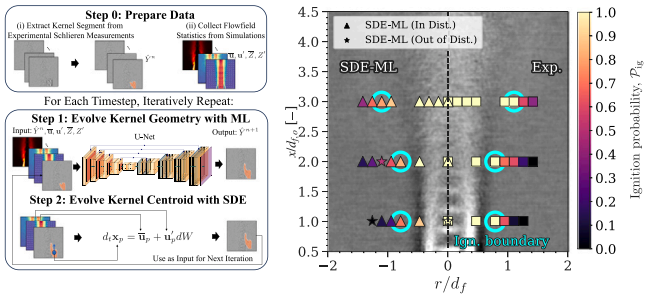

Fig. 9. Predictions of ignition probability using a physics-embedded stochastic differential equation (SDE)-ML model, showing (left) SDE-ML architecture and (right) ensemble-averaged experimental Schlieren measurements. Adapted from [140], Copyright 2024, with permission from Elsevier.

predictive accuracy and computational cost) from five different deep-learning based super-resolution models with a 2 TB public turbulent combustion dataset [144]. Through neural scaling analysis, the authors found that predictive benefits from a physics-based ML approach scale in log-linear fashion with computational power, as presented in Fig. 10. The results provide empirical evidence that disagrees with the notion that physics-based approaches are useful mostly in small data and model scenarios [145].

Despite these advances in deep-learning-based closure modeling, employing GPU-accelerated deep learning techniques within *a posteriori* simulations can be challenging. One challenge involves the need for heterogeneous (both CPU and GPU) numerical solvers. As such, there is still a need to explore ML techniques that can be easily deployed for inference within CPU-based solvers. In an *a posteriori* study involving turbulent spray flames, Yao et al. [146] demonstrated that existing gradient-boosted decision trees [125] could be deployed within numerical solvers with similar performances to MLPs. Given their accuracy compared to conventional algebraic closure models, these different ML approaches show a promising path for more accurate simulations of realistic combustion systems, while reducing the need for selecting and developing new closure models for under-explored conditions.

Manifold methods involving tabulated thermo-chemistry present another ML opportunity for integration with numerical solvers. This differs from the PCA-based manifold methods in Section 3.1, as the manifolds are typically constructed via domain knowledge in contrast to an automated data-driven approach. Since early work in ML-based tabulation [72,147], there have been several developments. Readshaw et al. [148] presented a robust data curation strategy for MLP-based tabulated chemistry within LES of a premixed swirl burner. The resulting MLP was shown to be approximately 14 times faster than conventional finite-rate chemistry with a low-memory footprint. Chi et al. [65] further investigated the computational benefits of MLP-based tabulated chemistry, and found in an *a posteriori* study that there was a trade-off compared to conventional tabulation: MLPs were more memory efficient but showed more time complexity. Nikolaou et al. [149] further formalized these trade-offs by developing a criterion for choosing between conventional and NN-based tabulation methods via the use of empirical scaling laws. These efforts in ML-based tabulated chemistry highlight a mature ML application within combustion, where systematic recommendations can be made for optimal deployment of ML tools.

Another challenge in the integration of ML methods with numerical solvers involves the presence of spurious non-physical errors that negatively impact numerical stability [143]. In an *a priori* study, Nakazawa et al. [150] demonstrated that the introduction of domain knowledge via physics-guided loss terms can diminish these spurious errors in source term modeling within a V-flame configuration.

In *a posteriori* LES of a premixed jet flame configuration, Ho et al. [151] demonstrated that the combined use of algebraic and NN-based closure models ameliorated spurious errors from the NN, see Fig. 11. For modeling stiff chemical source terms, Owoyele and Pal [83] proposed the employment of customized architectures for treating systems of ordinary differential equations (ODEs), and showed that neural ODE architectures were useful in mitigating these spurious errors within 0D calculations of a hydrogen/air mixture. Vijayarangan et al. [152] later extended the application of neural ODEs towards evolving non-linear manifolds (generated via autoencoders) of hydrogen/air and ethylene/air chemical systems within 0D calculations. Regression-based errors can also be avoided by focusing on prediction tasks that only indirectly influence flowfield quantities [153]. For example, Malpica Galassi et al. [154] employed NNs as surrogates for costly Jacobian eigenvector calculations required within a computational singular perturbation. Another approach for addressing these regression errors would involve improving the interpretability of ML-based models. However, these developments within this research direction remain an active research pursuit across most deep learning domains [155].

Many of the combustion ML applications mentioned in this section have focused on applications where errors do not result in safety or ethical concerns. In relation to this, combustion control is a domain where additional robust ML measures are required to ensure safe deployment [18]. As such, this field of combustion ML applications is relatively immature compared to virtual sensing, data post-processing, and computational modeling applications. In a computational study involving thermo-acoustic instabilities, Alhazmi and Sarathy [94] demonstrated that a recent development of safe RL [157] could be employed towards active control of combustors in a model-free manner. Zhan et al. [95] developed a safe RL scheme tailored for intelligently controlling real-world coal-fired power plants, which helps maintain high combustion efficiency during runtime. Despite the challenges posed by developing and testing safe ML techniques in potentially hazardous situations, further research into this largely under-explored, but important, area could lead to promising control techniques that increase the range of effective operating conditions in sustainable combustion technology.

To summarize this section, there are numerous nuances that require expert treatment when developing ML methods for complementing the different fields found within combustion science and engineering. This provides opportunities for further combustion ML research to develop improved autonomous tools for treating different constraints posed by data modalities, computational resources, interpretability requirements, and safety constraints.

4. Towards an AI-based combustion scientist

In Section 3, we reviewed progress of ML techniques for assisting combustion researchers and engineers in computational tasks related to distilling knowledge and performing tests, such as dimensionality reduction, surrogate modeling, and data partitioning. ML foundation models offer greater versatility by extending their applications to tasks typically reserved for human action as AI agents, as shown in Fig. 3. This provides opportunities for automating a broader range of tasks that could catalyze combustion research towards addressing its most pressing challenges, such as the discovery of new fuel properties, improving energy conversion efficiencies, and exploring novel combustion strategies.

Here, we introduce readers to the current leading AI paradigm, involving ML foundation models in Section 4.1. In addition, we demonstrate and evaluate state-of-the-art multi-modal foundation models in performing scientific tasks within combustion in Section 4.2. Note that since ML foundation models are currently the leading approach in performing human-like tasks, we will now begin to use the terms ML foundation models and AI agents interchangeably.

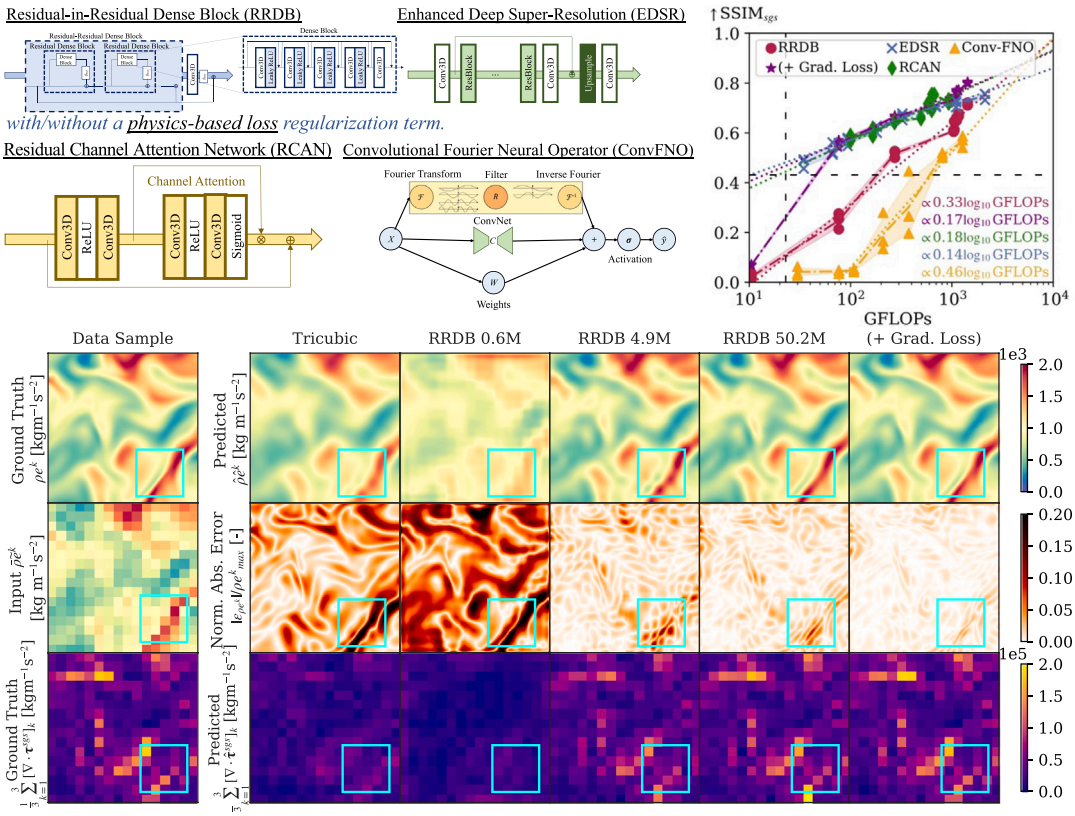

Fig. 10. Scaling behavior when scaling with computational power (measured in FLOPs) and number of model parameters in deep-learning based turbulent super-resolution application to represent subgrid-scale stresses.
Source: Adapted from [91].

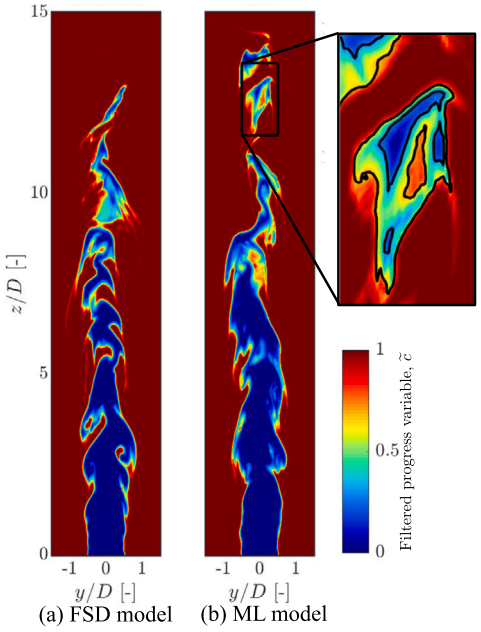

Fig. 11. LES of a premixed turbulent flame, showing instantaneous solutions of the filtered progress variable fields for simulations performed with (a) physics-based flame-surface-density (FSD) model [156] and (b) hybrid ML model; black isolines in zoomed view show regions in which the ML model is used. Reprinted from [151], Copyright 2024, with permission from Elsevier.

4.1. ML foundation models

An ML foundation model is a general purpose ML model (typically with >1B trainable parameters) that has been (i) pre-trained offline via self-supervised learning [158], *i.e.*, supervised learning with minimal labeling of large and diverse datasets (typically terabyte-scale), which can then be (ii) fine-tuned, *i.e.*, further offline training with supervised transfer learning involving smaller and more specific datasets for tailored downstream applications. Many foundation models also incorporate (iii) post-training treatment to improve their behavior beyond their offline training.

The most prominent example of foundation models are seen within language domains, where models such as GPT-3 [159] (with up to 175B parameters) are pre-trained with 570 GB of text data for next-token/word prediction. This model is later fine-tuned for performing human tasks, such as following instructions [160] and writing code [161]. Post-training, the model is continuously improved with RL from online feedback [160] to better respond to changing user behavior, and augmented with external tools and databases [96,162].

This section discusses potential developments in these three main components required to develop foundation models within combustion. Specifically, large scientific and engineering datasets that could be used for pre-training models are discussed in Section 4.1.1. Next, the benefits of transfer learning are demonstrated in Section 4.1.2. Finally, potential directions for integrating popular post-training within combustion is proposed in Section 4.1.3.

4.1.1. Large pre-training datasets for combustion

The first step in developing an ML foundation models involves pre-training on large volumes of minimally processed data via self-supervised learning techniques [158]. Examples of self-supervised learning tasks include next-word/token prediction in language domains [163],

image reconstruction in computer vision applications [164], and next-timestep prediction in dynamical systems [165]. Self-supervised learning enables ML models to harness large datasets in a cost-effective manner by focusing training on information inherent within data structures (such as smoothness or continuity within flowfields), instead of information provided by costly labeled datasets. The weights from the resulting pre-trained models can then be shared publicly and adapted for more complex tasks with smaller (and thus, more cost-effective) labeled datasets through transfer learning [128]. Thus, the key requirements of large-scale pre-training involve public repositories for data and model weights, as well as a collaborative scientific community that can maintain and contribute to public access of these weights and data.

For example, foundation models in language domains are typically pre-trained on data from the Common Crawl repository [166], which provides free access to petabytes of text data extracted from the World Wide Web. Weights of these models are then made available via an open-source platform [167], which also provides a benchmark leaderboard with up-to-date and transparent information on the capabilities of these models in a range of practical tasks. Both data and model weights, as well as auxiliary tools for improving access, currently heavily rely on contributions from the open-source community for continuous improvements and maintenance.

Beyond language, scientific domains have begun to adopt these practices in developing their own foundation models. For example, weather-related foundation models are pre-trained for next-timestep prediction problems using the petabyte-scale ERA5 dataset [168]. This has led to a growth in various weather-based pre-trained models such as GraphCast [165], Pangu [169], and FourCastNet [170]. Platforms and community contributions within scientific domains are not yet as mature as in language applications, as many of these datasets are stored in separate repositories with minimal tools for improving access. However, organized efforts such as WeatherBench [171] show promise in consolidating accessibility to these pre-trained models.

Combustion researchers have also long adopted community-based practices for accelerating scientific progress. For example, community contributions and shared resources have been prevalent within the chemical kinetics community. Here, finite-rate chemistry models, along with transport and thermodynamic properties, are easily accessible. In addition, open-source scientific codes such as Cantera [32] continuously improve access to chemical kinetics resources. Recently, these practices have started to gain traction within the turbulent combustion research community. This can be exemplified through BLASTNet [144], an open-source effort that aims to provide public access to ML-ready diverse DNS data by leveraging the combustion research community [91,172]. Currently, BLASTNet contains approximately 5 TB of data from nearly 40 different reacting and non-reacting flow configurations, and has been shown to be useful for benchmarking flowfield super-resolution [91], which is a combustion-related self-supervised learning task that has direct applications in closure modeling.

The motivation behind these different efforts in ML data and model sharing is linked to the discovery of scaling properties in deep learning-based models [78]. In essence, accuracy in deep learning models has been found to scale with an increasing number of diverse training samples and model size. We illustrate this through measuring the error of closure terms predicted by a residual-residual dense block deep learning architecture [164] (see Fig. 10), across different model sizes and training samples, in a super-resolution task with the aforementioned BLASTNet dataset [144]. The training target consists of 128^3 volumetric domains of density and 3D velocity components extracted from public DNS data, while the input consists of the corresponding Favre-filtered and downsampled (by eight-fold) data. All models are trained in a similar manner with the same adaptive learning rates ranging from `1e-4` to `1e-5` on the Adam optimizer [173] for 32,000 training iterations with data mini-batches of size 64. The resulting models are tested on an unseen test set of 173 volumetric data subsampled from 27 DNS configurations. Methods for sampling the training and test data are provided in further detail in Ref. [91].

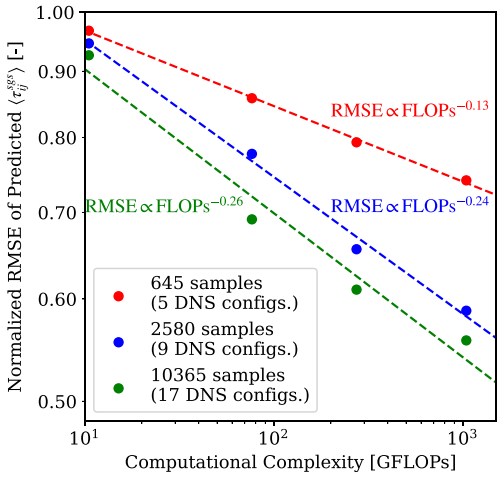

Fig. 12. Log–log plot of RMSE in ML test set predictions of the mean of six subgrid-scale stress components $\langle\tau_{ij}^{sgs}\rangle$ with computational complexity, measured in FLOPs, for three training sets of varying number of unique samples from different DNS configurations. RMSE is normalized by mean of $\langle\tau_{ij}^{sgs}\rangle$ across the test set.

Fig. 12 shows a log–log plot of normalized root mean square error (RMSE) in ML test set predictions of mean subgrid-scale stress $\langle\tau_{ij}^{sgs}\rangle$ with computational complexity, measured in floating point operations (FLOPs), for three training sets of varying number of unique samples from DNS configurations. For all training datasets, the ML error decreases with increasing computational complexity, following the power law RMSE $\propto$ FLOPs$^{-\alpha}$. This empirical power law property has also been observed in other ML domains [78] and was confirmed in a scientific ML application for reacting flows [91]. With an approximate 16-fold increase in unique training samples, α increases from 0.13 to 0.26. This indicates that the ML models exhibit improved scaling behavior when exposed to larger and more diverse datasets. This highlights the importance of scaling data availability for accelerating ML development within combustion.

4.1.2. Fine-tuning via supervised transfer learning

Even with access to large and diverse datasets, pre-training foundation models with state-of-the-art performance can incur large computational costs and requires laborious multi-GPU configuration for training. For instance, training the most prominent open-source AI model, *i.e.*, LLaMa-2 [43] (70B parameters) required 1.7M GPU-hours on Nvidia A100 GPUs. With recent advances in low-rank fine-tuning [174] and quantization [175], fine-tuning a 70B parameter model requires only approximately 100 GPU-hours. With this framework, one could easily develop foundation models towards a variety of language-based combustion tasks including writing chemistry/CFD code, summarizing scientific literature, and answering scientific inquiries.

In other domains, this is already being demonstrated in a growing number of studies [46,48]. For instance, Song et al. [46] developed a small dataset (consisting of 52k samples of material-science instruction and responses) that was used to fine-tune LLaMa-1 [176] models for performing tasks related to analyzing material science literature. Romera-Paredes et al. [48] augmented a 340B parameter foundation model [177] with a code testing module for discovering algorithms and heuristics for mathematical problems. These recent experiments demonstrate the potential of maturing ML technology in assisting with tasks beyond the scope described in Section 3, and show promise for greater automation within combustion science and engineering through AI agents.

We highlight the benefits of supervised fine-tuning via transfer learning by comparing predictions from two deep learning models

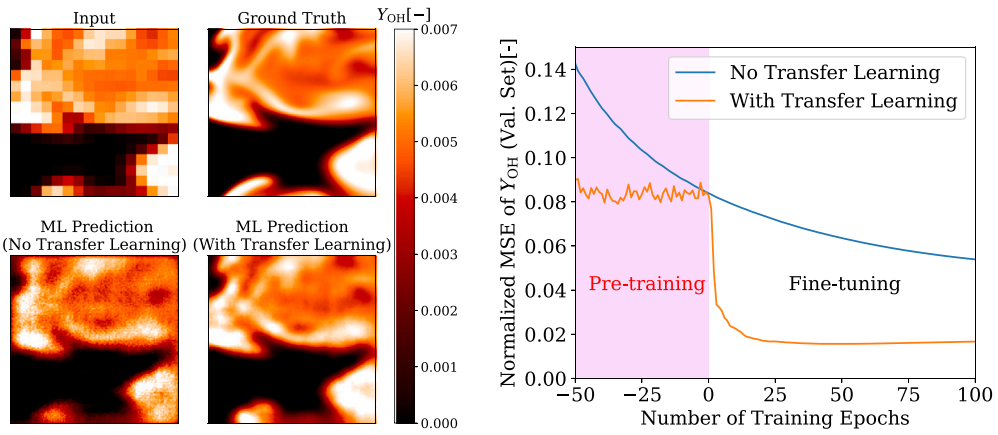

(a) Flowfields for Y_{OH} from ML predictions, along with corresponding input and ground truth target.

(b) Normalized MSE on validation set as a function of training epochs.

Fig. 13. Comparison of (a) flowfields and (b) normalized MSE of ML predictions on unseen inputs, without and with transfer learning, in super-resolution of OH mass fraction field.

(trained without and with transfer learning) in super-resolution of OH mass fraction Y_{OH}, as demonstrated with Fig. 13. In this ML task, the ground truth targets in the fine-tuning dataset are 17 samples of 2D slices (with 128^2 cells) obtained from public DNS data [91] of a premixed hydrogen/air flame configuration [178], with another 17 unseen samples kept for validating the models. These targets are Favre-filtered and downsampled by eight-fold to obtain corresponding inputs for the ML models. For both ML approaches, the Adam optimizer [173] was employed, with a learning rate of `1e−4` with batch size of 17, for training the same deep convolutional architecture [179] (with 2M parameters). Prior to fine-tuning, the model was pre-trained until the MSE loss plateaued, with 5528 total samples of velocity and density flowfields from 27 DNS configurations provided by the BLASTNet database [144]. In contrast, the model without transfer learning was pre-trained with 17 samples from the fine-tuning set, until the MSE loss matched the model with transfer learning. Information on the specific DNS configurations is discussed in further detail in Ref. [91]. During fine-tuning, both models are trained with 100 training epochs. Note that one epoch corresponds to the sum of training iterations required to loop through the training samples once, during gradient descent optimization.

Fig. 13(a) compares flowfield predictions for OH mass fraction, Y_{OH}, from the two ML models (without and with transfer learning), along with corresponding input and ground truth targets. When compared to the fine-tuned model, it can be seen that flowfield predictions from the ML model are noisier without transfer learning. A quantitative comparison of the prediction (normalized) mean squared error (MSE) from an unseen validation dataset, as a function of training epochs, between the two ML approaches is presented in Fig. 13(b). Here, it can be seen that transfer learning enables training to converge after approximately 20 epochs at a normalized MSE of 0.016, which is five-fold more computationally efficient and has 67% lower error, when compared to the model without transfer learning. These results, along with those from Fig. 12, highlight the importance of public pre-trained model weights and large pre-training datasets for optimizing computational and data efficiency, when developing combustion-related ML models. In this spirit, we make all ML models available through GitHub as supplementary material.

4.1.3. *Post-training treatment*

Even with large pre-training datasets and transfer learning-based approaches, ML foundation models can still behave erratically when predicting with new unseen inputs, often resulting in ML hallucinations [180] (*e.g.* outputting incoherent sentences, misrepresenting facts, inventing academic citations, and providing dangerous advice). In the case of language-based foundation models, these spurious errors in model predictions have led to practical and ethical concerns in information trustworthiness [181].

Similar concerns can arise if combustion ML models generate spurious predictions, especially since mathematical and data-driven models within combustion science and engineering are employed within safety critical systems in aerospace, power, and transportation sectors. An example of a spurious prediction in combustion ML is shown in Fig. 14(a), where nonphysical reaction predictions are seen in the inert region outside the central jet in a flowfield prediction from an ML-based solver [182]. Similar hallucinations have also been reported in other applications involving dynamical systems, including laminar [143] and turbulent [183] non-reacting flows.

Outside of combustion, several post-training strategies have been employed for ameliorating hallucinations within foundation models, by incorporating domain knowledge via integration with external tools [96] (where fine-tuned foundation models are used as an interface for more reliable tools to perform specific actions) and knowledge retrieval systems [162] (where fine-tuned foundation models interface with external databases to supplement information beyond offline training). Similar strategies in embedding physics domain knowledge via physics-informed ML [143] have been increasingly popular within combustion-related domains. External retrieval systems linking combustion knowledge databases with language-based foundation models have recently been explored [184]. However, post-training techniques involving the use of RL and online learning, for providing continuous updates to the foundation models from external feedback [160], have yet to be explored within combustion.

These different treatments could be employed within combustion applications to address spurious predictions illustrated in Fig. 14(a). As shown in Fig. 14(b), a gating ML model can be employed to identify subdomains where the ML-based solver is mispredicting. Then, the more challenging subdomains can be treated with a reliable external tool, such as a physics-based solver. This physics-based prediction can then be used to augment ML predictions — resulting in an overall improved prediction. Since numerical methods within this physics-based solver can incur large computational costs, the ML-based solver can be updated *in situ* via online learning with subdomain flowfields from the physics-based solver, in order to improve ML predictions with increasing number of timestep advancements. Similarly, the gating network can also be iteratively updated via online learning on the improved performance of the ML-based solver, which would reduce the reliance of this proposed post-training approach on costly numerical methods.

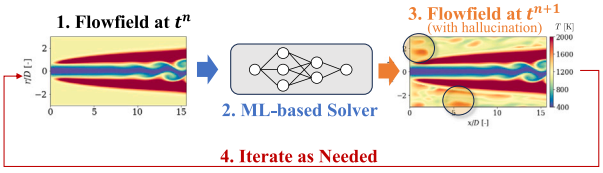

(a) ML hallucinations in autoregressive modeling resulting in spurious reactions and heat release (circled region).

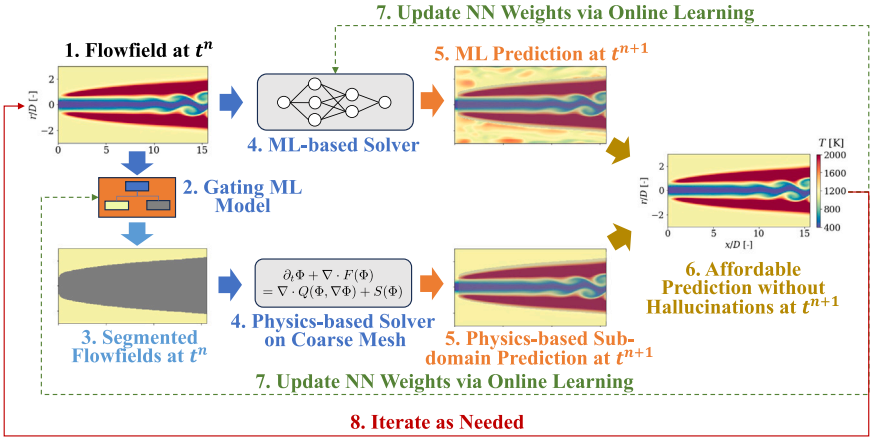

(b) Proposed strategy for aligning ML foundation models by combining domain knowledge and online learning.

Fig. 14. Spatio-temporal predictions of reacting flows via autoregressive modeling (a) resulting in spurious reactions and heat release, *i.e.*, ML hallucinations, in inert regions outside the central jet (circled region), and (b) proposed strategy for aligning ML foundation models by combining domain knowledge and online learning [182].

4.2. Assessing an AI combustion scientist

The previous section discusses the prospect of developing foundation models solely with combustion data. As indicated in Section 2, the progress of autonomous intelligent agents for performing general research tasks (including within combustion) might involve the development of multi-modal ML foundation models, where language-based models are integrated seamlessly with combustion-based models.

In order to assess the opportunities provided by AI agents, we perform an investigation of a computational combustion problem with the GPT-4 model [42]. This is performed by entering text prompts that request the AI agent to perform tasks related to the scientific process, *i.e.*, as listed in (i) formulating hypotheses, (ii) planning tests, (iii) performing tests, and (iv) distilling knowledge [18]. Selected results from this demonstration are shown in Fig. 15, while full prompts and responses are provided in the supplementary material.

For hypothesis formulation, when instructed to suggest computational combustion studies that can be feasibly performed with an AI agent, the agent provides six suggestions related to fuel mixtures, alternative fuels, emissions, fuel additives, knocking, and fuel pyrolysis. When providing these responses, the AI agent relies on the knowledge stored within the model weights during pre-training and fine-tuning. However, practical AI cannot solely rely on this stored knowledge, as this trained knowledge could be outdated for tasks that require more updated information. For this task, the AI agent uses tools (such as search engines) to augment its stored knowledge and to overcome its outdated training data. AI tool-use [96] and external knowledge retrieval systems [162] are popular solutions for reducing agent errors and boosting the range of agent actions, especially in tasks where next-token predictions are limited in accuracy and capabilities such as mathematical reasoning [185], and up-to-date knowledge retrieval [162].

When prompted to plan tests for a scientific investigation, the AI agent provides viable suggestions for fuel candidates involving ethanol-gasoline and bio-diesel mixtures, and lists correct sources [22,186] of chemical kinetic mechanisms for this study. However, we note that some links presented by the AI agent were inaccurate or broken, as seen in the supplementary material. This tendency of AI agents to hallucinate non-existent or incorrect information is an open research issue [180], which has so far been only partially ameliorated through a combination of tool-use [96], knowledge retrieval [162], and AI interpretability [155].

Another challenge that can arise within this step is the broad availability of chemical-kinetic models with different sensitivity ranges, as they could result in suboptimal solutions for emissions and complex configurations. This, however, can be partially addressed by more specific prompts and post-training treatments that provide the model with better context of the user request. As foundation models are typically trained on large corpora of publicly available code, the AI agent also has significant code development capabilities, as shown by the present AI agent's capabilities in writing starter code for investigating ethanol iso-octane fuel/air mixtures of different ratios in a homogeneous reactor. However, the initial versions of the code contained defects that arise from AI hallucinations. These issues can be partially resolved through iterative interactions with the AI agent when conducting tests.

When conducting tests, the AI agent is prompted to modify the starter code to output various quantities-of-interests (related to the thermo-chemical state) from the 0D calculations of ethanol and iso-octane fuel/air mixture. The generated code is then manually executed on a local computer, with code outputs fed into the AI agent for subsequent prompts. For more complex setups, this step can be automated with API access and tool-use mechanisms that involve AI software packages such as LangChain [187]. In this demonstration, we have restricted the configuration to only consider computational experiments, since the interaction of AI agents with the physical world through embodied AI [97] remains an open research issue. When performing this task, the AI agent is unable to create code to output all the suggested quantities of interest at once, and requested modifications are prompted in a piece-wise fashion. This dependency on prompting strategy is a feature of numerous ML foundation models [188]. As seen with previous tasks, the AI agent hallucinates several quantities-of-interests, such as flame speed and combustion stability, that are not of relevance to 0D calculations.

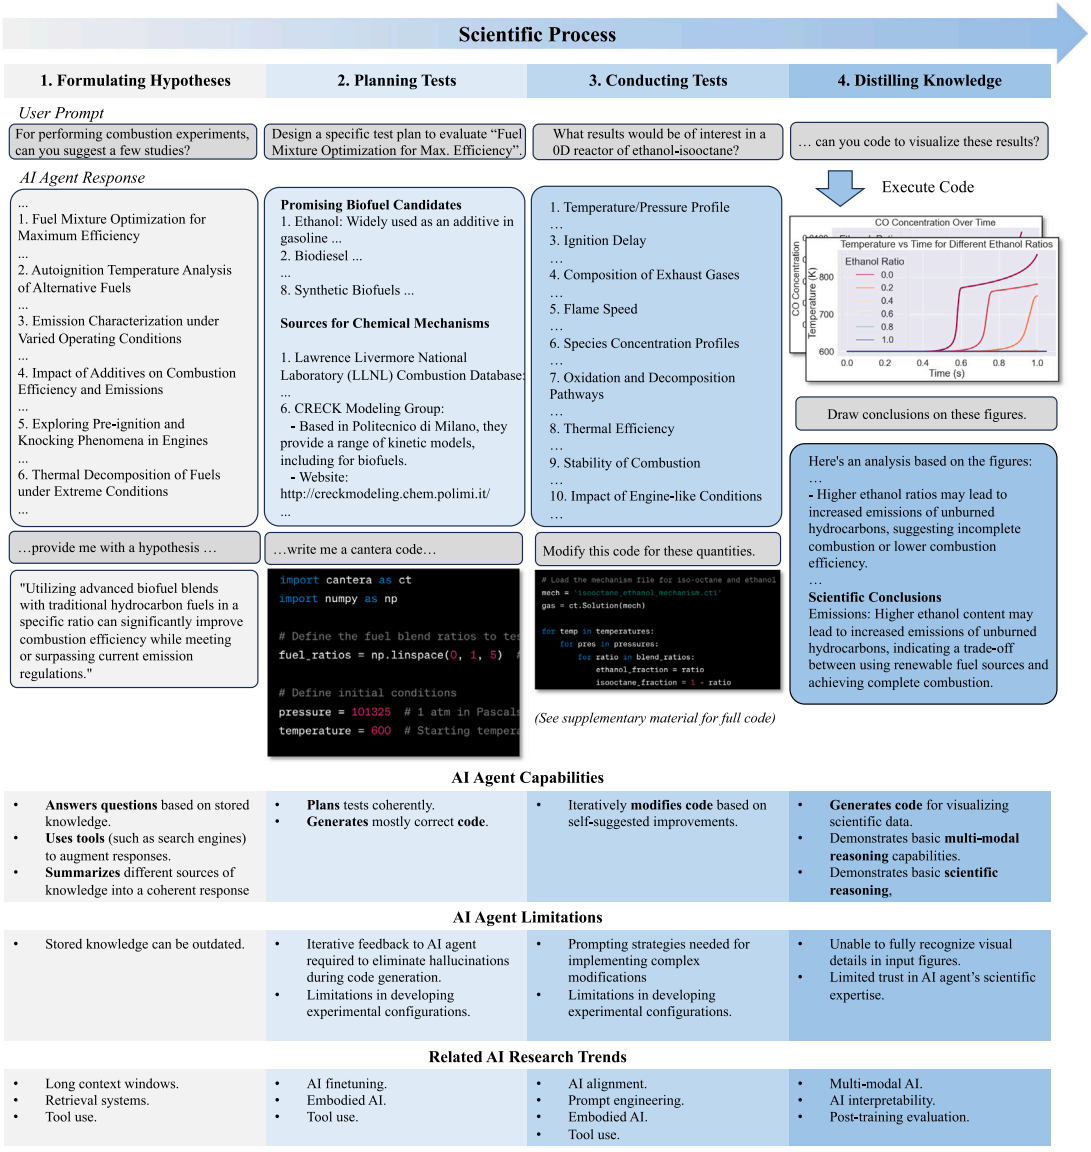

Fig. 15. Capabilities and limits of an AI agent identified during the present combustion-related scientific demonstration, with related research trends. Full prompts and responses are provided as supplementary material.

To distill knowledge from the 0D reactor calculations, the AI agent is prompted to write code for analyzing simulation results. The AI agent is seen to employ common `python` packages for visualizing the temporal behavior of temperature, pressure, and selected chemical. Next, we assess multi-modal capabilities. Here, the AI agent is correct in pointing out that the presence of unburned hydrocarbons in high ethanol fuel indicates a trade-off when employing bio-fuel-blends.

To further inspect the knowledge within the AI agents in solving combustion problems, we quantitatively evaluate the performance of the GPT-4 model in answering 26 short questions related to combustion fundamentals in thermo-chemistry, combustion, and chemical kinetics that can be used to evaluate combustion knowledge at undergraduate- and graduate-levels. We also compare GPT-4 performance with the latest language foundation model available to the present authors, *i.e.*, Gemini 1.5 [189], which uses an improved mixture-of-experts architecture that has been shown to improve the performance of language-based foundation models [190]. The results from this analysis are summarized in Table 1, with prompts, responses, and ground truth answers presented in the supplementary material. Here, we also compare these combustion scores with scores from public coding [161] and reading comprehension [191] evaluation suites. We choose to present these additional scores to further assess state-of-the-art AI performance in tasks that could be performed by an autonomous combustion scientist, specifically in combustion fundamentals, programming, and complex communication.

Table 1
Quantitative evaluation of two language foundation models on various tasks.

Question type	GPT-4 (2023)	Gemini 1.5 (2024)
Coding[a] [161]	67.0%	74.4%
Reading comprehension[a] [191]	80.9%	78.9%
Thermo-chemistry	46.7%	80.0%
1D flames	66.7%	83.3%
Chemical kinetics	60.0%	60.0%
Total combustion	53.8%	76.9%

[a] Indicates scores taken from the GPT-4 and Gemini 1.5 technical reports [42,189].

GPT-4 and Gemini 1.5 can answer 53.8% and 76.9% of the combustion questions correctly, respectively. Specifically, Gemini 1.5 can

User Prompt

What is the equilibrium reaction for the formation of nitrous oxide from nitrogen and oxygen? Express the equilibrium constant of nitric oxide reaction in terms of partial pressures of each species.

GPT-4 Response (Correct)

$$\mathrm{N_2(g) + O_2(g) \rightleftharpoons 2N_2O(g)}$$
$$K = \frac{[\mathrm{N_2O}]^2}{[\mathrm{N_2}][\mathrm{O_2}]}$$

Gemini 1.5 Response (Correct)

$$\mathrm{N_2(g) + O_2(g) \rightleftharpoons 2N_2O(g)}$$
$$K = \frac{[\mathrm{N_2O}]^2}{[\mathrm{N_2}][\mathrm{O_2}]}$$

(a) Prompt on emissions.

User Prompt

Provide an algebraic expression for the heat of formation of methanol using the group additivity model.

GPT-4 Response (Wrong)

$$\Delta \mathrm{H_f}^\circ(\mathrm{CH_3OH}) = \Delta \mathrm{H_f}^\circ(\mathrm{CH_3-}) + \Delta \mathrm{H_f}^\circ(\mathrm{-OH}) + \mathrm{I}$$

Gemini 1.5 Response (Wrong)

$$\Delta \mathrm{H_f}^\circ(\mathrm{CH_3OH}) = \Delta \mathrm{H_f}^\circ(\mathrm{C-(H)_3}) + \Delta \mathrm{H_f}^\circ(\mathrm{O-(C)(H)}) + \mathrm{I}$$
$$\Delta \mathrm{H_f}^\circ(\mathrm{CH_3OH}) \approx -42.2\,\mathrm{kJ/mol} - 158.6\,\mathrm{kJ/mol} = -200.8\,\mathrm{kJ/mol}$$

(b) Prompt on Benson group additivity.

User Prompt

Ozone (O3) decomposes exothermically to produce molecular O2 via a simple chain mechanism: $\mathrm{O_3 + M \longrightarrow O + O_2 + M}$, $\mathrm{O + O_3 \longrightarrow 2\,O_2}$

(i) Derive an ODE expression for the time rate of change of concentration of O atom.

(ii) What happens when O is assumed to be under quasi-steady state?

GPT-4 Response (Wrong)

$$\frac{d[\mathrm{O}]}{dt} = k_1[\mathrm{O_3}][\mathrm{M}] - k_2[\mathrm{O}][\mathrm{O_3}]$$
$$k_1[\mathrm{O_3}][\mathrm{M}] = k_2[\mathrm{O}][\mathrm{O_3}]$$
$$[\mathrm{O}] = \frac{k_1[\mathrm{M}]}{k_2}[\mathrm{O_3}]$$

Gemini 1.5 Response (Wrong)

$$\frac{d[\mathrm{O}]}{dt} = k_1[\mathrm{O_3}][\mathrm{M}] - k_2[\mathrm{O}][\mathrm{O_3}]$$
$$0 \approx k_1[\mathrm{O_3}][\mathrm{M}] - k_2[\mathrm{O}][\mathrm{O_3}]$$
$$[\mathrm{O}] = k_1/k_2$$

(c) Prompt on quasi-steady-state assumption.

Fig. 16. Abridged responses from both GPT-4 and Gemini 1.5 models, when prompted with fundamental questions on combustion. Full prompts and responses are provided as supplementary material.

answer combustion questions with similar performance to comprehension (78.9%) and coding-based (74.4%) tasks, while GPT-4 lacks combustion understanding when compared to comprehension (80.9%) and coding (67.0%). Nevertheless, these results show that state-of-the-art ML foundation models are continuously improving towards human level capabilities in understanding combustion.

Three prompts and corresponding responses from both GPT-4 and Gemini 1.5 are presented in Fig. 16. Fig. 16(a) shows that both models understand nitrous oxide reactions, and corresponding knowledge regarding chemical equilibria. While both models demonstrate understanding of heat of formation in Fig. 16(b), both models produce errors in different manners. GPT-4 answers with the groups $\mathrm{CH_3-}$ and $-\mathrm{OH}$, instead of the correct $\mathrm{C-(H)_3(O)}$ and $\mathrm{O-(C)(H)}$ [192], respectively, while Gemini 1.5 only mispredicts with $\mathrm{C-(H)_3}$. In Fig. 16(c), both models are shown to relate chemical reactions with chemical rates of production, and understand the implications of the quasi-steady-state (QSS) assumption. However, both models fail in different manners to provide the correct expression for $[\mathrm{O}] = k_1[\mathrm{M}]/k_2$. This inability to reason mathematically, even for simple operations, is well known within foundation models [180].

From the quantitative and qualitative behavior of these state-of-the-art models in Table 1 and Fig. 16, we can infer that AI models might be useful in solving combustion questions, coding, and communicating, but will still require significant intervention via human interaction, especially with the inconsistent errors resulting from these black-box models. As such, AI agents today are not yet ready for full automation within scientific research and engineering, but can, in their present form, serve as co-pilots and knowledge engines for assisting combustion researchers and engineers. However, the use of ML foundation models for automation of scientific tasks is still an under-explored topic within combustion. In Section 5, we highlight several research directions that aim towards developing autonomous AI agents.

5. Future opportunities for combustion AI

5.1. Discovery of new ML and AI tasks

As discussed in Section 3, researchers have successfully discovered new ML methods for various combustion problems. Methods have been developed for automatically extracting data-driven manifolds from unlabeled data for scientific interpretation [113,114] and computational acceleration [107,108] tasks. Labeled data has also been employed by off-the-shelf and combustion-tailored approaches for surrogate modeling purposes, particularly for complementing experimental sensors [134,136,137,139], closure modeling [91,142,146,150], tabulated chemistry [65,148,149], and chemical optimization [129, 131,132]. Future efforts could be directed towards discovering self-supervised learning techniques that would advance pre-training of ML foundation models for combustion, as previous studies [193–195] have shown that the predictive performance of these models can benefit from the discovery of new domain-specific pre-training tasks. Within combustion, self-supervised learning has begun to be explored, as seen with experimental reconstruction [137,139], turbulent super-resolution [91, 142], and denoising [134] applications. Further discovery of these tasks would be essential for developing effective, broadly accessible, and general purpose tools for accelerating combustion science through task automation.

5.2. Dataset curation

As shown in Section 4.1, deep learning models scale in predictive performance with increasing size and diversity of data (see Fig. 12 for ML-modeling of subgrid-scale stresses). Compared to other ML domains involving language and vision, the development of large datasets within

scientific domains can be a challenging task. Fortunately, our combustion community has previously shown success in building large repositories of publicly accessible data. This is evidenced from early works involving JANAF thermo-chemical data [29] and the NIST database [28], and also within more recent efforts, as seen with the growing availability of measurements from chemical kinetic experiments that have led to the diverse ecosystem of publicly accessible chemical kinetic mechanisms for different fuels.

Recently, the availability of low-cost data storage has also prompted the creation of multi-dimensional combustion flowfield datasets for ML and reduced-order modeling purposes [90,91]. These datasets have focused on high-fidelity reacting flow simulations, with many opportunities in storing and disseminating multi-fidelity and heterogeneous data remaining under-explored [172]. With further support from research institutions, government, and private sectors, these resources can develop further in both diversity and accessibility. Beyond these scientific datasets, there is also the opportunity to extract combustion knowledge from vast text repositories containing combustion-relevant literature [196] and source code [32,197], which complements recent efforts in combining language models with combustion knowledge [184]. These could be used to improve the mediocre combustion knowledge shown in Table 1 via fine-tuning.

Despite the existing availability of these methods, these datasets still need to be curated, pre-processed, and formatted to ML-ready standards. The labor-intensive nature of these suggested tasks is one reason there has not yet been a fine-tuned foundation model specifically for combustion, even with manageable fine-tuning costs and widely available pre-trained weights. Nevertheless, these efforts are sufficiently tractable, and provide a relatively affordable approach for developing improved AI agents that could better help researchers in brainstorming sustainability solutions, summarizing combustion literature, or developing and testing basic chemistry code modules.

5.3. Evaluation and interpretability

If the aforementioned methods and resources reach sufficient maturity for developing combustion ML foundation models, there would be a need to develop techniques for assessing the predictive capabilities and trustworthiness of the resulting models. This is especially true given current limitations in prediction tasks that result in spurious errors and hallucinations [180]. The naïve approach for evaluation would involve metrics related to the loss function (such as cross-entropy or mean-squared error) [62].

However, as the developed foundation models possess more general capabilities in performing multiple tasks, there is a need to develop and employ empirical metrics on curated benchmarks to sample the model performance. This is analogous to measuring human intelligence and competency through a series of examinations across different topics. As partially demonstrated via Table 1, empirical benchmarks [198,199] are built to assess abstract behavior such as the ability to code competently, reason mathematically, and communicate honestly. These evaluation suites could be adapted towards measuring more relevant abilities, such as the ability of writing chemistry code, solving combustion problems, and controlling combustion systems safely.

The popularity of empirical evaluations within AI and ML stems from the difficulty in interpreting black-box deep learning models. However, we note that there have been several emerging attempts in interpreting deep learning models through the examination of their underlying mathematical operations, *i.e.*, mechanistic interpretation [155], which has shown some degree of success in small models and simple architectures. Given that (i) significant domain expertise is required to develop domain-specific empirical evaluations and (ii) combustion researchers have long been interested in developing interpretable reduced-order models (see Section 3), there will be opportunities for scientific contributions in both evaluation and interpretability approaches. The combination of these approaches would help researchers assess, inspect, and trust AI agents for performing scientific tasks. This is especially important for combustion problems due to potential health risks in noxious emission and safety concerns within any propulsion and power-generating combustion systems.

5.4. Towards complete autonomy

Combustion researchers will reach a similar level of AI maturity (see Fig. 2) as other scientific domains [45–47], if fine-tuned multimodal combustion foundation models were to be developed. As discussed in Section 4, this would offer partial automation of many combustion-related tasks. These tools would still require human intervention due to issues related to hallucination and model transparency. In addition, these tools would likely be applicable only within computational and theoretical domains with tasks (such as including scientific code development, mesh generation, and chemical kinetic mechanism generation), especially due to the large number of open challenges related to deploying embodied AI within the physical world [79], particularly when considering potential hazards that could arise from failure in real-world combustion systems. While the pathway to overcoming these technical challenge is only being partially addressed, exponential growth in computing capabilities and increased interdisciplinary collaborations can help pave a way towards autonomous combustion science and engineering.

6. Summary

In this article, we discuss past, present, and potential developments within ML and AI to catalyze combustion research, especially in the face of climate-related challenges. After providing a holistic overview of ML and AI techniques and discussing important contributions that led to their increasing maturity in Section 2, we review recent progress in utilizing ML methods across various combustion application tasks in Section 3. By employing these ML methods, combustion researchers and engineers have discovered new ways of automating tasks related to processing measurements, optimizing chemical-kinetic models, discovering thermo-chemical manifolds, and accelerating computations.

In many of these combustion ML applications, ML models with limited parameters and training data have been shown to generate erroneous predictions. Promising research directions for addressing this challenge have typically involved the injection of scientific and engineering knowledge, as well as the development of larger ML models and datasets. We explore potential improvements in predictive capabilities that could be obtained from pre-training and fine-tuning large ML models, via a demonstration involving super-resolution of a reacting flow configuration. In addition, we discuss potential developments in RL and online learning that could further address issues related to ML requirements on safety and interpretability.

Recently, general-purpose ML foundation models have shown success in completing a wide range of human tasks. This has motivated numerous scientific and engineering fields towards developing and adapting these foundation models as AI agents within their respective domains. To this end, in Section 4 we examine the capabilities of leading foundation models in completing scientific tasks and answering questions related to combustion. By considering specific combustion problems to quantitatively evaluate the ability of current AI agents in answering fundamental questions related to thermo-chemistry, combustion physics, and chemical kinetics, we show that ML foundation models are able to complete approximately 60% to 80% of assigned tasks correctly, with errors resulting from ML hallucinations.

To manage these ML hallucinations and treat ML deficiencies, we suggest research opportunities in Section 5 to further adapt foundation models for automating combustion science and engineering through (i) the discovery of new self-supervised learning tasks, (ii) the curation of additional datasets for pre-training and fine-tuning, (iii) providing public access to pre-trained model weights for reproducibility and optimizing computational and data efficiency, and (iv) the development

of interpretability and custom evaluation suites that can be employed towards understanding future combustion AI agents. While results from this demonstration show that a fully autonomous AI combustion scientist is still not yet possible, they illustrate opportunities for current AI agents to serve as co-pilots for assisting combustion researchers in facing pressing challenges within the field.

We hope that the recommendations in this article, along with discussions involving previous work, can inform readers on methods for developing AI techniques that can lead to multi-modal scientific agents for solving combustion problems. This would in turn pave a pathway towards autonomous combustion research.

In the words of the titan Prometheus in a famed ancient Greek play [200], "Fire has proved for men a teacher in every art, their grand resource", but whether fire can prove a teacher for the next-generation of AI systems is left to the efforts of combustion researchers.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We gratefully acknowledge financial support from the U.S. Department of Energy (DOE) Office of Basic Energy Sciences under award DE-SC0022222, and from the DOE National Nuclear Security Administration under award DE-NA0003968.

Appendix A. Supplementary data

See https://github.com/IhmeGroup/pretrain_finetune for code used for the demonstrations discussed in Section 4.1. Full prompts and responses used to assess the ML foundation model in Section 4.2 are provided as supplementary material.

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.proci.2024.105730.

References

[1] K.R. Hein, New challenges for research in a changing energy market, Proc. Combust. Inst. 29 (1) (2002) 393–398.
[2] R.F. Sawyer, Science based policy for addressing energy and environmental problems, Proc. Combust. Inst. 32 (1) (2009) 45–56.
[3] K. Kohse-Höinghaus, Combustion, chemistry, and carbon neutrality, Chem. Rev. 123 (8) (2023) 5139–5219.
[4] R. Cracknell, S. Ciatti, S. Dorofeev, R. Eggels, K. McManus, K. Nakata, Decarbonization of mobility, including transportation and renewable fuels, Proc. Combust. Inst. 39 (1) (2023) 1–9.
[5] M. Hermesmann, T. Müller, Green, turquoise, blue, or grey? Environmentally friendly hydrogen production in transforming energy systems, Prog. Energy Combust. Sci. 90 (2022) 100996.
[6] P. Su-ungkavatin, L. Tiruta-Barna, L. Hamelin, Biofuels, electrofuels, electric or hydrogen?: A review of current and emerging sustainable aviation systems, Prog. Energy Combust. Sci. 96 (2023) 101073.
[7] A. Valera-Medina, H. Xiao, M. Owen-Jones, W. David, P. Bowen, Ammonia for power, Prog. Energy Combust. Sci. 69 (2018) 63–102.
[8] H. Levinsky, Why can't we just burn hydrogen? Challenges when changing fuels in an existing infrastructure, Prog. Energy Combust. Sci. 84 (2021) 100907.
[9] A.A. Perpignan, A. Gangoli Rao, D.J. Roekaerts, Flameless combustion and its potential towards gas turbines, Prog. Energy Combust. Sci. 69 (2018) 28–62.
[10] A.K. Agarwal, A.P. Singh, A. García, J. Monsalve-Serrano, Challenges and opportunities for application of reactivity-controlled compression ignition combustion in commercially viable transport engines, Prog. Energy Combust. Sci. 93 (2022) 101028.
[11] T.F. Wall, Combustion processes for carbon capture, Proc. Combust. Inst. 31 (1) (2007) 31–47.
[12] N. Liu, J. Lei, W. Gao, H. Chen, X. Xie, Combustion dynamics of large-scale wildfires, Proc. Combust. Inst. 38 (1) (2021) 157–198.
[13] S.S. Goldsborough, S. Hochgreb, G. Vanhove, M.S. Wooldridge, H.J. Curran, C.-J. Sung, Advances in rapid compression machine studies of low- and intermediate-temperature autoignition phenomena, Prog. Energy Combust. Sci. 63 (2017) 1–78.
[14] K. Kohse-Höinghaus, Clean combustion: Chemistry and diagnostics for a systems approach in transportation and energy conversion, Prog. Energy Combust. Sci. 65 (2018) 1–5.
[15] H.J. Curran, Developing detailed chemical kinetic mechanisms for fuel combustion, Proc. Combust. Inst. 37 (1) (2019) 57–81.
[16] P. Domingo, L. Vervisch, Recent developments in DNS of turbulent combustion, Proc. Combust. Inst. 39 (2) (2023) 2055–2076.
[17] B. Yang, Towards predictive combustion kinetic models: Progress in model analysis and informative experiments, Proc. Combust. Inst. 38 (1) (2021) 199–222.
[18] M. Ihme, W.T. Chung, A.A. Mishra, Combustion machine learning: Principles, progress and prospects, Prog. Energy Combust. Sci. 91 (2022) 101010.
[19] C. Cavallotti, Automation of chemical kinetics: Status and challenges, Proc. Combust. Inst. 39 (1) (2023) 11–28.
[20] G.P. Smith, D.M. Golden, M. Frenklach, N.W. Moriarty, B. Eiteneer, M. Goldenberg, C.T. Bowman, R.K. Hanson, S. Song, W.C. Gardiner Jr., V.V. Lissianski, Z. Qin, GRI-Mech 3.0, 1999, http://www.me.berkeley.edu/gri_mech/.
[21] X. You, F.N. Egolfopoulos, H. Wang, Detailed and simplified kinetic models of n-dodecane oxidation: The role of fuel cracking in aliphatic hydrocarbon combustion, Proc. Combust. Inst. 32 (1) (2009) 403–410.
[22] M. Mehl, W.J. Pitz, M. Sjöberg, J.E. Dec, Detailed Kinetic Modeling of Low-Temperature Heat Release for PRF Fuels in an HCCI Engine, SAE Technical Paper 2009-01-1806, 2009.
[23] Y. Li, C.-W. Zhou, K.P. Somers, K. Zhang, H.J. Curran, The oxidation of 2-butene: A high pressure ignition delay, kinetic modeling study and reactivity comparison with isobutene and 1-butene, Proc. Combust. Inst. 36 (1) (2017) 403–411.
[24] W. Tang, Y. Zhai, Q. Wang, C. Shao, I. Gorbatenko, S. Gail, A. Nicolle, S.M. Sarathy, An experimental and kinetic modeling study of the impact of nitric oxide and hydrogen on iso-Octane combustion, Combust. Flame 260 (2024) 113196.
[25] A. Stagni, C. Cavallotti, S. Arunthanayothin, Y. Song, O. Herbinet, F. Battin-Leclerc, T. Faravelli, An experimental, theoretical and kinetic-modeling study of the gas-phase oxidation of ammonia, React. Chem. Eng. 5 (2020) 696–711.
[26] C.F. Curtiss, J.O. Hirschfelder, Integration of stiff equations, Proc. Natl. Acad. Sci. USA 38 (3) (1952) 235–243.
[27] D.J. Gardner, D.R. Reynolds, C.S. Woodward, C.J. Balos, Enabling new flexibility in the SUNDIALS suite of nonlinear and differential/algebraic equation solvers, ACM Trans. Math. Software (2022).
[28] P. Linstorm, NIST chemistry webbook, NIST standard reference database number 69, J. Phys. Chem. Ref. Data Monogr. 9 (1998) 1–1951.
[29] D.R. Stull, H. Prophet, JANAF Thermochemical Tables, Report NSRDS–NBS 37, National Bureau of Standards, Washington, DC, 1971.
[30] S. Gordon, B.J. McBride, Computer Program for Calculation of Complex Chemical Equilibrium Compositions, Rocket Performance, Incident and Reflected Shocks and Chapman-Jouguet Detonations, NASA Technical Report SP-273, 1971.
[31] R.J. Kee, J.A. Miller, T.H. Jefferson, CHEMKIN: A General-Purpose, Problem-Independent, Transportable, FORTRAN Chemical Kinetics Code Package, Sandia National Laboratories Report, SAND80-8003, 1980.
[32] D.G. Goodwin, H.K. Moffat, R.L. Speth, B.W. Weber, Cantera: An object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes, 2023, https://www.cantera.org. Version 2.4.0.
[33] T. Lu, C.K. Law, Toward accommodating realistic fuel chemistry in large-scale computations, Prog. Energy Combust. Sci. 35 (2) (2009) 192–215.
[34] S.-H. Lam, D.A. Goussis, Understanding complex chemical kinetics with computational singular perturbation, Proc. Combust. Inst. 22 (1) (1989) 931–941.
[35] T. Turanyi, Reduction of large reaction mechanisms, New J. Chem. 14 (11) (1990) 795–803.
[36] T. Lu, C.K. Law, A directed relation graph method for mechanism reduction, Proc. Combust. Inst. 30 (1) (2005) 1333–1341.
[37] P. Pepiot-Desjardins, H. Pitsch, An efficient error-propagation-based reduction method for large chemical kinetic mechanisms, Combust. Flame 154 (1) (2008) 67–81.
[38] B. Bhattacharjee, D.A. Schwer, P.I. Barton, W.H. Green, Optimally-reduced kinetic models: Reaction elimination in large-scale kinetic mechanisms, Combust. Flame 135 (3) (2003) 191–208.
[39] T. Jaravel, H. Wu, M. Ihme, Error-controlled kinetics reduction based on non-linear optimization and sensitivity analysis, Combust. Flame 200 (2019) 192–206.
[40] T. Poinsot, Prediction and control of combustion instabilities in real engines, Proc. Combust. Inst. 36 (1) (2017) 1–28.
[41] R. Bommasani, D.A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M.S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill, et al., On the opportunities and risks of foundation models, 2021, arXiv preprint arXiv:2108.07258.

[42] OpenAI, GPT-4 technical report, 2023, arXiv preprint arXiv:2303.08774.

[43] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, et al., LLaMa 2: Open foundation and fine-tuned chat models, 2023, arXiv preprint arXiv:2307.09288.

[44] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H.W. Chung, C. Sutton, S. Gehrmann, et al., PaLM: Scaling language modeling with pathways, J. Mach. Learn. Res. 24 (240) (2023) 1–113.

[45] R. Luo, L. Sun, Y. Xia, T. Qin, S. Zhang, H. Poon, T.-Y. Liu, BioGPT: Generative pre-trained transformer for biomedical text generation and mining, Briefings Bioinform. 23 (6) (2022) bbac409.

[46] Y. Song, S. Miret, H. Zhang, B. Liu, HoneyBee: Progressive instruction finetuning of large language models for materials science, in: Proc. Conf. Empir. Methods Nat. Lang. Process., 2023, pp. 5724–5739.

[47] N.J. Szymanski, B. Rendy, Y. Fei, R.E. Kumar, T. He, D. Milsted, M.J. McDermott, M. Gallant, E.D. Cubuk, A. Merchant, et al., An autonomous laboratory for the accelerated synthesis of novel materials, Nature (2023) 1–6.

[48] B. Romera-Paredes, M. Barekatain, A. Novikov, M. Balog, M.P. Kumar, E. Dupont, F.J. Ruiz, J.S. Ellenberg, P. Wang, O. Fawzi, et al., Mathematical discoveries from program search with large language models, Nature (2023) 1–3.

[49] C.M. Bishop, Pattern Recognition and Machine Learning, Springer-Verlag, Berlin, Heidelberg, 2006.

[50] S. Russell, P. Norvig, Artificial Intelligence: A Modern Approach, Prentice Hall Press, Upper Saddle River, NJ, 1995.

[51] L. Zhou, Y. Song, W. Ji, H. Wei, Machine learning for combustion, Energy AI 7 (2022) 100128.

[52] M. Alirameza, C.R. Koch, M. Shahbakhti, Modeling, diagnostics, optimization, and control of internal combustion engines via modern machine learning techniques: A review and future directions, Prog. Energy Combust. Sci. 88 (2022) 100967.

[53] A. Nicolle, S. Deng, M. Ihme, N. Kuzhagaliyeva, E.A. Ibrahim, A. Farooq, Mixtures recomposition by neural nets: A multidisciplinary overview, J. Chem. Inf. Model. 64 (3) (2024) 597–620.

[54] A.M. Legendre, Nouvelles Méthodes Pour La Détermination Des Orbites Des Comètes, Courcier, Paris, France, 1805.

[55] D.E. Rumelhart, G.E. Hinton, R.J. Williams, Learning representations by back-propagating errors, Nature 323 (6088) (1986) 533–536.

[56] S. Arrhenius, Über die dissociationswärme und den einfluss der temperatur auf den dissociationsgrad der elektrolyte, Z. Phys. Chem. 4 (1) (1889) 96–116.

[57] B.J. McBride, Coefficients for Calculating Thermodynamic and Transport Properties of Individual Species, NASA Technical Report TM-4513, 1993.

[58] K. Pearson, On lines and planes of closest fit to systems of points in space, Lond. Edinb. Dublin Philos. Mag. J. Sci. 2 (11) (1901) 559–572.

[59] S. Lloyd, Least squares quantization in PCM, IEEE Trans. Inf. Theory 28 (2) (1982) 129–137.

[60] J.C. Sutherland, A. Parente, Combustion modeling using principal component analysis, Proc. Combust. Inst. 32 (1) (2009) 1563–1570.

[61] C.J.C.H. Watkins, Learning from Delayed Rewards (Ph.D. thesis), King's College, Oxford, 1989.

[62] I. Goodfellow, Y. Bengio, A. Courville, Deep Learning, The MIT Press, Cambridge, MA, 2016.

[63] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine learning in python, J. Mach. Learn. Res. 12 (2011) 2825–2830.

[64] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, ImageNet: A large-scale hierarchical image database, in: Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2009, pp. 248–255.

[65] C. Chi, X. Xu, D. Thévenin, Efficient premixed turbulent combustion simulations using flamelet manifold neural networks: A priori and a posteriori assessment, Combust. Flame 245 (2022) 112325.

[66] B. Ruscic, R.E. Pinzon, G. von Laszewski, D. Kodeboyina, A. Burcat, D. Leahy, D. Montoy, A.F. Wagner, Active thermochemical tables: Thermochemistry for the 21st century, J. Phys.: Conf. Ser. 16 (2005) 561–570.

[67] M. Frenklach, Transforming data into knowledge—process informatics for combustion chemistry, Proc. Combust. Inst. 31 (1) (2007) 125–140.

[68] G.L. Goteng, N. Nettyam, S.M. Sarathy, CloudFlame: Cyberinfrastructure for combustion research, in: International Conference on Information Science and Cloud Computing Companion, 2013, pp. 294–299.

[69] T. Varga, T. Turányi, E. Czinki, T. Furtenbacher, A.G. Császár, ReSpecTh: A joint reaction kinetics, spectroscopy, and thermochemistry information system, in: Proceedings of the 7th European Combustion Meeting, 2015, pp. 1–5.

[70] E. Ramalli, T. Dinelli, A. Nobili, A. Stagni, B. Pernici, T. Faravelli, Automatic validation and analysis of predictive models by means of big data and data science, Chem. Eng. J. 454 (2023) 140149.

[71] A. Krizhevsky, I. Sutskever, G.E. Hinton, ImageNet classification with deep convolutional neural networks, Adv. Neural Inf. Process. Syst. 25 (2012).

[72] F. Christo, A. Masri, E. Nebot, Artificial neural network implementation of chemistry with PDF simulation of H_2/CO_2 flames, Combust. Flame 106 (4) (1996) 406–427.

[73] M. Frenklach, H. Wang, M. Goldenberg, G.P. Smith, D.M. Golden, C.T. Bowman, R.K. Hanson, W.C. Gardiner, V. Lissianski, GRI-Mech: An Optimized Detailed Chemical Reaction Mechanism for Methane Combustion, Report GRI-95/0058, Gas Research Institute, 1995.

[74] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G.S. Corrado, A. Davis, J. Dean, M. Devin, et al., TensorFlow: Large-scale machine learning on heterogeneous systems, 2015, https://www.tensorflow.org.

[75] A. Paszke, S. Gross, G.S. Chintala, Pytorch: An open source machine learning framework, 2016, https://pytorch.org.

[76] J. Bradbury, R. Frostig, P. Hawkins, M.J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, Q. Zhang, JAX: composable transformations of python+numpy programs, 2018, http://github.com/google/jax.

[77] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, Adv. Neural Inf. Process. Syst. 30 (2017).

[78] J. Hoffmann, S. Borgeaud, A. Mensch, E. Buchatskaya, T. Cai, E. Rutherford, D.d.L. Casas, L.A. Hendricks, J. Welbl, A. Clark, et al., An empirical analysis of compute-optimal large language model training, Adv. Neural Inf. Process. Syst. 35 (2022) 30016–30030.

[79] Y. Bisk, A. Holtzman, J. Thomason, J. Andreas, Y. Bengio, J. Chai, M. Lapata, A. Lazaridou, J. May, A. Nisnevich, et al., Experience grounds language, in: Proc. Conf. Empir. Methods Nat. Lang. Process, 2020, pp. 8718–8735.

[80] R. Messenger, L. Mandell, A modal search technique for predictive nominal scale multivariate analysis, J. Amer. Statist. Assoc. 67 (340) (1972) 768–772.

[81] M. Raissi, P. Perdikaris, G.E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, J. Comput. Phys. 378 (2019) 686–707.

[82] C.J. Lapeyre, A. Misdariis, N. Cazard, D. Veynante, T. Poinsot, Training convolutional neural networks to estimate turbulent sub-grid scale reaction rates, Combust. Flame 203 (2019) 255–264.

[83] O. Owoyele, P. Pal, Chemnode: A neural ordinary differential equations framework for efficient chemical kinetic solvers, Energy AI 7 (2022) 100118.

[84] A. Halevy, P. Norvig, F. Pereira, The unreasonable effectiveness of data, IEEE Intell. Syst. 24 (2) (2009) 8–12.

[85] N.M. Marinov, C.K. Westbrook, W.J. Pitz, Detailed and global chemical kinetics model for hydrogen, in: S.H. Chan (Ed.), Transport Phenomena in Combustion, Taylor & Francis, 1996, pp. 118–127.

[86] E. Ranzi, P. Gaffuri, T. Faravelli, P. Dagaut, A wide-range modeling study of n-heptane oxidation, Combust. Flame 103 (1995) 91–106.

[87] H.J. Curran, P. Gaffuri, W.J. Pitz, C.K. Westbrook, A comprehensive modeling study of n-heptane oxidation, Combust. Flame 114 (1998) 149–177.

[88] H. Wang, X. You, A.V. Joshi, S.G. Davis, A. Laskin, F. Egolfopoulos, C.K. Law, USC mech version II. High-temperature combustion reaction model of H2/CO/C1-C4 compounds, 2017, http://ignis.usc.edu/USC_Mech_II.htm.

[89] E. Ranzi, A. Frassoldati, R. Grana, A. Cuoci, T. Faravelli, A. Kelley, C. Law, Hierarchical and comparative kinetic modeling of laminar flame speeds of hydrocarbon and oxygenated fuels, Prog. Energy Combust. Sci. 38 (4) (2012) 468–501.

[90] C.M. Huang, 2D benchmark reacting flow dataset for reduced order modeling exploration [Data set], 2020, University of Michigan - Deep Blue Data.

[91] W.T. Chung, B. Akoush, P. Sharma, A. Tamkin, K.S. Jung, J. Chen, J. Guo, D. Brouzet, M. Talei, B. Savard, et al., Turbulence in focus: Benchmarking scaling behavior of 3D volumetric super-resolution with BLASTNet 2.0 data, Adv. Neural Inf. Process. Syst. 36 (2023) 77430–77484.

[92] V. Mnih, K. Kavukcuoglu, D. Silver, A.A. Rusu, J. Veness, M.G. Bellemare, A. Graves, M. Riedmiller, A.K. Fidjeland, G. Ostrovski, et al., Human-level control through deep reinforcement learning, Nature 518 (7540) (2015) 529–533.

[93] H. Kim, M. Jordan, S. Sastry, A. Ng, Autonomous helicopter flight via reinforcement learning, Adv. Neural Inf. Process. Syst. 16 (2003).

[94] K. Alhazmi, S.M. Sarathy, Adaptive phase shift control of thermoacoustic combustion instabilities using model-free reinforcement learning, Combust. Flame 257 (2023) 113040.

[95] X. Zhan, H. Xu, Y. Zhang, X. Zhu, H. Yin, Y. Zheng, DeepThermal: Combustion optimization for thermal power generating units using offline reinforcement learning, Proc. AAAI Conf. Artif. Intell. 36 (4) (2022) 4680–4688.

[96] T. Schick, J. Dwivedi-Yu, R. Dessi, R. Raileanu, M. Lomeli, E. Hambro, L. Zettlemoyer, N. Cancedda, T. Scialom, Toolformer: Language models can teach themselves to use tools, Adv. Neural Inf. Process. Syst. 36 (2023) 68539–68551.

[97] D. Driess, F. Xia, M.S. Sajjadi, C. Lynch, A. Chowdhery, B. Ichter, A. Wahid, J. Tompson, Q. Vuong, T. Yu, et al., PaLM-E: An embodied multimodal language model, in: Proc. Mach. Learn. Res., Vol. 202, 2023, pp. 8469–8488.

[98] F. Bacon, Novum Organum, Clarendon Press, Oxford, United Kingdom, 1878.

[99] Z. Zhou, C. Huang, K. Lin, Y. Wang, C.K. Law, B. Yang, High-efficiency design of combustion kinetic experiments based on genetic algorithm, Proc. Combust. Inst. 39 (4) (2023) 5219–5228.

[100] P. Orzechowski, W. La Cava, J.H. Moore, Where are we now? A large benchmark study of recent symbolic regression methods, in: Proc. Genet. Evol. Comput. Conf. Companion, 2018, pp. 1183–1190.

[101] S.L. Brunton, J.L. Proctor, J.N. Kutz, Discovering governing equations from data by sparse identification of nonlinear dynamical systems, Proc. Natl. Acad. Sci. USA 113 (15) (2016) 3932–3937.

[102] M. Cranmer, Interpretable machine learning for science with PySR and SymbolicRegression.jl, 2023, arXiv preprint arXiv:2305.01582.

[103] R.S. Freitas, A. Péquin, R.M. Galassi, A. Attili, A. Parente, Model identification in reactor-based combustion closures using sparse symbolic regression, Combust. Flame 255 (2023) 112925.

[104] W.T. Chung, A.A. Mishra, M. Ihme, Interpretable data-driven methods for subgrid-scale closure in LES for transcritical LOX/GCH_4 combustion, Combust. Flame 239 (2022) 111758.

[105] K. Zdybał, J.C. Sutherland, A. Parente, Manifold-informed state vector subset for reduced-order modeling, Proc. Combust. Inst. 39 (4) (2023) 5145–5154.

[106] I.T. Jolliffe, J. Cadima, Principal component analysis: A review and recent developments, Philos. Trans. R. Soc. A 374 (2065) (2016) 20150202.

[107] A. Kumar, M. Rieth, O. Owoyele, J.H. Chen, T. Echekki, Acceleration of turbulent combustion DNS via principal component transport, Combust. Flame 255 (2023) 112903.

[108] K.S. Jung, A. Kumar, T. Echekki, J.H. Chen, On the application of principal component transport for compression ignition of lean fuel/air mixtures under engine relevant conditions, Combust. Flame 260 (2024) 113204.

[109] A. Jonnalagadda, S. Kulkarni, A. Rodhiya, H. Kolla, K. Aditya, A co-kurtosis based dimensionality reduction method for combustion datasets, Combust. Flame 250 (2023) 112635.

[110] S. Abdelwahid, M.R. Malik, H.A. Al Kader Hammoud, F.E. Hernández-Pérez, B. Ghanem, H.G. Im, Large eddy simulations of ammonia-hydrogen jet flames at elevated pressure using principal component analysis and deep neural networks, Combust. Flame 253 (2023) 112781.

[111] B.A. Perry, M.T. Henry de Frahan, S. Yellapantula, Co-optimized machine-learned manifold models for large eddy simulation of turbulent combustion, Combust. Flame 244 (2022) 112286.

[112] G.E. Hinton, R.R. Salakhutdinov, Reducing the dimensionality of data with neural networks, Science 313 (5786) (2006) 504–507.

[113] Z. Han, X. Tang, M.M. Hossain, C. Xu, Assessment of flame stability through a convolutional denoising autoencoder and statistical analysis, Combust. Flame 258 (2023) 113069.

[114] K. Iemura, M. Saito, Y. Suganuma, M. Kikuchi, Y. Inatomi, H. Nomura, M. Tanabe, Analysis of spatial-temporal dynamics of cool flame oscillation phenomenon occurred around a fuel droplet array by using variational auto-encoder, Proc. Combust. Inst. 39 (2) (2023) 2523–2532.

[115] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, J. Mach. Learn. Res. 15 (56) (2014) 1929–1958.

[116] M. Ullman, S. Barwey, G.S. Lee, V. Raman, Segmentation of high-speed flow fields using physics-informed clustering, Appl. Energy Combust. Sci. 15 (2023) 100181.

[117] H. Dave, N. Swaminathan, A. Parente, Interpretation and characterization of MILD combustion data using unsupervised clustering informed by physics-based, domain expertise, Combust. Flame 240 (2022) 111954.

[118] Y. Wang, C. Tao, Z. Zhou, K. Lin, C.K. Law, B. Yang, Clustering algorithm for experimental datasets using global sensitivity-based affinity propagation (GSAP), Combust. Flame 259 (2024) 113121.

[119] B.J. Frey, D. Dueck, Clustering by passing messages between data points, Science 315 (5814) (2007) 972–976.

[120] M.R. Malik, R. Khamedov, F.E. Hernández Pérez, A. Coussement, A. Parente, H.G. Im, Dimensionality reduction and unsupervised classification for high-fidelity reacting flow simulations, Proc. Combust. Inst. 39 (4) (2023) 5155–5163.

[121] G. D'Alessio, S. Sundaresan, M.E. Mueller, Automated and efficient local adaptive regression for principal component-based reduced-order modeling of turbulent reacting flows, Proc. Combust. Inst. 39 (4) (2023) 5249–5258.

[122] K. Zdybał, G. D'Alessio, A. Attili, A. Coussement, J.C. Sutherland, A. Parente, Local manifold learning and its link to domain-based physics knowledge, Appl. Energy Combust. Sci. 14 (2023) 100131.

[123] W.J. Murdoch, C. Singh, K. Kumbier, R. Abbasi-Asl, B. Yu, Definitions, methods, and applications in interpretable machine learning, Proc. Natl. Acad. Sci. USA 116 (44) (2019) 22071–22080.

[124] J.C. Saldinger, P. Elvati, A. Violi, A machine learning framework to predict the aggregation of polycyclic aromatic compounds, Proc. Combust. Inst. 39 (1) (2023) 1009–1017.

[125] T. Chen, C. Guestrin, XGBoost: A scalable tree boosting system, in: Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Min., Vol. 22, 2016, pp. 785–794.

[126] L. Grinsztajn, E. Oyallon, G. Varoquaux, Why do tree-based models still outperform deep learning on typical tabular data? Adv. Neural Inf. Process. Syst. 35 (2022) 507–520.

[127] J. Yu, S. Ruan, H. Song, L. Zhang, M. Yang, Machine learning rate constants of hydrogen abstraction reactions between ester and H atom, Combust. Flame 255 (2023) 112901.

[128] S. Thrun, Lifelong learning algorithms, in: Learning To Learn, Springer, 1998, pp. 181–209.

[129] Y. Zhang, W. Dong, L.A. Vandewalle, R. Xu, G.P. Smith, H. Wang, Neural network approach to response surface development for reaction model optimization and uncertainty minimization, Combust. Flame 251 (2023) 112679.

[130] J.A.C. Kildare, W.T. Chung, M.J. Evans, Z. Tian, P.R. Medwell, M. Ihme, Predictions of instantaneous temperature fields in jet-in-hot-coflow flames using a multi-scale U-Net model, Proc. Combust. Inst. 40 (2025) 105330.

[131] X. Su, W. Ji, J. An, Z. Ren, S. Deng, C.K. Law, Kinetics parameter optimization of hydrocarbon fuels via neural ordinary differential equations, Combust. Flame 251 (2023) 112732.

[132] W. Ji, F. Richter, M.J. Gollner, S. Deng, Autonomous kinetic modeling of biomass pyrolysis using chemical reaction neural networks, Combust. Flame 240 (2022) 111992.

[133] Y. Kim, J. Cho, N. Naser, S. Kumar, K. Jeong, R.L. McCormick, P.C. St. John, S. Kim, Physics-informed graph neural networks for predicting cetane number with systematic data quality analysis, Proc. Combust. Inst. 39 (4) (2023) 4969–4978.

[134] T. Yoon, S.W. Kim, H. Byun, Y. Kim, C.D. Carter, H. Do, Deep learning-based denoising for fast time-resolved flame emission spectroscopy in high-pressure combustion environment, Combust. Flame 248 (2023) 112583.

[135] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, Proc. IEEE 86 (11) (1998) 2278–2324.

[136] M. Gharib, P. Tischer, O. Schulze, M. Gräbner, A. Richter, Flame lift-off detector based on deep learning neural networks, Combust. Flame 260 (2024) 113215.

[137] M. Dai, B. Zhou, J. Zhang, R. Cheng, Q. Liu, R. Zhao, B. Wang, B. Gao, 3-D soot temperature and volume fraction reconstruction of afterburner flame via deep learning algorithms, Combust. Flame 252 (2023) 112743.

[138] S. Hochreiter, J. Schmidhuber, Long short-term memory, Neural Comput. 9 (8) (1997) 1735–1780.

[139] X. Cheng, F. Ren, Z. Gao, L. Wang, L. Zhu, Z. Huang, Predicting 3D distribution of soot particle from luminosity of turbulent flame based on conditional-generative adversarial networks, Combust. Flame 247 (2023) 112489.

[140] W.T. Chung, C. Laurent, D. Passiatore, M. Ihme, Ensemble predictions of laser ignition with a hybrid stochastic physics-embedded deep-learning framework, Proc. Combust. Inst. 40 (2024) 105304.

[141] R. Strelau, M. Frederick, W.C. Senior, R. Gejji, C.D. Slabaugh, Modes of laser spark ignition of a model rocket combustor, in: AIAA Paper 2023-2377, 2023.

[142] L. Nista, C. Schumann, T. Grenga, A. Attili, H. Pitsch, Investigation of the generalization capability of a generative adversarial network for large eddy simulation of turbulent premixed reacting flows, Proc. Combust. Inst. 39 (4) (2023) 5279–5288.

[143] P. Sharma, W.T. Chung, B. Akoush, M. Ihme, A review of physics-informed machine learning in fluid mechanics, Energies 16 (5) (2023) 2343.

[144] W.T. Chung, M. Ihme, K.S. Jung, J.H. Chen, J. Guo, D. Brouzet, M. Talei, B. Jiang, B. Savard, et al., BLASTNet simulation dataset version 2.1, 2024, https://blastnet.github.io/.

[145] G.E. Karniadakis, I.G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, Physics-informed machine learning, Nat. Rev. Phys. 3 (2021) 422–440.

[146] S. Yao, A. Kronenburg, O. Stein, Efficient modeling of the filtered density function in turbulent sprays using ensemble learning, Combust. Flame 237 (2022) 111722.

[147] M. Ihme, C. Schmitt, H. Pitsch, Optimal artificial neural networks and tabulation methods for chemistry representation in LES of a bluff-body swirl-stabilized flame, Proc. Combust. Inst. 32 (1) (2009) 1527–1535.

[148] T. Readshaw, L.L. Franke, W. Jones, S. Rigopoulos, Simulation of turbulent premixed flames with machine learning - tabulated thermochemistry, Combust. Flame 258 (2023) 113058.

[149] Z. Nikolaou, L. Vervisch, P. Domingo, Criteria to switch from tabulation to neural networks in computational combustion, Combust. Flame 246 (2022) 112425.

[150] R. Nakazawa, Y. Minamoto, N. Inoue, M. Tanahashi, Species reaction rate modelling based on physics-guided machine learning, Combust. Flame 235 (2022) 111696.

[151] J.Z. Ho, M. Talei, D. Brouzet, W.T. Chung, P. Sharma, M. Ihme, Augmenting filtered flame front displacement models for LES using machine learning with *a posteriori* simulations, Proc. Combust. Inst. 40 (2024) 105311.

[152] V. Vijayarangan, H.A. Uranakara, S. Barwey, R. Malpica Galassi, M.R. Malik, M. Valorani, V. Raman, H.G. Im, A data-driven reduced-order model for stiff chemical kinetics using dynamics-informed training, Energy AI 15 (2024) 100325.

[153] W.T. Chung, A.A. Mishra, N. Perakis, M. Ihme, Data-assisted combustion simulations with dynamic submodel assignment using random forests, Combust. Flame 227 (2021) 172–185.

[154] R. Malpica Galassi, P.P. Ciottoli, M. Valorani, H.G. Im, An adaptive time-integration scheme for stiff chemistry based on computational singular perturbation and artificial neural networks, J. Comput. Phys. 451 (2022) 110875.

[155] C. Olah, N. Cammarata, L. Schubert, G. Goh, M. Petrov, S. Carter, Zoom in: An introduction to circuits, Distill 5 (3) (2020) e00024–001.

[156] F. Charlette, C. Meneveau, D. Veynante, A power-law flame wrinkling model for LES of premixed turbulent combustion Part I: Non-dynamic formulation and initial tests, Combust. Flame 131 (2002) 159–180.

[157] G. Dalal, K. Dvijotham, M. Vecerik, T. Hester, C. Paduraru, Y. Tassa, Safe exploration in continuous action spaces, 2018, arXiv preprint arXiv:1801.08757.
[158] R. Balestriero, M. Ibrahim, V. Sobal, A. Morcos, S. Shekhar, T. Goldstein, F. Bordes, A. Bardes, G. Mialon, Y. Tian, et al., A cookbook of self-supervised learning, 2023, arXiv preprint arXiv:2304.12210.
[159] T. Brown, B. Mann, N. Ryder, M. Subbiah, J.D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, Adv. Neural Inf. Process. Syst. 33 (2020) 1877–1901.
[160] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, et al., Training language models to follow instructions with human feedback, Adv. Neural Inf. Process. Syst. 35 (2022) 27730–27744.
[161] M. Chen, J. Tworek, H. Jun, Q. Yuan, H.P.d. Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman, et al., Evaluating large language models trained on code, 2021, arXiv preprint arXiv:2107.03374.
[162] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, et al., Retrieval-augmented generation for knowledge-intensive NLP tasks, Adv. Neural Inf. Process. Syst. 33 (2020) 9459–9474.
[163] I. Sutskever, O. Vinyals, Q.V. Le, Sequence to sequence learning with neural networks, Adv. Neural Inf. Process. Syst. 27 (2014).
[164] X. Wang, K. Yu, S. Wu, J. Gu, Y. Liu, C. Dong, Y. Qiao, C. Change Loy, ESRGAN: Enhanced super-resolution generative adversarial networks, in: Proc. Eur. Conf. Comput. Vis. Workshop, 2018, pp. 63–79.
[165] R. Lam, A. Sanchez-Gonzalez, M. Willson, P. Wirnsberger, M. Fortunato, F. Alet, S. Ravuri, T. Ewalds, Z. Eaton-Rosen, W. Hu, et al., Learning skillful medium-range global weather forecasting, Science 382 (6677) (2023) 1416–1421.
[166] Common Crawl, Open repository of web crawl data, 2007, https://commoncrawl.org/.
[167] Hugging Face, The AI community building the future, 2016, https://huggingface.co/.
[168] H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, et al., The ERA5 global reanalysis, Q. J. R. Meteorol. Soc. 146 (730) (2020) 1999–2049.
[169] K. Bi, L. Xie, H. Zhang, X. Chen, X. Gu, Q. Tian, Accurate medium-range global weather forecasting with 3D neural networks, Nature 619 (7970) (2023) 533–538.
[170] T. Kurth, S. Subramanian, P. Harrington, J. Pathak, M. Mardani, D. Hall, A. Miele, K. Kashinath, A. Anandkumar, FourCastNet: Accelerating global high-resolution weather forecasting using adaptive Fourier neural operators, in: Proc. Platf. Adv. Sci. Comput. Conf., 2023.
[171] S. Rasp, S. Hoyer, A. Merose, I. Langmore, P. Battaglia, T. Russel, A. Sanchez-Gonzalez, V. Yang, R. Carver, S. Agrawal, et al., WeatherBench 2: A benchmark for the next generation of data-driven global weather models, 2023, arXiv preprint arXiv:2308.15560.
[172] W.T. Chung, K.S. Jung, J.H. Chen, M. Ihme, BLASTNet: A call for community-involved big data in combustion machine learning, Appl. Energy Combust. Sci. 12 (2022) 100087.
[173] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: Y. Bengio, Y. LeCun (Eds.), Proc. Int. Conf. Learn. Represent., Vol. 3, 2015.
[174] E.J. Hu, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, et al., LoRA: Low-rank adaptation of large language models, in: Proc. Int. Conf. Learn. Represent., Vol. 9, 2021.
[175] T. Dettmers, A. Pagnoni, A. Holtzman, L. Zettlemoyer, QLoRA: Efficient finetuning of quantized LLMs, Adv. Neural Inf. Process. Syst. 36 (2023) 10088–10115.
[176] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, et al., LLaMa: Open and efficient foundation language models, 2023, arXiv preprint arXiv:2302.13971.
[177] R. Anil, A.M. Dai, O. Firat, M. Johnson, D. Lepikhin, A. Passos, S. Shakeri, E. Taropa, P. Bailey, Z. Chen, et al., PaLM 2 technical report, 2023, arXiv preprint arXiv:2305.10403.
[178] B. Savard, E.R. Hawkes, K. Aditya, H. Wang, J.H. Chen, Regimes of premixed turbulent spontaneous ignition and deflagration under gas-turbine reheat combustion conditions, Combust. Flame 208 (2019) 402–419.
[179] B. Lim, S. Son, H. Kim, S. Nah, K. Mu Lee, Enhanced deep residual networks for single image super-resolution, in: Proc. IEEE Conf. Comput. Vision Pattern Recogn. Workshop, 2017, pp. 136–144.
[180] Z. Ji, N. Lee, R. Frieske, T. Yu, D. Su, Y. Xu, E. Ishii, Y.J. Bang, A. Madotto, P. Fung, Survey of hallucination in natural language generation, ACM Comput. Surv. 55 (12) (2023) 1–38.
[181] W. Blau, V.G. Cerf, J. Enriquez, J.S. Francisco, U. Gasser, M.L. Gray, M. Greaves, B.J. Grosz, K.H. Jamieson, G.H. Haug, et al., Protecting scientific integrity in an age of generative AI, Proc. Natl. Acad. Sci. USA 121 (22) (2024) e2407886121.
[182] P. Sharma, W.T. Chung, M. Ihme, A physics-informed machine learning approach for predicting dynamic behavior of reacting flows with application to combustion, 2024, Under review.
[183] K. Stachenfeld, D.B. Fielding, D. Kochkov, et al., Learned simulators for turbulence, in: Proc. Int. Conf. Learn. Represent., Vol. 10, 2022.
[184] V. Sharma, V. Raman, A reliable knowledge processing framework for combustion science using foundation models, Energy AI 16 (2024) 100365.
[185] D. Hendrycks, C. Burns, S. Kadavath, A. Arora, S. Basart, E. Tang, D. Song, J. Steinhardt, Measuring mathematical problem solving with the MATH dataset, in: Proc. Neur. Inf. Proces. Syst. Datasets Benchmarks, Vol. 1, 2021.
[186] A. Stagni, A. Frassoldati, A. Cuoci, T. Faravelli, E. Ranzi, Skeletal mechanism reduction through species-targeted sensitivity analysis, Combust. Flame 163 (2016) 382–393.
[187] H. Chase, LangChain, 2022, https://github.com/langchain-ai/langchain.
[188] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q.V. Le, D. Zhou, et al., Chain-of-thought prompting elicits reasoning in large language models, Adv. Neural Inf. Process. Syst. 35 (2022) 24824–24837.
[189] M. Reid, N. Savinov, D. Teplyashin, D. Lepikhin, T. Lillicrap, J.-b. Alayrac, R. Soricut, A. Lazaridou, O. Firat, J. Schrittwieser, et al., Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context, 2024, arXiv preprint arXiv:2403.05530.
[190] A.Q. Jiang, A. Sablayrolles, A. Roux, A. Mensch, B. Savary, C. Bamford, D.S. Chaplot, D.d.l. Casas, E.B. Hanna, F. Bressand, et al., Mixtral of experts, 2024, arXiv preprint arXiv:2401.04088.
[191] D. Dua, Y. Wang, P. Dasigi, G. Stanovsky, S. Singh, M. Gardner, DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs, in: Proc. N. Am. Assoc. Comput. Linguist., Vol. 1, 2019, pp. 2368–2378.
[192] S.W. Benson, Thermochemical Kinetics: Methods for the Estimation of Thermochemical Data and Rate Parameters, second ed., Wiley, New York, 1976.
[193] Y. Li, S. Bubeck, R. Eldan, A. Del Giorno, S. Gunasekar, Y.T. Lee, Textbooks are all you need II: Phi-1.5 technical report, 2023, arXiv preprint arXiv:2309.05463.
[194] B. Wallace, B. Hariharan, Extending and analyzing self-supervised learning across domains, in: Proc. Eur. Conf. Comput. Vis., Vol. 16, 2020, pp. 717–734.
[195] A. Tamkin, G. Banerjee, M. Owda, V. Liu, S. Rammoorthy, N. Goodman, DABS 2.0: Improved datasets and algorithms for universal self-supervision, Adv. Neural Inf. Process. Syst. 35 (2022) 38358–38372.
[196] T. Saier, M. Färber, unarXive: A large scholarly data set with publications' full-text, annotated in-text citations, and links to metadata, Scientometrics 125 (3) (2020) 3085–3108.
[197] H. Jasak, OpenFOAM: Open source CFD in research and industry, Int. J. Nav. Archit. Ocean Eng. 1 (2) (2009) 89–94.
[198] P. Liang, R. Bommasani, T. Lee, D. Tsipras, D. Soylu, M. Yasunaga, Y. Zhang, D. Narayanan, Y. Wu, A. Kumar, et al., Holistic evaluation of language models, Trans. Mach. Learn. Res. (2023).
[199] A. Srivastava, A. Rastogi, A. Rao, A.A.M. Shoeb, A. Abid, A. Fisch, A.R. Brown, A. Santoro, A. Gupta, A. Garriga-Alonso, et al., Beyond the imitation game: Quantifying and extrapolating the capabilities of language models, Trans. Mach. Learn. Res. (2023).
[200] J. Scully, C.J. Herington, Prometheus Bound, Oxford University Press, Oxford, United Kingdom, 1989.