

Machine Learning Possibilities for Nucleosynthesis Studies

Michael Smith^{1,*}, Dan Lu²

^{1,*} *Physics Division, Oak Ridge National Laboratory, Oak Ridge, TN, USA*

² *Computational Sciences and Engineering Division, Oak Ridge National Laboratory, Oak Ridge, TN, USA*

Correspondence*:

Michael Smith, Physics Division, MS-6354, Bldg. 3500, Oak Ridge National Laboratory, Oak Ridge, TN, 37831-6354, USA
smithms@ornl.gov

ABSTRACT

Nuclear astrophysics is an interdisciplinary field focused on exploring the impact of nuclear physics on the evolution and explosions of stars and the cosmic creation of the elements. While researchers in astrophysics and in nuclear physics are separately using machine learning approaches to advance studies in their fields, there is currently little use of machine learning in nuclear astrophysics. We briefly describe the most common categories of machine learning algorithms, and then detail their numerous possible uses to advance nuclear astrophysics, with a focus on simulation-based nucleosynthesis studies. We show that machine learning offers novel, complementary, creative approaches to address many important nucleosynthesis puzzles, with the potential to initiate a new frontier in nuclear astrophysics research.

Keywords: nuclear astrophysics, nucleosynthesis, simulations, machine learning, neural nets

1 INTRODUCTION

Machine learning (ML) is now broadly used in the field of astrophysics to study a wide variety of phenomena. A selection of recent examples includes projects in observational cosmology [1], galactic evolution [2], cosmic ray measurement interpretation [3], star formation with dark matter [4], galaxy spectral energy distributions [5], blazer observations [6], Ba star abundances [7], and metallicity-dependent abundances [8]. While not as widely used in nuclear physics, ML utilization there is growing, as reviewed in Ref. [9]. Some more recent examples in nuclear physics include the use of ML for studies at low- and medium-energies [10], for neutron-induced reaction cross section evaluations [11], for heavy-ion fusion cross sections [12], for nuclear mass predictions [13, 14, 15, 16, 17], for active-target time projection chamber data analysis [18], for analysis of time-of-flight data [19], and for reaction cross section predictions [20].

From these examples, it would seem likely that ML would also be widely utilized in the interdisciplinary research that spans these two fields – that is, in *nuclear astrophysics* – but this is not the case. Researchers in nuclear astrophysics [21, 22] explore the critical leverage that physics at the femtometer scale has on stellar systems that are 10^{24} times larger, as well as on the origin and evolution of the Universe. At the core of many of these studies are simulations that combine the results nuclear measurements, nuclear theory, and astrophysical observations to expand our understanding of the cosmic creation of the elements, and to

decipher why systems like our sun, red giants, white dwarfs, variable stars, and others evolve peacefully while other stars end their lives in nova, supernova, hypernova, or even more exotic explosions. In this era of multi-messenger astronomy [23], improved predictions of such simulations are needed more than ever to connect measurements at new and planned accelerator facilities [24, 25, 26, 27] and advances in nuclear theory driven by exascale supercomputers [28] with observations from the latest generation of ground- and space-based telescopes [29].

While astrophysics and nuclear physics have separately used ML approaches to advance their frontiers, there is very little use of ML approaches in nuclear astrophysics. This is evident from recent journal articles, from presentations at the largest recent international symposia (*e.g.*, [30, 31]), and from the latest strategic planning documents for nuclear physics [32]; there are two exceptions (at the time of writing this article), Refs. [33, 34], which will be discussed below in Section 3. To address this paucity of ML studies, this article explores the possibilities that ML approaches have to advance the field of nuclear astrophysics, with a focus on simulation-based studies of the cosmic synthesis of the elements [35, 36]. We show that ML offers novel, complementary, creative approaches to address many important nucleosynthesis puzzles. Since traditional approaches in this work have not changed in decades, ML has the potential to initiate a new frontier in nuclear astrophysics research.

To begin, we first give a brief description of a variety of widely-utilized ML algorithm types in Section 2. Some important nucleosynthesis puzzles are then described in Section 3, along with suggestions of methods to employ ML algorithms that may advance our knowledge. We then briefly discuss some of the challenges of ML approaches in Section 4, and then give a summary in Section 5.

2 WIDELY-UTILIZED MACHINE LEARNING APPROACHES

ML approaches have been shown to be very effective in addressing data-centric problems in a wide variety of fields. Common uses for ML algorithms are to classify data, make decisions, predict values, identify outliers or anomalies, find patterns, interpret large datasets, quantify uncertainties, efficiently map inputs to outputs, reduce dimensionalities, and find hidden functional relationships. Below we give a brief description of widely utilized ML algorithm *categories* that are routinely used for these (and other) tasks. These categories will then be discussed in Section 3 as possible approaches to address important challenges in nucleosynthesis research.

Linear Regression (LinR) [37] – fitting a data set to a (single- or multivariate) linear function by least squares minimization is well understood and widely utilized for data sets with a linear input-output mapping. ML use emphasizes accurate predictions; usage in statistics emphasizes the correctness of the linear model.

Kernel Ridge Regression (KRR) [38] – produces fits (predictions) over multiple variables that have high correlations (multicollinearities) that can cause problems in standard regressions; *KRRs* map the original data into a more complex kernel-defined feature space, but do not generate prediction uncertainties.

Logistic Regression (LogR) [39] – this widely used classification algorithm maps continuous quantities to (usually two) discrete quantities (*e.g.*, “Yes/No”, “On/Off”) by fitting a sigmoid (logit) function to the data. This approach is easily understood (explainable) and works with non-linear data sets.

Classification Tree (CT) [40] – these employ a set of cascading rule-based tests with a tree-like structure to sort (classify) labelled data into categories. A variation – Classification and Regression Trees (*CARTs*) – can be used for regression by partitioning data into groups with similar values of a dependent variable.

Decision Tree (DT) [41] – structurally similar to *CTs*, these trees make decisions (*i.e.*, give answers) based on rule-based tests (*i.e.*, questions). *DTs* can function as “expert systems” that drill down to a recommendation based on multivariate input.

Random Forest (RF) [42] – instead of using rule-based data tests, *RFs* randomly generate many *DTs* that each “vote” on a classification, in order to overcome limitations of single *DTs* and add features like weighting and error estimation. *RFs* can also be used for regression when testing on a continuous variable.

Gradient Boost (GB) [43] – in contrast to *RFs* which combine results of different *DTs* as a final step, a *GB* combines *DT* results *in series* to make successively stronger (better predictive) models. *GBs* have great flexibility for tuning and loss functions, and can give highly accurate predictions.

Isolation Forest (IF) [44] – since anomalies (outliers) in data tend to be “few and different” from other data, a tree structure can be used to separate them after only a few tests (*i.e.*, near the “root” of the tree). *IFs* are characterized by fast execution and high performance for a wide range of anomalies.

Support Vector Machine (SVM) [45] – for data with n features, an n -dimensional vector is drawn with values on each coordinate corresponding to its feature value; classifications are made via relative data point distances in the n -space. *SVM* algorithms are robust against outliers and computationally efficient.

K Means Clustering (KM) [46] – unlabeled data points are grouped via distances in feature space to k randomly-assigned centroids; data is classified by iteratively reassigning centroids until obtaining optimally separated clusters. Care is needed with this approach to avoid local minima.

K Nearest Neighbor (KNN) [47] – used to classify (group) data points based on commonalities with the majority of its k nearest neighbors in feature space. *KNNs* can also be used for regression where commonalities are treated as numerical distances.

Principal Component Analysis (PCA) [48] – reduce the dimensionality of large multivariate data sets by finding fewer (sometimes new) parameters to “represent” the data collection with minimal information loss. *PCA* works well with highly correlated data sets that have many parameters.

Naïve Bayes (NB) [49] – uses information in the data to estimate Bayesian posterior probabilities with the (naïve) assumption that attributes are conditionally independent. Classifications and decisions are made by setting thresholds on probabilities.

Variational Autoencoder (VAE) [50] – uses a neural net to compress/encode data as parameters of a distribution over random variables in a continuous lower-dimensional latent space, then generatively reconstruct/decode the data. This reduces noise, adds probabilities, and focuses on critical data features.

Symbolic Regression (SR) [51] – used to optimally specify a mathematical formula to input-to-output data mapping, by altering both the structure and parameters of an analytical model. This produces a model that can be more easily explained (*i.e.*, interpreted).

Gaussian Processes (GP) [52] – use a collection of normally-distributed random variables to specify distributions over complex functions without knowing the exact form. *GPs* are useful for accurate predictions (regression) with uncertainties, as well as for classifications.

Neural Network (NN) [53] – layers of interconnected nodes (neurons), each with its own weighting, bias, and activation function, serve to process inputs to outputs; backpropagation is used to iteratively adjust weights by comparing outputs to training data, after which predictions can be made from new inputs. Widely used to model complex functions, deep *NNs* (*DNNs*) with many hidden layers are the basis of “deep learning” (*DL*).

Convolutional Neural Net (CNN) [54] – these NNs extract features by “sliding” (*i.e.*, convolving) a set of filters (kernels) over data that has a grid-like structure. Filter outputs are subsequently collected and combined by fully connected NN layers for classification. CNNs are widely used for image analyses.

Bayesian Neural Net (BNN) [55] – these NNs have stochastic weights to simulate, using a Bayesian inference framework, the predictions of multiple possible models and the probability distribution associated with each. In this way, BNNs are used to quantify the prediction uncertainties of NN-based models.

Recurrent Neural Net (RNN) [56] – by structuring a NN with repeating layers (loops) that link in forward and reverse (recurrent) directions, information can be stored (a “memory”) as inputs of arbitrary length are sequentially processed. RNNs are widely used to analyze time series data, speech, music, and text.

Graph Neural Net (GNN) [57] – graph-like data (*e.g.*, entities plus their relationships, which can be images, texts, molecular structures, and more) are analyzed by an optimizable transformation of all graph attributes to find missing elements or relationships, identify/characterize subcomponents, or other prediction tasks.

Radial Basis Function Neural Net (RBFNN) [58] – useful for regression of non-linear functions or associated classification problems, these 3-layer NNs have a unique hidden layer with neuron weights determined by the (kernel-calculated) distance from a central point; they train quickly but can be difficult to set up.

Emulators [59] – fast-executing ML models trained to approximately reproduce the results of (*i.e.*, produce the same input - output mapping as) complex simulations. They facilitate exploring uncertainties, sensitivities, parameter spaces, and more. NNs, DNNs, CNNs, GPs, and RFs are often used as emulators.

Generative Adversarial Network (GAN) [60] – these use a CNN to generate (create) new data resembling training data, and an opposing NN to discriminate between training and new data. GANs can generate images from text, reword text with new keywords, create training data for algorithms, and much more.

Large Language Model (LLM) [61] – trained on up to $\sim 10^{12}$ tokens using $\sim 10^{12}$ parameters, these models include recurrent, generative, and other NN layers to recognize, predict, and generate text. LLMs use statistical approaches in contrast to traditional rule-based approaches of natural language processing.

Transformers [62] – a widely used, highly scalable DNN architecture that revolutionized LLMs by using the concept of *attention* to comprehend contextual relationships within text and sequential data. It excels at summarizing and translating text, answering questions, analyzing sequential data, and much more.

Generative Pre-trained Transformer (GPT) [63] – adding generative capabilities to a Transformer foundation trained on internet-scale data, this is a popular foundation for LLMs. ChatGPT is a chat interface app for GPT queries (prompts). Fine-tuned GPTs are specialized to work with text, images, music, video, and more.

As described in Section 3 below, nucleosynthesis studies may possibly be advanced via use of algorithms in the above categories. Our algorithm category summary may also inspire the use of ML approaches in other research fields.

3 MACHINE LEARNING FOR NUCLEOSYNTHESIS STUDIES

In any given astrophysical system, the complex mechanisms responsible for the cosmic synthesis of nuclei involve over 10^{50} nuclei spread over $\sim 10^3$ nuclear species (isotopes) together with sequences of up to $\sim 10^4$ interconnecting thermonuclear reactions that transmute nuclides of one species to another. To simulate nucleosynthesis in such a system, the specified initial abundances are numerically evolved, in short time steps, into their final abundances over an appropriate set of hydrodynamic conditions [35, 64].

The compositional changes predicted by simulations are often given as one-dimensional ($1D$) plots of abundance histories (Fig. 1 [65, 66, 67, 68]). Because there are too many histories to show at once on these plots, $2D$ visualizations (Fig. 2) are often used which display abundance values (mapped to colors) of each isotope at its location set by neutron number N and proton number Z ; such an N - Z plane is referred to as the “nuclide chart”. At any time step in the simulation, such an image shows the cumulative effect of all reactions that have both created and destroyed (“burned”) each isotope up to that point in the simulation. A complementary $2D$ nuclide chart visualization (Fig.3) represents the amount of material changing from one isotope to another via individual reactions (the reaction “flux” [69]) as the width/pattern/color of arrows connecting the isotopes. Such “flow diagrams” are especially appropriate given that nucleosynthesis is often conceptualized as the flow of stellar material (abundances) from lighter- to heavier-mass nuclides.

While the overall approach of most simulation-based nucleosynthesis studies are similar, there are differences depending on the investigation goals, available computational power, and acceptable approximations. For studies exploring the importance thermonuclear reactions on simulation predictions, a full treatment of thermonuclear burning is often used wherein the abundances of all relevant isotopes are solved numerically considering all interconnecting reactions. To speed execution, many studies employ a “post-processing” simulation approach where the full thermonuclear burning problem is computed over predetermined temperature and density vs. time trajectories (*e.g.*, Refs. [68, 70]). Additional execution speed is obtained by assuming spherical symmetry in the hydrodynamic trajectories, so calculations are made in one dimension ($1D$) along the system radius.

Some studies more realistically couple a full thermonuclear burn treatment to a $1D$ hydrodynamics code (*e.g.*, Refs. [71, 72]). This coupling, missing in post-processing studies, is critical because it produces self-consistent solutions. Further model enhancements require more complex hydrodynamics. For example, mixing length theory [73] is often used to approximate the complex effects of convection in $1D$ studies, but more realistic convection treatments require $2D$ or $3D$ hydrodynamics codes [74, 75]. For certain effects like standing accretion shock instabilities (*SASI*) [76] and stellar rotation, $3D$ hydrodynamic approaches are required. Because of the extreme computational demands of the $2D$ and $3D$ codes [77], however, these simulations often employ a very truncated (approximate) treatment of thermonuclear burning that includes only the reactions and isotopes that most influence the hydrodynamics. To more accurately time-evolve the full isotopic inventory in such approaches, hydrodynamic trajectories are then extracted from the simulation for use in separate post-processing nucleosynthesis studies. This extraction is often done with a “tracer particle” approach, discussed further below in Section 3.1. There are, however, known issues with the use of truncated thermonuclear burning treatments, including problems with energy generation, neutrino heating, and nucleosynthesis in core-collapse supernovae [78]. A major goal in nucleosynthesis studies is therefore to develop simulations that couple full thermonuclear burning with multi-dimensional hydrodynamics.

This goal has not yet been practically realized, however, due to the daunting computational requirements. For this reason, and because every relevant thermonuclear reaction and every set of hydrodynamic conditions cannot be investigated, it is critical to carry out studies that guide researchers on where to focus their efforts to make the most progress in understanding nucleosynthesis. Some such efforts, for example, strive to understand details of nucleosynthesis flows – a challenge as these flows reflect the interplay of the hydrodynamic conditions and the underlying relevant nuclear physics. Other efforts focus on devising approximations to nucleosynthesis flows – desirable to reveal underlying structures or symmetries, as well as to speed simulation execution and thereby accelerate scientific discoveries. There are also studies focusing on quantifying the uncertainties of model predictions to enable robust comparisons

of predictions to observations. Sensitivity Analyses are another important approach, where changes in model predictions (outputs) caused by systematic variations of inputs are examined; this flags inputs that significantly impact critical simulation outputs for further investigation. And finally, there are efforts to improve models, such as by identifying and correcting anomalous inputs.

The following subsections give some details on possible studies in each of these areas mentioned above – nucleosynthesis flows, complexity reduction, uncertainty quantification, sensitivity analyses, and improving models – using ML algorithms in the categories mentioned in Section 2. Some of these are novel and creative approaches that hold the promise of enabling significant progress for simulation-based nucleosynthesis studies, where the overall solution scheme has not changed in decades. A brief discussion of the utility of ML (especially *LLMs*) for speeding scientific workflows is also given. Furthermore, additional advances in the field may also be achieved by combining the approaches discussed below in innovative ways.

3.1 Nucleosynthesis Flows

Flow Patterns – Identifying and analyzing patterns in complex nucleosynthesis flows may provide insights, decouple overlapping (*e.g.*, thermonuclear and hydrodynamic) effects, and pinpoint critical nuclides and reactions for future study. For example, many experimental efforts have been driven [79] by the identification of (nearly) identical flow patterns over different portions of the nuclide chart (*e.g.*, the *Hot CNO*, *NeNa*, *MgAl*, and *SiP* cycles (Fig.4)) in simulations of nucleosynthesis in nova explosions and X-ray bursts; see, for example, Refs. [69, 79]. Repetitive flow patterns may also arise from the use of thermonuclear rates derived from statistical reaction models, since these models generate very similar reaction cross sections for target nuclides separated by an alpha particle (*i.e.*, two units in each of N and Z) on the nuclide chart. Other interesting flow patterns include: interruptions in the flow from lower to higher masses at particular nuclides (“waiting point nuclei”) [80]; a many-to-one reduction in the number of possible flow paths to higher masses (“bottleneck reactions”) [80]; a rapid cessation of flow to higher masses, indicating expansive cooling of the system or exhaustion of the thermonuclear “fuel” of an exploding star [81]; and a balancing of flows through particle captures and their inverse photodisintegrations, indicating the condition of Nuclear Statistical Equilibrium (*NSE*) [82].

The complexity of simulation predictions has, however, limited systematic studies of the flow patterns mentioned above, and has hampered searches for novel flow patterns that could provide other important signatures of aspects of thermonuclear burning or the onset of certain hydrodynamic conditions. This is where ML approaches could be very useful. For example, by treating each isotope as a graph node, and each reaction flux arrow as a directed graph edge, a *GNN* could be employed to identify and analyze patterns in flow diagrams like Fig.3. For such a “flow graph”, there are a limited number of edges directed into, and directed out of, each node, representing the possible nuclear reactions (Fig.5). Because of these limited possible edges, and because these edges are localized to reach nearest (or near) neighbors, these graphs are much less complex than many routinely analyzed by *GNNs*. For a *GNN*-based nucleosynthesis study, the edges should be weighted by the reaction flux, and the nodes should be indexed by their (N, Z) values and (if necessary) weighted by their abundance values.

A *GNN* could take a flow graph as input, generate a representation in a lower-dimensional latent (feature) space, and then identify clusters of nodes, objects and their connections, and region classifications. These could facilitate studies of the numerous flow effects mentioned above, and perhaps identify some new effects as well. *GNNs* can also generate new graph visualizations which could reveal structures and anomalies not readily apparent in the original flow diagram. By weighting nodes with, for example, accelerator

beam intensities, a *GNN* could be used to rank identified features (*e.g.*, nucleosynthesis waiting points, bottlenecks, repetitive patterns) to prioritize experiments that match facility capabilities. By weighting graph edges with reaction energy release (Q -values), a *GNN* could identify a truncated selection of isotopes and reactions that generate nearly the same thermonuclear energy as the original simulation; this could subsequently be used for accelerated nucleosynthesis simulations coupling hydrodynamics and approximate (truncated) thermonuclear burning.

In a complementary manner, ML approaches may also help identify flow patterns in abundance diagrams like Fig.2; this can be especially useful because not all simulation codes generate flow diagrams like Fig.3. Noting that abundance diagrams are coarsely “pixelated” over the nuclide chart, *CNNs* are a natural approach to search for patterns using filters that “slide” over the image. Since the pixel values (abundances) change in time, ML video analysis approaches can be used on a series of sequential pixelated abundance images, like those shown in Fig.6. For example, a *CNN* could be used to extract high-level features from individual frames that are then fed to an *RNN* that keeps a memory of the frame-to-frame temporal correlations [83]. Such a scheme could be used as a novel nucleosynthesis simulation emulator, as described below in Section 3.2.

Adjusting the color palette could also aid in analyzing these images. First, in the image creation stage, abundances are usually continuously mapped to colors, but using a coarser discrete color binning [80] as in Fig.2 could accelerate *CNN* image analyses by reducing the dimension of the color space. Additionally, such binning could facilitate the use of *KNN* or *KM* clustering algorithms to find abundance patterns via groupings of isotopes in a combined color- and (Z , N)-space. After image generation, colors can be “quantized” using a variety of techniques (*CNNs*, *RNNs*, *GANs*, and more) that preserve visual structures over a smaller color palette for faster subsequent analyses [84]. Alternatively, *KNN* and *KM* clustering algorithms could process (pre-image) numerical abundance values to determine optimal non-uniform color maps for improved *CNN* image analysis, or be used to reveal otherwise hidden patterns by bypassing image generation altogether.

Finally, since flows can be numerically approximated by the change of the abundances in Fig. 2 in time (*e.g.*, between sequential images), *abundance time derivative* images could be analyzed using *CNNs*, *CNNs* in combination with *RNNs*, or clustering algorithms to generate new insights, especially when flow diagrams are not available. Fig. 7 shows an example of the evolution of abundance time derivatives for an energetic nova explosion. As the temperature rises to the peak, the abundances at nuclides with higher (lower) Z values are increasing (decreasing) in time, with corresponding positive (negative) values for their abundance time derivatives; this is consistent with reaction flow towards higher Z (*i.e.*, upwards) on the N - Z plane. After a transition near the peak temperature, these abundance time derivatives flip in sign, consistent with the reaction flow dominated by positron decays which flow diagonally downward and to the right on the nuclide chart. New thermonuclear burning insights may be derived by examining effects including: correlations between abundance time derivative sign changes and changes in hydrodynamic conditions; rapid fluctuations in abundance time derivative signs; patterns and groupings of isotopes with similar abundance time derivative signs and sign changes; and more.

Flow Correlations – Instances where two or more isotopes have nearly identical predicted abundance vs time histories [85] (Fig. 8) may reveal underlying nucleosynthesis structures, especially for isotopes that have significant mass differences (*i.e.*, are well separated from one another on the nuclide chart). In such cases, the flow correlations could be connecting groups of isotopes in a *localized NSE*, a condition called nuclear quasi-statistical equilibrium (*QSE*) [86]. Alternatively, correlations could be due to flows through a sequence of intervening reactions that connect these distant isotopes. However, because simulations track

the abundances of so many isotopes, and because there are numerous general abundance time evolution trends (e.g., abundances increasing during rapid temperature rises), such flow correlations could merely be random. It is therefore important not only to search for such correlations, but also to determine if a causal relationship is present.

There are numerous ML approaches that could be used to search for such causal flow correlations. For example, abundance histories could be pre-processed (in this case, labeled) by selecting their values (and/or time derivatives) over a coarse time grid and then using a *KM* or *SVM* clustering algorithm or a *NN* to group similar histories together. *NNs* could also be used to facilitate Dynamic Time Warping algorithms [87] for picking out similar curves. A wider range of variations of abundance histories can likely be handled by “encoding” abundances histories with an *RNN* (or a *1D CNN*) and then using a *DNN* for feature extraction and subsequent clustering and classification.

To show causality, analyses of *2D* flow diagrams (like Fig. 3) with a *GNN* could be used to identify cases where sequences of strong flows “connect” distant isotopes with correlated abundance histories, or more generally to identify new possible causal connections. This would involve weighting graph edges by reaction flux and graph nodes by abundance values and comparing nearby edges and nodes to find sequences of strong flows. For cases of *QSE*, analyses of the pixelated *2D* abundance plots (like Fig. 2) are appropriate, because it was noted in Ref. [86] that the shape of these groups in the *N-Z* plane can critically impact final abundance predictions. *CNNs* would be useful for such analyses, as well as *KNN* and *SVM* clustering algorithms.

Tracer Particles – In a hydrodynamics simulation, tracer particles (*TPs*) are passive Lagrangian mass elements that move along with the fluid; in spite of the name, they are not actually individual “particles” like a proton. By tracing (recording) the time-dependence of *TP* properties – position, velocities, angular momentum, temperature, density, composition – a characterization of complex fluid flows can be obtained. *TPs* are widely utilized in nucleosynthesis research to extract hydrodynamic profiles from simulations coupling multi-dimensional hydrodynamics with a truncated treatment of thermonuclear burning; by later following a full thermonuclear burn simulation over *each TP* profile and combining the results, a fuller treatment of thermonuclear burning can be obtained. Examples of this *TP* approach include studies of core collapse supernovae [88, 89], binary neutron star mergers [90], and Type Ia supernovae [91, 92]. A few thousand *TPs* are typically used, in some cases evenly distributed across the entire spatial grid [89] and in other cases strategically located to track critical spatial regions.

There are a number of issues, however, that arise from the use of *TPs*. These include [88]: the impact of initial tracer positions and velocities; the precise times when tracer particles are initiated and terminated in the simulation; the number of tracer particles to deploy; the challenges of obtaining convergence of post-processed abundances; problems with methods to add more tracer particles *after* the simulation is completed (which can aid in convergence); and that *TP* approaches do not generate uncertainties. There are other issues including possible discontinuities in velocity assignments [93] as well as inadequate spatial resolution and inconsistent thermodynamic evolution. For core-collapse supernovae simulations, questions can arise as to which tracers are ejected from the explosion, and the very divergent nucleosynthesis results in asymmetrical explosions [89].

ML approaches can address some of issues arising from *TP* use. For example, outlier trajectories that may prevent convergence of post-processing nucleosynthesis abundances could be identified and removed using *IF*, *SVM*, and *NNs*, or clustering algorithms like *KNN* or *KM*. *GPs* could be used to smooth tracer particle trajectories to make them more generic or “representative” of the astrophysical environment, removing

complex structure that may be tied to a specific model; this could help deal with velocity discontinuities or convergence issues. Going further, it could be advantageous to determine a smaller set of “pseudo-tracers” that could each represent many individual *TPs*. These could be centroids of a cluster of *TPs* identified by *KNN* or *KM*, especially in combination with a *GP* that assigned uncertainties to *TP* trajectories. Runs with 10000 pseudo-tracers could then enable analyses with one to two orders of magnitude improved spatial resolution with little increase in computational power; this could help with abundance convergence and spatial resolution problems. Alternatively, an ensemble of pseudo-tracer runs with initial variations in parameters such as positions/velocities, start/stop times, or others could be used to address these and other issues in a more computationally efficient manner. Pseudo-tracers could also help reveal hidden flow structures or mark transitions between different flow regions (*e.g.*, strong outflows vs. convective regions).

For uncertainties, *GPs* could be utilized to add uncertainties to tracer particle trajectories, which could then be propagated through post-processing simulations with a Monte Carlo uncertainty quantification (UQ) approach as discussed below in Section 3.3. Finally, regarding nucleosynthesis in asymmetrical systems, clustering algorithms may be useful to divide *TP* trajectories into groups that, when appropriately mass weighted, could be used to determine final abundances in the system weighted over respective contributions from (for example) polar vs. equatorial trajectories. Combined with *GPs*, this could also determine uncertainties in these weighted abundances.

3.2 Complexity Reduction

Emulators – Nucleosynthesis simulations can be considered functions that map inputs (initial abundances, thermonuclear reaction rates, hydrodynamic conditions) to outputs (final abundances, nuclear energy generation). As discussed in Section 3.1, simulations with the most realistic hydrodynamics and a full treatment of thermonuclear burning are not yet computationally viable. By replacing such simulations with fast-executing approximations (*i.e.*, emulators), the complexity of the problem is reduced, and more realistic simulations become viable. In this way, emulators may accelerate the pace of scientific discovery, and also enable the multiple runs of more realistic simulations as required for determining their uncertainties (see Section 3.3) as well as the sensitivities of their outputs to inputs (see Section 3.4).

NNs are a popular foundation for emulators as they are known to be universal function approximators. Specifically, a *NN* with fixed depth and arbitrary width can, to any specified accuracy, approximate *any* continuous function when the activation functions are continuous and nonpolynomial [94, 95]. Numerous studies of different *NN* depths, widths, and activations have since been studied in this regard; for a review, see Ref. [96]. By training a *DNN* on a set of nucleosynthesis simulation inputs and outputs, for example, that simulation would be effectively reverse engineered: loading a new set of inputs (within range of the training data) into a trained emulator would produce a new and consistent set of outputs. As with all ML approaches, running emulators with inputs outside of the training data range (*i.e.*, making extrapolations) can produce problematic outputs; retraining the emulator is usually required for extrapolations. Besides *DNNs*, other ML algorithms have been used as emulators, including *CNNs*, *GPs*, *SVMs*, and *RFs*. The choice of algorithm depends on the complexity of the problem, the size of the training data set, the available computational resources, the desired accuracy of the approximation, and the goals of the study.

For nucleosynthesis codes that couple hydrodynamics and thermonuclear burning, a natural first goal could be to replace either the hydrodynamics or thermonuclear burn with an emulator. To replace the hydrodynamics, an approach similar to that used in Ref. [97] could be used, wherein a *CNN*-based architecture trained on only 16 simulations was found to calculate turbulent fluid dynamics more accurately than classical numerical solvers on a comparably low resolution spatial grid. Their modest system was

also able to capture the behavior generated by the *athena++* solver [98], a state-of-the-art magnetohydrodynamics code used for high-performance computing astrophysics simulations.

An alternative approach is to emulate the thermonuclear burn (rather than hydrodynamics) calculations in a fully coupled nucleosynthesis code: at each time step, the emulator would give approximate results for the change in abundances of all the tracked species. This approach can offer improved performance if the emulator runs faster than the traditional linearized thermonuclear burn solution. As discussed below, the execution time of traditional simulations scales as the square of the number of tracked isotopes, so emulators will give more performance gains for simulations tracking many hundreds to thousands of isotopes. This approach was first attempted for supernova nucleosynthesis in [33] with a *DNN* emulating a (very approximate) three-isotope system, and careful attention was given in that study to integrating the hydrodynamics with the emulator.

There are alternatives to using emulators that precisely mimic the approach used by standard nucleosynthesis codes, wherein all isotopic abundances are evolved through short time steps from initial to final values. Examples include: studies that predict only final abundances for all tracked isotopes in a system; studies to determine final abundances of a few particular isotopes to compare to observations (*e.g.*, nova contributions to the galactic ${}^7\text{Li}$ abundance [99]); studies that track only isotopes with abundance values above a certain threshold; or studies to predict an observable light curve (*e.g.*, an X-ray burst [100]), especially valuable when there is little if any ejected material to observe. An example of the first can be found in a recent preprint [34], where an 80-isotope simulation of nucleosynthesis in a massive star was emulated with a *NN* with two hidden layers of 256 nodes each. They trained their emulator on a set of final abundances of 6×10^4 simulations (with variable initial temperature and density conditions) and obtained – in some cases – qualitative agreement between the emulator and the traditional simulation. They plan to enlarge the training data set and utilize alternate *NN* architectures to explore the possibilities of quantitative agreement with the stellar simulations.

Moving farther away from traditional approaches, emulators could be constructed to approximate a series of images of isotopic abundances on a nuclide chart, which essentially form the frames of a simulation animation as sketched in Fig. 6. As briefly mentioned above in Section 3.1, such an emulator may be constructed with a *CNN* for image analysis combined with an *RNN* for frame-to-frame memory. Similarly, a *GNN* could be combined with an *RNN* to emulate a nucleosynthesis flow diagram animation. With these approaches, the time dependence of abundances or flows could be extracted from the emulated animation frames, thereby providing a novel nucleosynthesis solver.

An important approach for emulators is to utilize physics equations into the loss function of a *NN*. Such “physics-driven” ML approaches have demonstrated improved performance over traditional *NN* systems, and typically require far less data for training. This approach was successfully used in Ref. [101] to predict fluid flows (using the Navier-Stokes equations), in Ref. [102] to predict seismic responses (using equations of motion), and in Ref. [103] to solve a nonlinear inverse problem in geological drilling (using a parameterized Earth model).

Given these successes, it is likely that the first ML emulation of a realistic nucleosynthesis simulation will be created in the near future.

Reduce Simulation Dimensionality – The execution time for thermonuclear burn simulations roughly scales with the *square* of the number of nuclear species (isotopes) N whose abundances are followed. This is because at each simulation time step, an N by N matrix – constructed to detail how each species transmutes into others – must be inverted to determine a linearized solution of the abundance

changes [36, 35, 64]. As discussed previously, realistic nucleosynthesis simulations with multi-dimensional hydrodynamics necessarily truncate their thermonuclear burning treatment – by following fewer nuclei – to be computationally viable. “Alpha nuclei” (e.g., ^{12}C , ^{16}O , ^{20}Ne ... ^{56}Fe) are typically chosen for this as their interlinking alpha-capture reactions (e.g., $^{12}\text{C}(\alpha,\gamma)^{16}\text{O}$...) generate most of the thermonuclear energy in the system; a full treatment of the evolution of the complete isotopic inventory then requires an extraction of hydrodynamic profiles and a subsequent post-processing simulation.

For some astrophysical scenarios, it may be possible to use a similar approach of following fewer isotopes – specifically, those that can represent *much* of the isotopic inventory evolution (as well as the energy generation, if desired). Specifically, it may be possible to use a *PCA* or other dimension-reducing ML algorithm [104, 105] to help determine a small set of *pseudo-isotopes*, each of which represents multiple nuclide species, and track only the changes in their respective *pseudo-abundances*. The smaller number of tracked species will result in a significant reduction in complexity of the system and in the execution time of a standard nucleosynthesis simulation. The choice of pseudo-isotopes could be verified by running an ensemble of traditional simulations and identifying (a) the most abundant isotopes in a traditional simulation and (b) other “nearby” isotopes whose abundances ratios to them are similar throughout the ensemble. Such an analysis could be done with *KNN*, *KM*, or other clustering algorithms.

A more ML-centric approach would be to use a *VAE* as a simulation emulator, where the smaller set of *pseudo-abundances* (determined by the algorithm) form the lower dimensional space. The simplest approach would be to train the algorithm on final abundances only. In cases where it may be desirable to more closely mimic a traditional nucleosynthesis simulation, training could be done on the abundances at each time step (e.g., Fig. 6). For analyses of such time-series data, *VAEs* (which use *NNs*) have some advantages in dimension reduction [106] over more traditional approaches such as wavelet decomposition or *PCAs*. A complementary approach would be to use an *RNN* to analyze the time-series data.

Overall, the advantages of using *pseudo-isotopes* is that their presence and identity could reveal hidden structures in the nucleosynthesis process that are not evident in traditional simulations. The *pseudo-isotopes* determined for one simulation could be reused for computationally-efficient parameter space explorations, or potentially reused in simulations of other astrophysical environments. Furthermore, a comparison of such *pseudo-isotopes* across simulations of different astrophysical environments could give a new perspective that aids in deconvolving nuclear and hydrodynamic effects that work together to create the elements of the Universe.

Approximating Flows – In the above discussion of “alpha nuclei” like ^{12}C and ^{16}O , the interlinking alpha-capture reactions like $^{12}\text{C}(\alpha,\gamma)^{16}\text{O}$ were mentioned as being critical to track energy generation in this particular approximation to thermonuclear burning. This approach can be generalized to consider *pseudo-reactions* which link together *pseudo-isotopes* (or perhaps regions of the nuclide chart) for nucleosynthesis studies needing to approximate reaction flow and/or energy generation. ML techniques like those described above – *PCAs*, *VAEs* – could be used to identify the most critical reactions during particular time windows of a simulation, then the time evolution could be studied with *VAEs* or *RNNs*. Because the focus here is on reaction flow, it is natural to use a *GNN* (as described in Section 3.1) in combination with a *VAE* to reduce the dimensionality of the problem. Furthermore, by combining treatments of both *pseudo-isotopes* and interlinking *pseudo-reactions*, both reaction flows and isotopic abundances could be approximated simultaneously.

Some nucleosynthesis studies focus on reaction flow patterns that repeat over different portions of the nuclide chart. As mentioned in Section 3.1 above, Fig. 4 gives an example of repetitive thermonuclear

burning cycles that occur in a nucleosynthesis simulation of a nova explosion: the *NeNa*, *MgAl*, and *SiP* cycles, named for the isotopes involved. When such cycles occur in an explosive astrophysical environment, they are considered a “trap” or “sink” for material that otherwise would be processed by thermonuclear burning up to higher mass isotopes. These traps can be characterized by their “leakage rate” of material that escapes the trap and resumes a flow to higher masses. By representing all the isotopes in each cycle as a single *pseudo-isotope*, and interlinking cycles with *pseudo-reactions* that characterize the leakage from one to the next, the complexity of the reaction flow can be greatly reduced. A similar technique was used in a traditional nucleosynthesis simulation of an X-ray burst [107], primarily to speed up the simulation execution. However, this novel work was never replicated, perhaps in part because advances in computing hardware made this approximation less necessary for post-processing nucleosynthesis studies. It could, however, prove fruitful to explore possible insights that may arise from this approach when applied to more realistic nucleosynthesis simulations.

Functional Representations – Current approaches to solving the time-evolution of abundances, and the accompanying thermonuclear energy generation, in a given astrophysical environment involves numerically solving a first order set of coupled differential equations, usually through linearizing the problem and employing an implicit differencing approach for numerical stability (see, e.g., [36]). For 1D post-processing approaches, these simulations can be completed rather quickly with very modest computing requirements. In the distant past, however, computational power was much more limited, and significant efforts were spent devising analytical approximations to estimate abundance changes and energy generation. In some cases, valuable scaling laws were devised (see, e.g., [108, 109, 110]) to approximate how predictions of these quantities changed with temperature or other variables, thereby significantly reducing the complexity of the model and allowing further detailed investigations.

While the field has largely moved away from these approaches, scaling laws may offer valuable insights – perhaps into the complex interplay of hydrodynamics and thermonuclear burning, or perhaps for “back-of-the-envelope” approximations – not readily apparent from numerical solutions. They may also facilitate large-scale explorations of parameter spaces because these approximations are so fast to calculate.

ML approaches could revitalize such scaling studies, by first creating an emulator to approximate the input-output mapping (as described above), then using a symbolic regression (*SR*) algorithm with a set of temperature-dependent functions. Such an approach could generate, for example, approximations for the values of abundance ratios, or the total nuclear energy generation, as a function of peak temperature in an astrophysical system. It is possible that such ML techniques could yield nucleosynthesis scaling laws as valuable as those used to explain the revival of a stalled shock wave in a core collapse supernova by neutrino heating [111] – a work that revolutionized studies of the core collapse mechanism. There are, in fact, recent suggestions that using *NNs* in combination with *SRs* are not just a way to approximate the behavior of a physical system, but also a way to uncover previously hidden physical laws [112].

3.3 Uncertainty Quantification

The quantitative comparison of nucleosynthesis simulation predictions to observations is critical for advancing studies of many astrophysical environments. One example comes from the Big Bang, where comparisons of the primordial ${}^7\text{Li}$ abundance to model predictions [113], the “cosmic lithium problem”, has driven studies of diverse effects including lithium depletion in stars, dark matter, exotic particles, thermonuclear reaction rates, and observational techniques [114]. The material ejected from Type Ia supernovae provide another example: observations of isotopic nickel abundances from these explosions are compared to nucleosynthesis predictions to discriminate between different explosion mechanisms [92].

Measurements of isotopic abundance ratios in grains of meteorites provide a third example, where comparisons to nucleosynthesis predictions are used to attribute the origin of some meteorites to Asymptotic Giant Branch stars and Carbon stars [115], and others to novae, supernovae, or elsewhere in the cosmos [116, 117].

Uncertainties in simulation predictions are needed to make robust comparisons to observations. While lacking in many nucleosynthesis research efforts, uncertainty quantification (*UQ*) treatments are now becoming routine for many simulation-based physics studies [118]. A widely-utilized *UQ* treatment involves propagating input uncertainties through a simulation [119] using Monte Carlo input sampling. Specifically, a large (1000 - 10000) ensemble of simulations are executed, each of which has small independent random variations of the input parameters over their respective probability distribution functions; the ensemble outputs are then analyzed to determine the prediction uncertainties. This approach has previously been used for nucleosynthesis studies of the big bang [113], novae [120], X-ray bursts [81], and red giant stars and other scenarios [121]. However, it has not been used for more realistic multi-dimensional nucleosynthesis simulations due to the long execution time required for each run of the ensemble.

This could change by utilizing any of the ML complexity-reduction approaches mentioned above in Section 3.2. Specifically, the fast execution times of simulation emulators, pseudo-abundances, flow approximations, or functional representations could make viable Monte Carlo *UQ* approaches with multi-dimensional simulations that couple hydrodynamics with thermonuclear burning. This approach could also facilitate *ID* coupled and post-processing simulations that employ the larger thermonuclear reaction networks needed for more complex problems, especially in cases where larger ensemble sizes needed for more precise uncertainty determinations.

Another ML-based *UQ* approach is based on *BNNs*. Gaining use in theoretical nuclear physics [122], *BNNs* employ stochastic weights for *NNs* used in a Bayesian framework. By constructing a simulation emulator with a *BNN*, uncertainties of the *NN* predictions are naturally generated. A related approach is based on deep ensembles [123] wherein numerous *DNNs*, here used as emulators, are trained with random initializations of their model parameters. Deep ensembles can be viewed as a Bayesian approach using delta-function posteriors, but have the added flexibility to capture different posterior modes if needed [124].

Deep ensembles are just one of many approaches to *UQ* that can be used for deep learning (*i.e.*, *DNN*-based) studies [125]; others include Monte Carlo dropout, Bootstrapping, and Gaussian Mixture Models [126]. There are also specialized *UQ* approaches developed for *DNNs* that model time-series data [127]; an application of this for nucleosynthesis would be to use an *RNN* to predict a set of abundance vs. time values with uncertainties. Finally, most of the approaches mentioned above can be configured to incorporate uncertainties arising from training data as well as those from the *DNN* architecture itself; in this manner, a more complete uncertainty characterization can be obtained.

3.4 Sensitivity Analyses

Sensitivity Analyses (*SA*) are examinations of changes in one or more target outputs of a simulation resulting from variation of an individual input, while keeping all other inputs and model parameters fixed. *SA* are widely utilized to understand the relationships of inputs to outputs in complex systems, and especially to identify those inputs which have the strongest impact on outputs. *SA* have been widely used in nuclear astrophysics, for example to identify which thermonuclear reaction inputs or hydrodynamic conditions – when varied – have a significant impact on predicted isotopic abundances or nuclear energy generation (see, *e.g.*, Refs. [128, 129, 130]).

SA typically involve less than 100 simulation runs, far fewer than those of the Monte Carlo-based *UQ* study described above in Section 3.3. In spite of this, *SA* are often computationally prohibitive for nucleosynthesis models with multi-dimensional hydrodynamics and full thermonuclear burning. Similar to aiding *UQ* studies, the ML complexity-reducing strategies discussed in Section 3.2, such as *NN*-based emulators, can be used to facilitate *SA* of more realistic nucleosynthesis models. In this way, better guidance can be given for setting priorities for future studies of simulation inputs.

Since the variation of inputs are much larger in *SA* (orders of magnitude) compared to that in *UQ* studies (usually a few standard deviations), the corresponding ML algorithm training parameter space is much larger for *SA* than for *UQ* determinations. A recent example from reactor physics illustrated how a *DNN* constructed to emulate a high-fidelity reactor simulation was used for both *UQ* and *SA* [131]. This study utilized a Group Method of Data Handling approach [132] with a *DNN* for parametric optimization in high-dimensional spaces in a manner that may be useful for emulating nucleosynthesis simulations.

There are other ways in which ML algorithms can help with *SA*. One approach is to include a “feature importance” score as an added domain knowledge in *DTs* [133] (or with other categorization algorithms like *RFs*). By identifying important features, this is complementary to traditional *SA* approaches. Another technique is to use *GPs* or *BNNs* to identify which simulation predictions have the largest uncertainties, and set these as target outputs for a traditional *SA*. A third approach is active learning [134], where an iterative procedure (either manual or automated) is used to determine which new data (or features) would reduce prediction uncertainties (for example, from a *GP* or *BNN*) when added to the training data set. While active learning is primarily used to boost efficiency via the use of smaller training data sets, it can also help identify critical data that change predictions and/or reduce prediction uncertainties.

These examples show that there are a number of ways that traditional *SA* can be advanced, or complemented, by ML approaches, and this can be very useful for identifying the most critical nucleosynthesis simulation inputs. There are also a number of deeper connections between *SA* and ML, which are discussed in Ref. [135]. These include how *SA* can be used for ML interpretability, and how some ML approaches are redevelopments of earlier work in *SA* (e.g., using Gaussian processes as emulators to speed up *SA* [136]).

3.5 Improved Models

Some of the ML techniques mentioned above have focused on working with existing model outputs to find patterns or correlations in nucleosynthesis flows that can improve our understanding, while others focused on speeding up model execution (with emulators, approximations, or dimension reductions) to aid *UQ* and *SA* that can guide future studies. This subsection describes how ML approaches can be used to improve models, independently or in combination with those efforts mentioned above, specifically by examining model inputs and outputs.

Thermonuclear reaction rates are critical simulation inputs, with some astrophysical environments requiring thousands of input rates to fully describe the relevant nucleosynthesis. These rates are temperature dependent and are collected in large libraries; see Ref. [130] for a detailed discussion. Rates are determined by convoluting an energy-dependent nuclear reaction cross section with the temperature-dependent Maxwell-Boltzmann distribution of relative energies of nuclei in an astrophysical environment [137, 36]. For the widely utilized *REACLIB* library [138] containing 55000 rates, this convolution is performed numerically for each rate and then fit to a 7-parameter analytical temperature-dependent function, and the fit parameters for each reaction are then stored in the library and subsequently input into a simulation. It is very challenging, however, to obtain precise parameter fits because the rates vary by up to 30 orders of

magnitude over temperatures relevant for nucleosynthesis; precision is needed because fit deviations of a few percent can significantly alter nucleosynthesis predictions. ML approaches for regression including *LinR*, *KRR*, *GP*, *GB*, *SVM*, and *NNs* could be very useful to better determine these reaction rate fit parameters.

Rate improvements can also be realized by examining the underlying reaction cross sections. While some cross sections are determined from experimental measurements, the majority (nearly 90%) are determined from theoretical reaction models [130]. A single flawed rate can distort model predictions, but it is challenging to individually check thousands of rates for anomalies. Theoretically based rates, however, usually exhibit smooth variations across the N - Z plane, making some ML approaches well suited to look for outliers based on comparisons of the shapes of cross section *vs.* energy (or rate *vs.* temperature) curves for reactions spanning the nuclide chart. Specifically, *IF*, *SVM*, and *NNs* may help with this, as well as clustering algorithms like *KNN* and *KM*.

In addition to finding outliers of theoretical rates, ML approaches can aid in improving cross sections determined from measurements. A recent example is the use of *KNN* and *DT* algorithms to improve evaluations of the $^{233}\text{U} + n$ and $^{35}\text{Cl}(n,p)$ reactions [139] that are important for nuclear security and nuclear energy applications. Another is the use of deep ensembles to extract features from resonances in nuclear scattering [140], as an alternative to phenomenological R-matrix analyses. Approaches similar to these could be useful for reactions needed for nuclear astrophysics. There is also a long-term effort to revamp the workflow for producing evaluated cross section libraries using nuclear data elements (datasets, evaluated cross sections, validation benchmarks) wrapped in software containers, linked in a Bayesian network, and updated via *GPs*; first results are reported in Ref. [141]. Additional examples of ML utilizations for improving nuclear data and nuclear physics can be found in Ref. [122].

ML approaches can also be utilized to improve the hydrodynamic inputs for nucleosynthesis simulations. As discussed in Section 3.1, tracer particle input for nucleosynthesis simulations can contain outliers that could be identified and removed with *IF*, *SVM*, and *NNs* or with *KNN*, *KM*, or other clustering algorithms. Also, *GPs* could be used to smooth tracer particle trajectories to make them more “representative” of the astrophysical environment rather than tied to the specific hydrodynamics model that generated them.

The same *GP*-based smoothing approach could also be utilized for modifying hydrodynamic profiles extracted from multiple spatial zones. Fig. 9 shows such (unsmoothed) temperature profiles of zones extracted from a nova simulation [142]; *GP* smoothing could be done in a manner that removes significant zone-to-zone discontinuities that may skew post-processing nucleosynthesis predictions. ML outlier approaches (as mentioned above) could also be helpful to search for anomalous profiles extracted from “full” (coupled hydrodynamics + thermonuclear burning) models, whereas clustering (*KNN* and *KM*) and other ML techniques could improve the consistency of profile extractions from the full coupled model output.

While there are other numerical and techniques that could be used for the above input manipulations, the ML approaches mentioned could produce excellent results that are complementary to traditional techniques use for regression, smoothing, and outlier detection.

Another approach to improving models is to identify outliers in model predictions or outputs. For example, when running a Monte Carlo ensemble of simulations for *UQ* as described above in Section 3.3, ML outlier approaches could be used to flag any anomalous outputs; the inputs for that particular simulation run could then be examined for anomalies and appropriately corrected. This “back-tracing” of outlying outputs to find anomalous inputs with ML could also be invaluable for parameter space explorations which are so

widely used in many fields including nucleosynthesis (e.g., Ref. [143]); more discussions of parameter space explorations are given in the next subsection.

3.6 Scientific Workflows

Exploring Parameter Space

The above subsections have detailed some of the discovery potential of ML to improve our understanding and approximations of nucleosynthesis flows, to perform SA, to determine prediction uncertainties, and to improve inputs for nucleosynthesis simulations. Many of these ML-based investigations likely begin by developing an appropriate model, obtaining or generating an extensive training dataset, and training and adjusting the model. For some projects, it is then necessary to execute the model hundreds or thousands of times to explore some particular input parameter space, followed by the critical steps of processing, analyzing, visualizing, and comparing results of each run before choosing the next set of inputs. The immense parameter space of some problems – which for nuclear astrophysics could be initial abundances, hydrodynamic histories over a multi-dimensional spatial grid, and thermonuclear reaction rate values – along with the complexity of operations at each execution makes many such projects computationally intractable without extensive parameter space truncations and/or model approximations. As an example, the parameter study of core-collapse supernova models in Ref. [143] used approximations for generating an explosion (artificially depositing energy in the core) as well as for exploring nucleosynthesis (tracking only 13 nuclides, followed by post-processing calculations that tracked 211 species).

In some fields, ML approaches are now removing the need for such approximations when exploring large parameter spaces. An example is the high-profile 3D protein folding problem, where new protein discoveries can tremendously impact drug development, medical research, environmental remediation, and many other fields. *AlphaFold* [144], a DNN protein folding model designed for these searches, was able to predict over 350000 stable 3D protein folding structures in 2021; this number grew to 2×10^7 one year later. In comparison, traditional techniques had yielded only 70000 structures in 60 years of work, a pace requiring $\sim 10^9$ years to match the performance of *AlphaFold*. Similarly, the *GNoME* Project [145] used GNNs to explore the vast parameter space of crystalline structures for materials science research. It predicted 2.2 million new structures, of which about one quarter are thought to be especially stable; this work thereby expanded the number of such materials known to man by an order of magnitude.

A promising approach for parameter space explorations may lie reinforcement learning, an ML approach in which an “agent” interacts with an environment and learns in stages to make decisions that maximize a reward function [146]. LLMs are now being used in this context to generate rewards that outperform those engineered by human experts [147], without the need to train on task-specific examples. As the use of reinforcement learning grows in physics [148], it is quite possible that LLM-powered reinforcement learning could soon be utilized in nucleosynthesis studies to, for example, efficiently steer explorations of the expansive parameter spaces that were described above.

Such examples give a glimpse into the rapidly-expanding capabilities of ML for exploring large parameter spaces – a proficiency that could be extremely useful for nucleosynthesis research.

Transforming Workflows

The workflow of most research projects includes many important, but rote, tasks such as literature searches/summaries and generating reports and presentations. For many, these and related efforts are serious productivity bottlenecks that can limit time spent generating new insights and discoveries. Fortunately,

Large Language Models (*LLMs*) can be used to reduce some of these burdens, as well as streamline and accelerate many other critical aspects of scientific workflows.

Triggered by the development of the *Transformers* architecture [62] and the use of (up to) $\sim 10^{11}$ *DNN* model parameters trained on enormous textual datasets of $\sim 10^{12}$ tokens (words), *LLMs* have rapidly (since 2022) gained wide acceptance through their ability to adeptly answer text queries (denoted as “prompts”) via a familiar chat interface (*e.g.*, *ChatGPT* [63]). *LLMs* can quickly provide detailed background information, summarize and synthesize literature results, refine text passages, outline and draft reports and papers, generate computer codes, process and visualize data, recommend and prioritize research problems, and perform many other important research tasks. In essence, *LLMs* can function as research assistants.

There is currently a race to develop larger and ever-more capable *LLMs* and *LLM*-based tools, especially those with multi-modal (text, audio/video, data analysis, coding, and more) capabilities. As with any new approach, *LLMs* have limitations and problems, both technical [149] and conceptual [150]; some of these are discussed in Section 4 below along with recent approaches to their solution. In spite of their issues, the use of *LLMs* in research, and their acceptance as an invaluable tool, is rapidly growing. It would be advantageous for nucleosynthesis researchers to capitalize on the rapidly-developing capabilities of *LLMs* for their work.

4 CHALLENGES

Given the widespread utilization of ML in many fields, much attention has been given to enumerating and addressing the challenges and limitations of these approaches. We briefly discuss some critical issues (and their possible solutions) below; for more details, see the review on ML for physics in Ref. [151].

Training Datasets

The success of any ML model depends in large part to the availability of large, high-quality, low-noise, bias-free data sets for training. For many research problems, these are not easy to obtain or generate. The first step is to determine the training data contents needed for a particular study. For simulation-based studies of nucleosynthesis, the contents may likely consist of collections of (*input*, *output*) pairs, where the *input* would describe the astrophysical environment type (*e.g.*, core-collapse supernova) and properties (*e.g.*, mass distribution, entropy, and electron fraction in the core), initial abundances, a thermonuclear reaction rate library, and thermodynamics information such as initial conditions for hydrodynamics calculations or temperature and density histories for post-processing simulations. The *output* depends on the nature of the problem being addressed. For the studies discussed in Section 3, these include: the final abundances for all relevant isotopes; a collection of abundances at various time steps during the simulation; those time-dependent abundances with added hydrodynamic and/or position information; images of final abundances (*e.g.*, Fig. 2) or reaction fluxes (*e.g.*, Fig. 3) on the nuclide chart; time-dependent nuclide chart images of abundances (*e.g.*, Fig. 6), reaction fluxes, or abundance derivatives (*e.g.*, Fig. 7) along with hydrodynamics information; one-dimensional plots of abundance vs. time (*e.g.*, Fig. 1); abundances and hydrodynamics information from tracer particles (see Section 3.1); or combinations of the above information.

Once the data contents for a new study are determined, a dataset can be established by processing or augmenting datasets from previous studies or external sources, or can be generated through many new executions of the relevant simulation(s). For simulation-based studies of nucleosynthesis, this latter option – running many simulations with a variety of inputs to generate “synthetic” data – is ideal, if the researcher has the appropriate astrophysical codes. One critical factor here is to ensure that the distribution of input values are representative of the actual distribution so as to avoid biased results. There is an extensive

general-purpose literature on this topic and methods to reduce biases; for details, see Refs. [152, 153]). Another key point is to include a sufficiently wide variation of inputs for these simulation runs so that they span the range needed for future use of the trained ML model. That is, the ML model should be used to generate interpolations rather than (less reliable) extrapolations whenever possible. Other well-studied, important training techniques include methods to avoid over-fitting [154], and partitioning datasets for training/testing/validation [155].

Since some nucleosynthesis researchers do not have access to simulation codes, and few researchers have access to codes for *all* astrophysical environments, a set of curated archives of nucleosynthesis simulation results would be invaluable to enable many more in the community to launch ML-based studies.

Model Choice, Execution, and Performance

For a given problem, selecting an appropriate ML algorithm to use among the many choices can be a challenge in this quickly-evolving field. For example, some algorithms adapt to larger or changing data sets better than others, and some are flexible enough to be used for multiple projects with minimal changes. There is also a steep learning curve for some ML algorithms, which constrains the number of researchers adopting these approaches – and which suggests the utility of collaborating with ML experts. Additionally, the lack of “interpretability” of many ML models – especially *NNs* – limit their acceptance by many researchers (including peer reviewers). The complex nature of the algorithms often makes it nearly impossible to determine which features of the data are responsible for the predictions [156]. However, some models are more interpretable than others, suggesting that careful model choice may help with collaborative research projects as well as with publication peer review. The topic of ML model interpretability is reviewed in Ref. [157] and is discussed in terms of sensitivity analyses in Ref. [135].

Regarding execution, most ML approaches are computationally intensive, so the availability of adequate compute resources (particularly those with graphical processing units or tensor processing units) is essential. Large storage systems with fast I/O are needed for handling the large data sets needed for model training. Knowledge of the python coding language is essential for some projects, as is access to python libraries such as pyTorch, TensorFlow, Matplotlib, NumPy, Scikit-learn, and others [158, 159].

Regarding model performance, extrapolating models to parameter regions beyond those in the training data can provide significant limitations on the potential for scientific insights from ML techniques. This can be partially addressed by interpretable models [160] and by physics-driven models [161], but is best handled by generating predictions within the bounds of the training data. Another issue is over-fitting [154] and under-fitting, both of which can lead to poor model performance. Because this is a widely known issue, there are many established mechanisms to monitor and avoid such effects [162].

Large Language Model Challenges

The extremely rapid development of *LLMs* makes it likely that any list of challenges with their use will be quickly outdated, because many current issues will be solved while new problems will arise as capabilities increase. For this reason, Ref. [163] systematically formulate 16 major *categories* of open problems for *LLMs* and then comprehensively discuss, for each, the latest examples and solution approaches. Below, we will describe some of the issues with *LLMs* that are the most relevant for scientific research: flawed output, lack of reasoning, and data handling.

For the first issue, there are many different examples of incorrect or nonsensical *LLM* output, including: when models have incomplete, biased, or no training on data in a particular scientific domain; when the training data is outdated by recent research; when their (statistically-generated) responses differ from

the “ground truth”; and when responses are nonsensical or are fabricated lies that are presented in an authoritative manner (“hallucinations” [164]). Given these numerous flaws, users must be cautious when using *LLM* output for research purposes. There are many approaches being pursued to rectify these issues. The first is better training, both in general and especially domain-specific, which could be done by some users with smaller *LLMs* on their own data. Another is requiring the *LLM* to utilize “ground truth” information – via manually-uploaded curated data sets [165], accessing a database via knowledge trees [166, 167] or (especially) knowledge graphs [168], attaching a web browser to the *LLM* [169], or using a retrieval-augmented generation (*RAG*) framework [170]. In this latter approach, relevant “expert” content is first retrieved from a source (*e.g.*, a database or the internet) and then fed to an *LLM* along with the query to generate a response along with the expert reference. Additionally, improved prompting (shorter, multi-step queries) has been shown to improve response quality, as have “mixture of experts”/“mixture of models” approaches wherein multiple models are given the same prompt and answers are polled and combined.

The second major issue – a lack of scientific reasoning – arises because *LLMs* are trained on internet-scale amounts of textual data to generate the most probable response to a query, rather than understanding the underlying scientific principles (*e.g.*, cause and effect). As a result, *LLMs* often cannot generate a step-by-step reasoning process for a response, have difficulty in verifying or reproducing results, have limited capabilities to review scientific papers, and generally perform more poorly than humans on reasoning benchmarks [171]. The overwhelming need to use and display reasoning in *LLM* responses has led to much research activity in this topic. One approach that has improved reasoning abilities is to use “chain-of-thought” prompting [172] wherein exemplars of intermediate reasoning steps are provided as *LLM* input. Another approach is “visualization of thought” [173] wherein reasoning steps are spatially visualized and followed; this approach could be useful for interpreting causality in reaction flow diagrams. A third promising approach is the ReAct framework [174] that combines tasks devoted to reasoning with those for action, in an effort to mimic human learning. Specifically, an *LLM* is prompted to generate a “chain-of-thought” reasoning response, and then to generate an appropriate action, get feedback from the environment (an observation), and repeat the process. As a hypothetical example for nucleosynthesis, if the task is to answer the question “Has the approach to uncertainty quantification used in this nucleosynthesis journal article been used for studies of other astrophysical scenarios?”, then the first reasoning result could be “I need to search this journal article and find which uncertainty quantification method was used”. The first action could then be to carry out the search, and if the observation (the search result) is null (*e.g.*, no mention of “uncertainty quantification method” in the article), then the next round of reasoning could be “I need to determine other terminologies for ‘uncertainty quantification’”, leading to a next round of action (a second search) and observation (that terms like “error analysis” or “probability distribution function determination” are possible alternatives). The iterative nature of reasoning/action/observation is a significant improvement over the return of a single response typical of most *LLMs*.

The third major issue – data handling – arises because *LLMs* are primarily trained on textual data rather than on tabular, time series, graphical, or other numerical data formats. While it may seem more appropriate to use other ML tools for data handling and analyses, *LLMs* are now becoming integrated into many scientific workflows, driving efforts to enhance their capabilities with non-textual data. One approach is to greatly expand *LLM* training datasets to include audio, video, and numerical data files. An example is provided in the recent release of *ChatGPT 4o* [175], which (among other advances) exhibits true multi-modal capabilities and can interpret user-provided tabular and graphical data. Another technique for numerical data handling is to use an “agent”-centered approach [176] wherein a larger problem is broken up into subtasks, each of which is handled by a different ML model (some of which may be *LLMs*)

playing a different role and collaborating in the overall solution. Even when all the agents are from a single *LLM*, this approach gives superior performance due to more targeted prompts. Including numerical-based models like those described in Section 2 as agents can boost performance with data handling, analysis, and interpretation tasks. This was demonstrated in *Data-Copilot* [177] where a custom workflow with multiple data-centered tools combined with an *LLM* is autonomously created in response to complex problems provided by users. For problems that require writing computer code, researchers could consider integrating a software developer agent like *Devin* [178] into their workflow.

Finally, while there are some who recommend that *LLMs* should play only a subsidiary role in research [150], there are others who are strident advocates of their use (*e.g.*, Ref. [179]). Given the current rate of development of *LLM* capabilities, it is likely that they will become an integral portion of the workflows of the next (and perhaps even the current) generation of researchers.

Overall, it is important to acknowledge that ML techniques are not a simple panacea for all research roadblocks. However, used with the appropriate caution, the examples given above illustrate the promise and possibilities of capitalizing on ML approaches to advance the methodology of simulation-based studies of nucleosynthesis and to improve our understanding of the cosmic creation of the elements.

5 SUMMARY

Machine learning approaches have proven extremely useful in many fields including astrophysics and nuclear physics but have an untapped potential in nuclear astrophysics. Some very promising utilizations of ML are for studies to advance our understanding of the complex processes that synthesize nuclides in astrophysical environments. We briefly summarize the characteristics over 25 widely-utilized ML algorithm categories, and then describe how they can be used for simulation-based nucleosynthesis studies. Specifically, we describe unexplored possibilities for ML to better understand and approximate nucleosynthetic flows, to quantify uncertainties, to perform sensitivity analyses, and to identify anomalous inputs. We also discuss how ML tools can speed up scientific workflows and improve research productivity. The use of ML to advance the decades-old methodology of simulation-based studies of nucleosynthesis has the potential to significantly improve our understanding of the cosmic creation of the elements and thereby open a new frontier in nuclear astrophysics research. This is especially the case given the rapid development of the capabilities of ML tools. Collaborations between nuclear astrophysicists and ML experts would be an excellent way to realize the promise of ML for nucleosynthesis studies.

CONFLICT OF INTEREST STATEMENT

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

AUTHOR CONTRIBUTIONS

MSS initiated this project, and MSS and DL wrote and revised the manuscript.

FUNDING

This work was supported by the U.S. Department of Energy Office of Science, Office of Nuclear Physics, under Contract Number DE-AC05-00OR22725 at ORNL.

ACKNOWLEDGMENTS

The authors wish to thank John Sparger and Austin Harris for useful discussions and comments on the manuscript.

CONTRIBUTION TO THE FIELD

This manuscript includes a summary of the functionalities of over 25 categories of popular ML algorithms, and then describes how they can be used to advance simulation-based studies of the cosmic synthesis of nuclei.

REFERENCES

- [1] Moriwaki K, Nishimichi T, Yoshida N. Machine learning for observational cosmology. *Rep. Prog. Phys.* **86** (2023) 076901. doi:10.1088/1361-6633/acd2ea.
- [2] Fraser TS, Tojeiro R, Chittenden HG. Applying unsupervised learning to resolve evolutionary histories and explore the galaxy-halo connection in IllustrisTNG. *Mon. Not. Roy. Astron. Soc.* **522** (2023) 5758–5774. doi:10.1093/mnras/stad015.
- [3] Arimura R. A machine learning approach for mass composition analysis with TALE-SD data. *Eur. Phys. J. Web of Conferences* (2023), vol. 283, 02011. doi:10.1051/epjconf/202328302011.
- [4] Hernández CA, González RE, Padilla ND. Not hydro: using neural networks to estimate galaxy properties on a dark-matter-only simulation. *Mon. Not. Roy. Astron. Soc.* **524** (2023) 4653–4669. doi:10.1093/mnras/stad2112.
- [5] González-Morán AL, Arrabal Haro P, Muñoz-Tuñón C, Rodríguez-Espinosa JM, Sánchez-Almeida J, Calhau J, et al. The PAU survey: classifying low-z SEDs using Machine Learning clustering. *Mon. Not. Roy. Astron. Soc.* **524** (2023) 3569–3581. doi:10.1093/mnras/stad2123.
- [6] Ding J, Huang Y, Li XD, Wang X, Wang Y, Yang L. Hybrid deep learning for blazar classification and correlation search with neutrinos. *Mon. Not. Roy. Astron. Soc.* **523** (2023) 4120–4135. doi:10.1093/mnras/stad1683.
- [7] den Hartogh JW, Yagüe López A, Cseh B, Pignatari M, Világos B, Roriz MP, et al. Barium stars as tracers of s-process nucleosynthesis in AGB stars. II. Using machine learning techniques on 169 stars. *Astron. Astrophys.* **672** (2023) A143. doi:10.1051/0004-6361/202244189.
- [8] Sun H. Machine Learning Refinements to Metallicity-Dependent Isotopic Abundances. *arXiv preprint arXiv:2403.02678* (2024).
- [9] Boehnlein A, Diefenthaler M, Sato N, Schram M, Ziegler V, Fanelli C, et al. Colloquium: Machine learning in nuclear physics. *Rev. Mod. Phys.* **94** (2022) 031003. doi:10.1103/RevModPhys.94.031003.
- [10] He W, Li Q, Ma Y, Niu Z, Pei J, Zhang Y. Machine learning in nuclear physics at low and intermediate energies. *Science China Physics, Mechanics & Astronomy* **66** (2023) 282001.
- [11] Xu R, Ge Z, Tian Y, Tao X, Jin Y, Zhang Y, et al. Covariance evaluation of neutron cross sections in cendl. *Eur. Phys. J. Web of Conferences* (EDP Sciences) (2023), vol. 281, 00029.
- [12] Li Z, Gao Z, Liu L, Wang Y, Zhu L, Li Q. Importance of physical information on the prediction of heavy-ion fusion cross sections with machine learning. *Phys. Rev. C* **109** (2024) 024604.
- [13] Le XK, Wang N, Jiang X. Nuclear mass predictions with multi-hidden-layer feedforward neural network. *Nucl. Phys. A* **1038** (2023) 122707.

- [14] Wu XH, Pan C, Zhang KY, Hu J. Nuclear mass predictions of the relativistic continuum Hartree-Bogoliubov theory with the kernel ridge regression. *Phys. Rev. C* **109** (2024) 024310. doi:10.1103/PhysRevC.109.024310.
- [15] Zhang X, Li W, Fang J, Niu Z. Nuclear mass predictions with the naive Bayesian model averaging method. *Nucl. Phys. A* **1043** (2024) 122820.
- [16] Li M, Sprouse TM, Meyer BS, Mumpower MR. Atomic masses with machine learning for the astrophysical r process. *Phys. Lett. B* **848** (2024) 138385.
- [17] Yüksel E, Soydaner D, Bahtiyar H. Nuclear mass predictions using machine learning models. *arXiv preprint arXiv:2401.02824* (2024).
- [18] Wu H, Wang Y, Wang Y, Deng X, Cao X, Fang D, et al. Machine learning method for ^{12}C event classification and reconstruction in the active target time-projection chamber. *Nucl. Inst. Meth. A* **1055** (2023) 168528. doi:10.1016/j.nima.2023.168528.
- [19] Sanchez-Caballero A, Alcayne V, Cano-Ott D, Mendoza E, de Rada AP. A Case Study on Deep Learning applied to Capture Cross Section Data Analysis. *Eur. Phys. J. Web of Conferences* (EDP Sciences) (2023), vol. 284, 16001.
- [20] Gargouri R, Akkoyun S, Maalej R, Damak K. Performance of machine learning algorithms on neutron activations for Germanium isotopes. *Rad. Phys. Chem.* **208** (2023) 110860.
- [21] Schatz H, Becerril Reyes AD, Best A, Brown EF, Chatziioannou K, Chipps KA, et al. Horizons: nuclear astrophysics in the 2020s and beyond. *J. Phys. G* **49** (2022) 110502. doi:10.1088/1361-6471/ac8890.
- [22] Arcones A, Thielemann FK. Origin of the elements. *Astron. Astrophys. Rev.* **31** (2023) 1. doi:10.1007/s00159-022-00146-x.
- [23] Neronov A. Multi-Messenger Astronomy. *J. Phys. Conf. Ser.* (2021), vol. 2156, 012006. doi:10.1088/1742-6596/2156/1/012006.
- [24] Wei J, Ao H, Behar S, Bultman N, Casagrande F, Cogan S, et al. Advances of the FRIB project. *Int. J. Mod. Phys. E* **28** (2019) 193003.
- [25] Motobayashi T, Sakurai H. Research with fast radioactive isotope beams at RIKEN. *Prog. Theor. Expt. Phys.* **2012** (2012) 03C001. doi:10.1093/ptep/pts059.
- [26] Scheidenberge C. NUSTAR experiments on the way from GSI to FAIR. *Eurasian J. Phys. Funct. Matl.* **1** (2017) 19. doi:10.29317/ejpfm.2017010103.
- [27] Hong B. Status of the RAON project in Korea. *Assoc. Asia Pacific Phys. Soc. Bull.* **33** (2023) 3. doi:10.1007/s43673-022-00074-z.
- [28] [Dataset] NUCLEI. Nuclear Computational Low Energy Initiative. https://nuclei.mps.ohio-state.edu/nuclei_home.php (2022).
- [29] National Academies of Sciences, Engineering, and Medicine. *Pathways to Discovery in Astronomy and Astrophysics for the 2020s* (Washington, DC: The National Academies Press) (2021). doi:10.17226/26141.
- [30] [Dataset] NICXVII. 17th International Symposium on Nuclei in the Cosmos. <https://indico.ibs.re.kr/event/548/timetable/#all> (2023).
- [31] [Dataset] OMEG. Origins of Matter and Evolution of Galaxies 16. <https://indico.maygiatoc.com/event/1/timetable/#all.detailed> (2022).
- [32] Aidala C, Aprahamian A, Bedaque P, Bernstein L, Carlson J, Carpenter M, et al. A New Era Of Discovery: The 2023 Long-Range Plan For Nuclear Science. Tech. rep., Lawrence Livermore National Laboratory (LLNL), Livermore, CA (United States) (2023).

- [33] Fan D, Willcox DE, DeGrendele C, Zingale M, Nonaka A. Neural Networks for Nuclear Reactions in MAESTROeX. *Astrophys. J.* **940** (2022) 134. doi:10.3847/1538-4357/ac9a4b.
- [34] [Dataset] Grichener A, Renzo M, Kerzendorf W, Bellinger E, Justham S, Farmer R, et al. Nuclear neural networks-implementing machine-learning methods for stellar nucleosynthesis. https://kspa.soe.ucsc.edu/sites/default/files/Aldana_Grichener___KSPA_final_report.pdf (2024).
- [35] Hix WR, Thielemann FK. Computational methods for nucleosynthesis and nuclear energy generation. *J. Comp. Appl. Math.* **109** (1999) 321–351.
- [36] Arnett D. *Supernovae and nucleosynthesis: an investigation of the history of matter, from the big bang to the present*, vol. 7 (Princeton University Press) (1996).
- [37] Kumar S, Bhatnagar V. Review of Regression Models in Machine Learning. *J. Intelligent Sys. and Computing* **3** (2022) 40–47.
- [38] Hastie T, Tibshirani R, Friedman J, Hastie T, Tibshirani R, Friedman J. Kernel smoothing methods. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (2009) 191–218.
- [39] Bisong E, Bisong E. Logistic regression. *Building Machine Learning and Deep Learning Models on Google Cloud Platform: A Comprehensive Guide for Beginners* (2019) 243–250.
- [40] Breiman L. *Classification and regression trees* (Routledge) (2017).
- [41] Quinlan JR. Induction of decision trees. *Machine learning* **1** (1986) 81–106.
- [42] Cutler A, Cutler DR, Stevens JR. *Random Forests* (New York, NY: Springer New York), chap. 5 (2012), 157–175. doi:10.1007/978-1-4419-9326-7_5.
- [43] Friedman JH. Greedy function approximation: A gradient boosting machine. *Annals of Statistics* **29** (2001) 1189 – 1232. doi:10.1214/aos/1013203451.
- [44] Liu FT, Ting KM, Zhou ZH. Isolation Forest. *2008 Eighth IEEE Int. Conf. Data Mining* (2008), 413–422. doi:10.1109/ICDM.2008.17.
- [45] Steinwart I, Christmann A. *Support vector machines* (Springer Science & Business Media) (2008).
- [46] Xu J, Lange K. Power k-means clustering. Chaudhuri K, Salakhutdinov R, editors, *Proc. 36th Int. Conf. Machine Learning* (PMLR) (2019), vol. 97, 6921–6931.
- [47] Zhang Z. Introduction to machine learning: k-nearest neighbors. *Ann. Translational Med.* **4** (2016) 218. doi:10.21037/atm.2016.03.37.
- [48] Kherif F, Latypova A. Chapter 12 - Principal Component Analysis. Mechelli A, Vieira S, editors, *Machine Learning* (Academic Press) (2020), 209–225. doi:https://doi.org/10.1016/B978-0-12-815739-8.00012-2.
- [49] Webb GI, Keogh E, Miikkulainen R. Naïve Bayes. *Encyclopedia of machine learning* **15** (2010) 713–714.
- [50] Kingma DP, Welling M, et al. An introduction to variational autoencoders. *Foundations and Trends in Machine Learning* **12** (2019) 307–392.
- [51] Cava WL, Orzechowski P, Burlacu B, de Franca FO, Virgolin M, Jin Y, et al. Contemporary symbolic regression methods and their relative performance. *Thirty-fifth Conf. on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)* (2021), 1.
- [52] Rasmussen CE, Williams CKI. *Gaussian Processes for Machine Learning* (The MIT Press) (2005). doi:10.7551/mitpress/3206.001.0001.
- [53] Suk HI. Chapter 1 - An Introduction to Neural Networks and Deep Learning. Zhou SK, Greenspan H, Shen D, editors, *Deep Learning for Medical Image Analysis* (Academic Press) (2017), 3–24. doi:https://doi.org/10.1016/B978-0-12-810408-8.00002-X.

- [54] Ankile LL, Heggland MF, Krange K. Deep convolutional neural networks: A survey of the foundations, selected improvements, and some current applications. *arXiv preprint arXiv:2011.12960* (2020).
- [55] Jospin LV, Laga H, Boussaid F, Buntine W, Bennamoun M. Hands-On Bayesian Neural Networks—A Tutorial for Deep Learning Users. *IEEE Computational Intelligence Mag.* **17** (2022) 29–48. doi:10.1109/MCI.2022.3155327.
- [56] Lipton ZC, Berkowitz J, Elkan C. A Critical Review of Recurrent Neural Networks for Sequence Learning. *arXiv preprint arXiv:1506.00019* (2015).
- [57] Wu Z, Pan S, Chen F, Long G, Zhang C, Yu PS. A comprehensive survey on graph neural networks. *IEEE Trans. Neural Networks and Learning Sys.* **32** (2021) 4–24. doi:10.1109/TNNLS.2020.2978386.
- [58] Lee CC, Chung PC, Tsai JR, Chang CI. Robust radial basis function neural networks. *IEEE Trans. Systems, Man, and Cybernetics, Part B (Cybernetics)* **29** (1999) 674–685.
- [59] Kasim MF, Watson-Parris D, Deaconu L, Oliver S, Hatfield P, Froula DH, et al. Building high accuracy emulators for scientific simulations with deep neural architecture search. *Machine Learning: Science and Technology* **3** (2021) 015013.
- [60] Gonog L, Zhou Y. A review: Generative adversarial networks. *2019 14th IEEE Conference on Industrial Electronics and Applications (ICIEA)* (2019), 505–510. doi:10.1109/ICIEA.2019.8833686.
- [61] Zhao WX, Zhou K, Li J, Tang T, Wang X, Hou Y, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223* (2023).
- [62] Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. *Adv. Neural Info. Proc. Sys.* **30** (2017).
- [63] [Dataset] OpenAI. ChatGPT. <https://chatgpt.com> (2022).
- [64] Hix WR, Meyer BS. Thermonuclear kinetics in astrophysics. *Nuc. Phys. A* **777** (2006) 188–207. doi:10.1016/j.nuclphysa.2004.10.009.
- [65] [Dataset] Smith, Michael S. Computational Infrastructure for Nuclear Astrophysics (CINA). <https://nucastrodata.org/infrastructure/> (2024).
- [66] Nesaraja C, Lingerfelt E, Scott J, Smith M, Hix W, Bardayan D, et al. A new computational infrastructure for nuclear astrophysics. *Nucl. Phys. A* **758** (2005) 174C.
- [67] Smith MS, Lingerfelt EJ, Scott JP, Nesaraja CD, Hix WR, Chae K, et al. Computational Infrastructure for Nuclear Astrophysics. *Proc. Origin of Matter and Evolution of Galaxies* (AIP Conf. Proc. 847) (2006), 470.
- [68] [Dataset] Hix WR. XNET. <https://eagle.phys.utk.edu/xnet/trac/>; <https://github.com/starkiller-astro/XNet> (2024).
- [69] van Wormer L, Görres J, Iliadis C, Wiescher M, Thielemann FK. Reaction Rates and Reaction Sequences in the rp-Process. *Astrophys. J.* **432** (1994) 326. doi:10.1086/174572.
- [70] [Dataset] NUGRID Collaboration. NUGRID. https://nugrid.github.io/content/codes_collab.html (2024).
- [71] Paxton B, Bildsten L, Dotter A, Herwig F, Lesaffre P, Timmes F. Modules for Experiments in Stellar Astrophysics (MESA). *Astrophys. J. Suppl. Ser.* **192** (2011) 3.
- [72] Weaver TA, Zimmerman GB, Woosley SE. Presupernova evolution of massive stars. *Astrophys. J.* **225** (1978) 1021–1029. doi:10.1086/156569.
- [73] Joyce M, Tayar J. A Review of the Mixing Length Theory of Convection in 1D Stellar Modeling. *Galaxies* **11** (2023) 75. doi:10.3390/galaxies11030075.

- [74] Fryxell B, Olson K, Ricker P, Timmes FX, Zingale M, Lamb DQ, et al. FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes. *Astrophys. J. Suppl. Ser.* **131** (2000) 273–334. doi:10.1086/317361.
- [75] Almgren AS, Beckner VE, Bell JB, Day MS, Howell LH, Joggerst CC, et al. CASTRO: A New Compressible Astrophysical Solver. I. Hydrodynamics and Self-gravity. *Astrophys. J.* **715** (2010) 1221–1238. doi:10.1088/0004-637X/715/2/1221.
- [76] Dunham SJ, Endeve E, Mezzacappa A, Blondin JM, Buffaloe J, Holley-Bockelmann K. A Parametric Study of the SASI Comparing General Relativistic and Non-Relativistic Treatments. *arXiv preprint arXiv:2307.10904* (2023).
- [77] Papatheodore T, Messer OB. Exposing hierarchical parallelism in the flash code for supernova simulation on summit and other architectures. *Exascale Scientific Applications: Scalability and Performance Portability* (2017) 95.
- [78] Navó G, Reichert M, Obergaulinger M, Arcones A. Core-collapse supernova simulations with reduced nucleosynthesis networks. *Astrophys. J.* **951** (2023) 112.
- [79] Smith MS, Rehm KE. Nuclear Astrophysics Measurements with Radioactive Beams. *Ann. Rev. Nucl. Part. Sci.* **51** (2001) 91. doi:10.1146/annurev.nucl.51.101701.132430.
- [80] Smith MS. Nuclear data for astrophysics research: A new online paradigm. *Journal of the Korean Physical Society* **59** (2011) 761–766. doi:10.3938/jkps.59.761.
- [81] Roberts LF, Hix WR, Smith MS, Fisker JL. Monte Carlo simulations of Type I X-ray burst nucleosynthesis. *Proc. of Sci.* (2006) 202.
- [82] Lippuner J, Roberts LF. SkyNet: A Modular Nuclear Reaction Network Library. *Astrophys. J. Suppl. Ser.* **233** (2017) 18. doi:10.3847/1538-4365/aa94cb.
- [83] Xu Z, Hu J, Deng W. Recurrent convolutional neural network for video classification. *2016 IEEE International Conference on Multimedia and Expo (ICME)* (2016), 1–6. doi:10.1109/ICME.2016.7552971.
- [84] Hou Y, Zheng L, Gould S. Learning to structure an image with few colors. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2020), 10116. doi:10.1109/ICME.2016.7552971.
- [85] Zhang L, Smith MS, Hix WR, Lingerfelt EJ. Abundance Correlations in Explosive Nucleosynthesis Simulations. *AIP Conf. Proc.* (2013), vol. 1594, 437.
- [86] Meyer BS, Krishnan TD, Clayton DD. Theory of Quasi-Equilibrium Nucleosynthesis and Applications to Matter Expanding from High Temperature and Density. *Astrophys. J.* **498** (1998) 808–830. doi:10.1086/305562.
- [87] Seshan A. Using Machine Learning to Augment Dynamic Time Warping Based Signal Classification. *arXiv preprint arXiv:2206.07200* (2022).
- [88] Sieverding A, Waldrop PG, Harris JA, Hix WR, Lentz EJ, Bruenn SW, et al. Tracer particles for core-collapse supernova nucleosynthesis: The advantages of moving backward. *Astrophys. J.* **950** (2023) 34. doi:10.3847/1538-4357/acc8d1.
- [89] Harris JA, Hix WR, Chertkow MA, Lee C, Lentz EJ, Messer OB. Implications for post-processing nucleosynthesis of core-collapse supernova models with Lagrangian particles. *Astrophys. J.* **843** (2017) 2.
- [90] Bovard L, Rezzolla L. On the use of tracer particles in simulations of binary neutron stars. *Classical and Quantum Gravity* **34** (2017) 215005. doi:10.1088/1361-6382/aa8d98.

- [91] Seitenzahl IR, Röpke FK, Fink M, Pakmor R. Nucleosynthesis in thermonuclear supernovae with tracers: convergence and variable mass particles. *Mon. Not. Roy. Astron. Soc.* **407** (2010) 2297–2304. doi:10.1111/j.1365-2966.2010.17106.x.
- [92] Seitenzahl IR, Pakmor R. Nucleosynthesis and Tracer Methods in Type Ia Supernovae. Tanihata I, Toki H, Kajino T, editors, *Handbook of Nuclear Physics* (Singapore: Springer Nature Singapore) (2020), 1–34. doi:10.1007/978-981-15-8818-1_87-1.
- [93] Tiede C, Zrake J, MacFadyen A, Haiman Z. How Binaries Accrete: Hydrodynamic Simulations with Passive Tracer Particles. *Astrophys. J.* **932** (2022) 24. doi:10.3847/1538-4357/ac6c2b.
- [94] Cybenko G. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems* **2** (1989) 303–314.
- [95] Park S, Yun C, Lee J, Shin J. Minimum width for universal approximation. *arXiv preprint arXiv:2006.08859* (2020).
- [96] DeVore R, Hanin B, Petrova G. Neural network approximation. *Acta Numerica* **30** (2021) 327–444.
- [97] Stachenfeld K, Fielding DB, Kochkov D, Cranmer M, Pfaff T, Godwin J, et al. Learned coarse models for efficient turbulence simulation. *arXiv preprint arXiv:2112.15275* (2021).
- [98] Stone JM, Tomida K, White CJ, Felker KG. The athena++ adaptive mesh refinement framework: Design and magnetohydrodynamic solvers. *Astrophys. J. Suppl. Ser.* **249** (2020) 4.
- [99] Starrfield S, Bose M, Iliadis C, Hix WR, Woodward CE, Wagner RM. Hydrodynamic Simulations of Oxygen–Neon Classical Novae as Galactic ${}^7\text{Li}$ Producers and Potential Accretion-induced Collapse Progenitors. *Astrophys. J.* **962** (2024) 191.
- [100] Galloway DK, Chenevez J, Wörpel H, Keek L, Ootes L, Watts AL, et al. The multi-instrument burst archive (MINBAR). *Astrophys. J. Suppl. Ser.* **249** (2020) 32.
- [101] Ma H, Zhang Y, Thuerey N, Hu X, Haidn OJ. Physics-Driven Learning of the Steady Navier-Stokes Equations Using Deep Convolutional Neural Networks. *Comm. Comp. Phys.* **32** (2022) 715–736. doi:https://doi.org/10.4208/cicp.OA-2021-0146.
- [102] Zhang R, Liu Y, Sun H. Physics-guided convolutional neural network (PhyCNN) for data-driven seismic response modeling. *Engineering Structures* **215** (2020) 110704.
- [103] Jin Y, Shen Q, Wu X, Chen J, Huang Y. A physics-driven deep-learning network for solving nonlinear inverse problems. *Petrophysics* **61** (2020) 86–98.
- [104] Fodor IK. A survey of dimension reduction techniques. Tech. rep., Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States) (2002).
- [105] Espadoto M, Martins RM, Kerren A, Hirata NS, Telea AC. Toward a quantitative survey of dimension reduction techniques. *IEEE Trans. Visualization and Computer Graphics* **27** (2019) 2153–2173.
- [106] Todo W, Laurent B, Loubes JM, Selmani M. Dimension Reduction for time series with Variational AutoEncoders. *arXiv preprint arXiv:2204.11060* (2022).
- [107] Rembges F, Freiburghaus C, Rauscher T, Thielemann FK, Schatz H, Wiescher M. An Approximation for the rp-Process. *Astrophys. J.* **484** (1997) 412–423. doi:10.1086/304300.
- [108] Caughlan GR, Fowler WA. The mean lifetimes of carbon, nitrogen, and oxygen nuclei in the cno bicycle. *Astrophys. Journal* **136** (1962) 453.
- [109] Parker PD, Bahcall JN, Fowler WA. Termination of the proton-proton chain in stellar interiors. *Astrophys. J.* **139** (1964) 602.
- [110] Clayton DD. *Principles of stellar evolution and nucleosynthesis* (University of Chicago Press) (1983).
- [111] Bethe HA, Wilson JR. Revival of a stalled supernova shock by neutrino heating. *Astrophys. J.* **295** (1985) 14–23.

- [112] Cranmer M, Sanchez Gonzalez A, Battaglia P, Xu R, Cranmer K, Spergel D, et al. Discovering symbolic models from deep learning with inductive biases. *Adv. Neural Info. Proc. Sys.* **33** (2020) 17429–17442.
- [113] Smith MS, Kawano LH, Malaney RA. Experimental, computational, and observational analysis of primordial nucleosynthesis. *Astrophys. J. Suppl. Ser.* **85** (1993) 219–247.
- [114] Fields B. Big Bang Nucleosynthesis: Nuclear Physics in the Early Universe. Tanihata I, Toki H, Kajino T, editors, *Handbook of Nuclear Physics* (Singapore: Springer Nature Singapore) (2023), 111. doi:10.1007/978-981-15-8818-1_111-1.
- [115] Liu N, Barosch J, Nittler LR, O'D Alexander CM, Wang J, Cristallo S, et al. New Multielement Isotopic Compositions of Presolar SiC Grains: Implications for Their Stellar Origins. *Ap. J. Lett.* **920** (2021) L26. doi:10.3847/2041-8213/ac260b.
- [116] Nittler LR. Presolar stardust in meteorites: recent advances and scientific frontiers. *Earth Plan. Sci. Lett.* **209** (2003) 259–273. doi:10.1016/S0012-821X(02)01153-6.
- [117] Zinner EK. Presolar Grains. *Treatise on Geochemistry* **1** (2003) 711. doi:10.1016/B0-08-043751-6/01144-0.
- [118] Ghanem R, Higdon D, Owhadi H, et al. *Handbook of uncertainty quantification*, vol. 6 (Springer New York) (2017).
- [119] Kroese DP, Taimre T, Botev ZI. *Handbook of Monte Carlo Methods* (John Wiley & Sons) (2013).
- [120] Hix WR, Smith MS, Starrfield S, Mezzacappa A, Smith DL. Impact of nuclear reaction rate uncertainties on nova models. *Nucl. Phys. A* **718** (2003) 620–622.
- [121] Rauscher T, Nishimura N, Cescutti G, Hirschi R, Murphy ASJ. Abundance uncertainties obtained with the PIZBUIN framework for Monte Carlo reaction rate variations. *AIP Conf. Proc.* (AIP Publishing) (2018), vol. 1947, 020015.
- [122] Boehnlein A, Diefenthaler M, Sato N, Schram M, Ziegler V, Fanelli C, et al. Colloquium: Machine learning in nuclear physics. *Reviews of Modern Physics* **94** (2022) 031003.
- [123] Lakshminarayanan B, Pritzel A, Blundell C. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems* **30** (2017).
- [124] Gustafsson FK, Danelljan M, Schon TB. Evaluating scalable bayesian deep learning methods for robust computer vision. *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition Workshops* (2020), 318–319.
- [125] Abdar M, Pourpanah F, Hussain S, Rezazadegan D, Liu L, Ghavamzadeh M, et al. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information Fusion* **76** (2021) 243–297.
- [126] Hubschneider C, Hutmacher R, Zöllner JM. Calibrating uncertainty models for steering angle estimation. *2019 IEEE intelligent transportation systems conference (ITSC)* (IEEE) (2019), 1511–1518.
- [127] Song T, Ding W, Liu H, Wu J, Zhou H, Chu J. Uncertainty quantification in machine learning modeling for multi-step time series forecasting: Example of recurrent neural networks in discharge simulations. *Water* **12** (2020) 912.
- [128] Beun J, Blackmon JC, Hix WR, McLaughlin GC, Smith MS, Surman R. Neutron capture on ^{130}Sn during r-process freeze-out. *J. Phys. G: Nucl. Part. Phys.* **36** (2008) 025201.
- [129] Mumpower MR, Surman R, McLaughlin G, Aprahamian A. The impact of individual nuclear properties on r-process nucleosynthesis. *Prog. Part. Nucl. Phys.* **86** (2016) 86–126.
- [130] Smith MS. Nuclear data resources and initiatives for nuclear astrophysics. *Front. Astron. Space Sci.* **10** (2023) 1243615. doi:10.3389/fspas.2023.1243615.

- [131] Radaideh MI, Kozlowski T. Analyzing nuclear reactor simulation data and uncertainty with the group method of data handling. *Nuclear Engineering and Technology* **52** (2020) 287–295.
- [132] Ivakhnenko AG. Polynomial theory of complex systems. *IEEE Trans. Syst., Man, and Cybernetics* (1971) 364–378.
- [133] Al Iqbal MR, Rahman S, Nabil SI, Chowdhury IUA. Knowledge based decision tree construction with feature importance domain knowledge. *2012 7th international conference on electrical and computer engineering (IEEE)* (2012), 659–662.
- [134] [Dataset] Settles B. Active learning literature survey. <https://minds.wisconsin.edu/bitstream/handle/1793/60660/TR1648.pdf> (2009).
- [135] Scholbeck CA, Moosbauer J, Casalicchio G, Gupta H, Bischl B, Heumann C. Position paper: Bridging the gap between machine learning and sensitivity analysis. *arXiv preprint arXiv:2312.13234* (2023).
- [136] Le Gratiet L, Marelli S, Sudret B. *Metamodel-Based Sensitivity Analysis: Polynomial Chaos Expansions and Gaussian Processes* (Springer International Publishing), chap. 38 (2017), 1289–1325. doi:10.1007/978-3-319-12385-1_38.
- [137] Rolfs CE, Rodney WS. *Cauldrons in the cosmos: Nuclear astrophysics* (University of Chicago press) (1988).
- [138] Cyburt RH, Amthor AM, Ferguson R, Meisel Z, Smith K, Warren S, et al. The JINA REACLIB database: its recent updates and impact on type-I X-ray bursts. *Astrophys. J. Suppl. Ser.* **189** (2010) 240.
- [139] Vicente-Valdez P, Bernstein L, Fratoni M. Nuclear data evaluation augmented by machine learning. *Annals of Nuclear Energy* **163** (2021) 108596.
- [140] Kim CH, Chae KY, Smith MS, Bardayan DW, Brune CR, deBoer RJ, et al. Eliminating unobservable parameters in physics models using deep learning. *arXiv preprint arXiv:2305.02623* (2023).
- [141] Schnabel G, Capote R, Koning A, Brown D. Nuclear data evaluation with Bayesian networks. *arXiv preprint arXiv:2110.10322* (2021).
- [142] Politano M, Starrfield S, Truran J, Weiss A, Sparks W. Hydrodynamic studies of accretion onto massive white dwarfs: ONeMg-enriched nova outbursts. I. Dependence on white dwarf mass. *Astrophys. J.* **448** (1995) 807.
- [143] Nakamura T, Umeda H, Nomoto K, Thielemann FK, Burrows A. Nucleosynthesis in type II supernovae and the abundances in metal-poor stars. *Astrophys. J.* **517** (1999) 193.
- [144] Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, et al. Highly accurate protein structure prediction with AlphaFold. *Nature* **596** (2021) 583–589.
- [145] Merchant A, Batzner S, Schoenholz SS, Aykol M, Cheon G, Cubuk ED. Scaling deep learning for materials discovery. *Nature* **624** (2023) 80–85.
- [146] Arulkumaran K, Deisenroth MP, Brundage M, Bharath AA. Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine* **34** (2017) 26–38.
- [147] Ma YJ, Liang W, Wang G, Huang DA, Bastani O, Jayaraman D, et al. Eureka: Human-level reward design via coding large language models. *arXiv preprint arXiv:2310.12931* (2024).
- [148] Martín-Guerrero JD, Lamata L. Reinforcement learning and physics. *Applied Sciences* **11** (2021) 8589.
- [149] Nejjar M, Zacharias L, Stiehle F, Weber I. LLMs for Science: Usage for Code Generation and Data Analysis. *arXiv preprint arXiv:2311.16733* (2023).
- [150] Birhane A, Kasirzadeh A, Leslie D, Wachter S. Science in the age of large language models. *Nature Reviews Physics* **5** (2023) 277–280.

- [151] Carleo G, Cirac I, Cranmer K, Daudet L, Schuld M, Tishby N, et al. Machine learning and the physical sciences. *Rev. Mod. Phys.* **91** (2019) 045002.
- [152] Jeong W, Lee K, Yoo D, Lee D, Han S. Toward reliable and transferable machine learning potentials: uniform training by overcoming sampling bias. *J. Physical Chem. C* **122** (2018) 22790–22795.
- [153] Mehrabi N, Morstatter F, Saxena N, Lerman K, Galstyan A. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)* **54** (2021) 1–35.
- [154] Ying X. An overview of overfitting and its solutions. *Journal of physics: Conference series* (IOP Publishing) (2019), vol. 1168, 022022.
- [155] Xu Y, Goodacre R. On splitting training and validation set: a comparative study of cross-validation, bootstrap and systematic sampling for estimating the generalization performance of supervised learning. *J. Analysis and Testing* **2** (2018) 249–262.
- [156] Barbierato E, Gatti A. The Challenges of Machine Learning: A Critical Review. *Electronics* **13** (2024).
- [157] Rudin C, Chen C, Chen Z, Huang H, Semenova L, Zhong C. Interpretable machine learning: Fundamental principles and 10 grand challenges. *Statistic Surveys* **16** (2022) 1–85.
- [158] Raschka S. *Python machine learning* (Packt publishing ltd) (2015).
- [159] Saabith AS, Vinothraj T, Fareez M. Popular python libraries and their application domains. *Int. J. Adv. Eng. Res. Dev.* **7** (2020).
- [160] Muckley ES, Saal JE, Meredig B, Roper CS, Martin JH. Interpretable models for extrapolation in scientific machine learning. *Digital Discovery* **2** (2023) 1425–1435.
- [161] Brahma I, Jennings R, Freid B. Using physics to extend the range of machine learning models for an aerodynamic, hydraulic and combusting system: The toy model concept. *Energy and AI* **6** (2021) 100113. doi:<https://doi.org/10.1016/j.egyai.2021.100113>.
- [162] Salman S, Liu X. Overfitting mechanism and avoidance in deep neural networks. *arXiv preprint arXiv:1901.06566* (2019).
- [163] Kaddour J, Harris J, Mozes M, Bradley H, Raileanu R, McHardy R. Challenges and applications of large language models. *arXiv preprint arXiv:2307.10169* (2023).
- [164] Yao JY, Ning KP, Liu ZH, Ning MN, Yuan L. LLM lies: Hallucinations are not bugs, but features as adversarial examples. *arXiv preprint arXiv:2310.01469* (2023).
- [165] [Dataset] Perplexity AI. Perplexity AI. <https://www.perplexity.ai/> (2024).
- [166] [Dataset] Dietterich TG. What's wrong with large language models and what we should be building instead. <https://web.engr.oregonstate.edu/~tgd/talks/dietterich-fixing-llms-cods-comad-2024.pdf> (2023).
- [167] Sarthi P, Abdullah S, Tuli A, Khanna S, Goldie A, Manning CD. RAPTOR: Recursive Abstractive Processing for Tree-Organized Retrieval. *arXiv preprint arXiv:2401.18059* (2024).
- [168] [Dataset] Diffbot. Diffbot. <https://www.diffbot.com/> (2024).
- [169] [Dataset] Google. Google Gemini. <https://gemini.google.com> (2024).
- [170] Lewis P, Perez E, Piktus A, Petroni F, Karpukhin V, Goyal N, et al. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. Larochelle H, Ranzato M, Hadsell R, Balcan M, Lin H, editors, *Adv. Neural Info. Proc. Syst.* (Curran Associates, Inc.) (2020), vol. 33, 9459–9474.
- [171] Yu P, Wang T, Golovneva O, AlKhamissi B, Verma S, Jin Z, et al. ALERT: Adapting language models to reasoning tasks. *arXiv preprint arXiv:2212.08286* (2022).
- [172] Wei J, Wang X, Schuurmans D, Bosma M, Ichter b, Xia F, et al. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. Koyejo S, Mohamed S, Agarwal A, Belgrave D, Cho K, Oh

A, editors, *Advances in Neural Information Processing Systems* (Curran Associates, Inc.) (2022), vol. 35, 24824–24837.

[173] Wu W, Mao S, Zhang Y, Xia Y, Dong L, Cui L, et al. Visualization-of-Thought Elicits Spatial Reasoning in Large Language Models. *arXiv preprint arXiv:2404.03622* (2024).

[174] Yao S, Zhao J, Yu D, Du N, Shafran I, Narasimhan K, et al. ReAct: Synergizing Reasoning and Acting in Language Models. *arXiv preprint arXiv:2210.03629* (2023).

[175] [Dataset] Open AI. Hello GPT-4o. <https://openai.com/index/hello-gpt-4o/> (2024).

[176] Wu Q, Bansal G, Zhang J, Wu Y, Zhang S, Zhu E, et al. Autogen: Enabling next-gen LLM applications via multi-agent conversation framework. *arXiv preprint arXiv:2308.08155* (2023).

[177] Zhang W, Shen Y, Lu W, Zhuang Y. Data-copilot: Bridging billions of data and humans with autonomous workflow. *arXiv preprint arXiv:2306.07209* (2023).

[178] [Dataset] Cognition Labs. Introducing Devin, the first AI software engineer. <https://www.cognition-labs.com/introducing-devin> (2024).

[179] Chen L, Ahmed NK, Dutta A, Bhattacharjee A, Yu S, Mahmud QI, et al. Position Paper: The Landscape and Challenges of HPC Research and LLMs. *arXiv preprint arXiv:2402.02018* (2024).




Figure 1. Nucleosynthesis simulation predictions shown as *ID* plots of abundances *versus* time for numerous different isotopes. The zero time is set at the peak temperature of the event. The calculation was made with the Computational Infrastructure for Nuclear Astrophysics (CINA) [65, 66, 67] running the XNET post-processing simulation code [68].




Figure 2. Nucleosynthesis abundance diagram where color indicates abundances for each isotope on the chart of the nuclides ($N-Z$ plane). The simulation was made with CINA [65].




Figure 3. Nucleosynthesis flow diagram where arrows indicate reaction flux from one isotope to the next through individual thermonuclear reactions on the chart of the nuclides. The simulation was made with CINA [65].

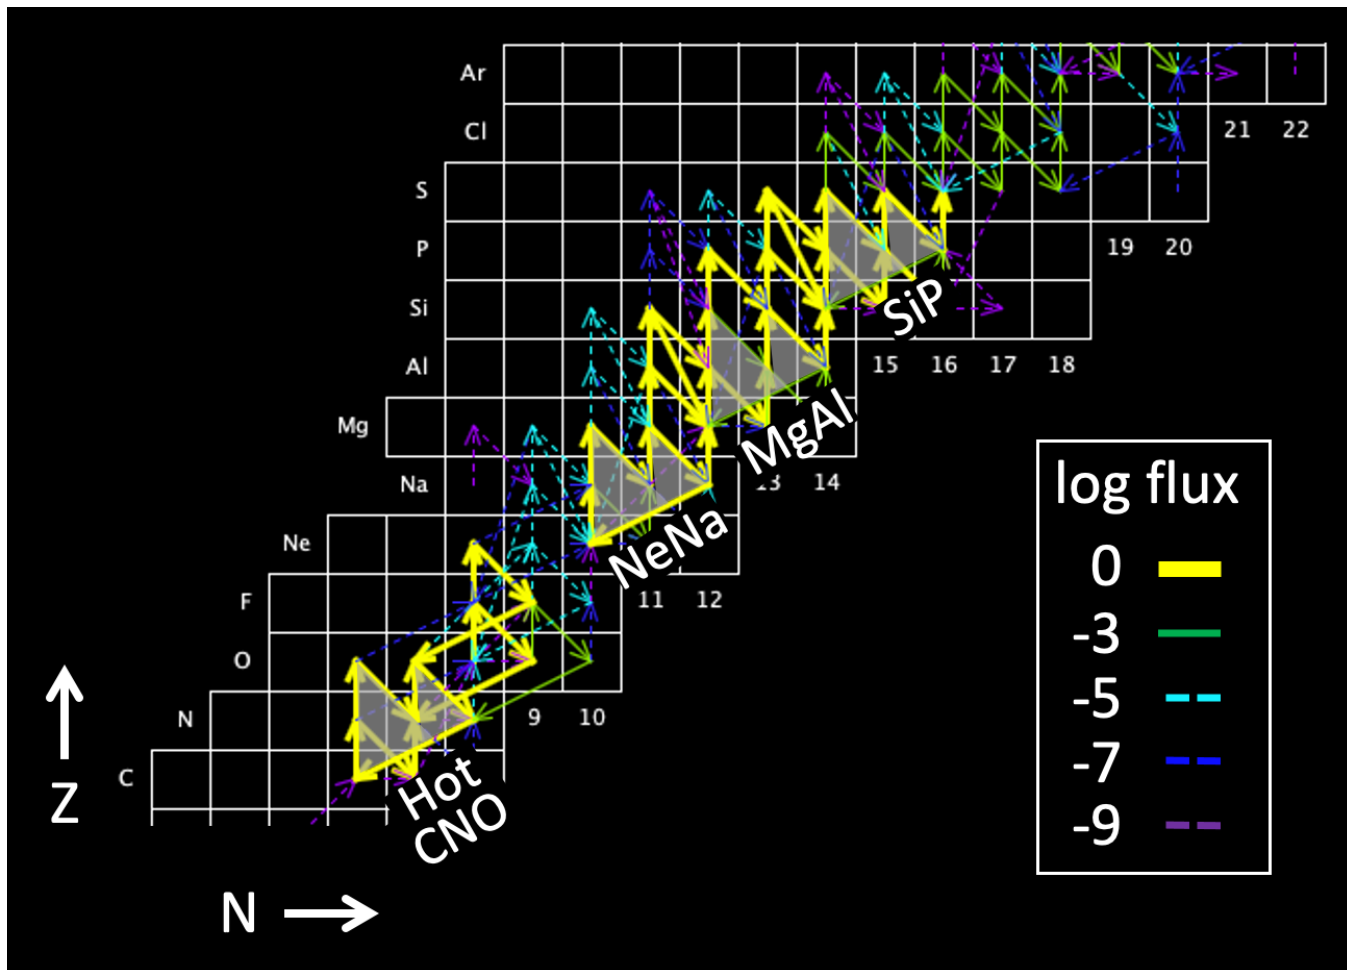


Figure 4. Repetitive reaction flow patterns – the *Hot CNO*, *NeNa*, *MgAl*, and *SiP* cycles – in a nucleosynthesis simulation of an energetic nova explosion. The simulation was made with CINA [65].




Figure 5. The prominent thermonuclear reactions that create (destroy) a given proton-rich nuclide are shown as orange (blue) arrows connecting isotopes on the N - Z plane. Inverse reactions and some other reactions are not shown for clarity. For stable and for neutron-rich nuclides, the prominent reactions are different.

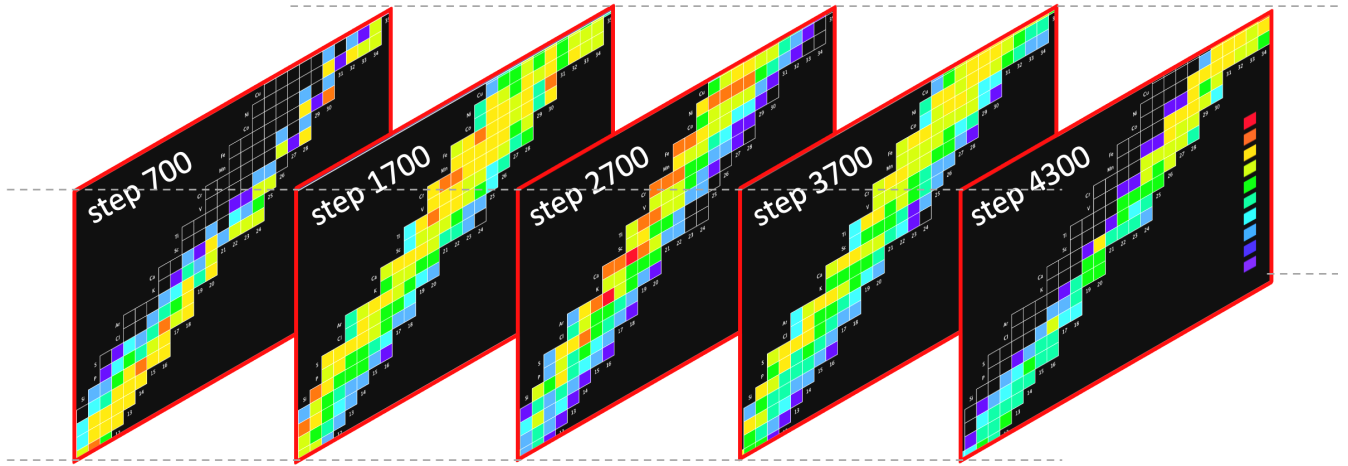


Figure 6. Snapshots of the evolution of isotopic abundances at different time steps of a simulation of a Type I X-ray burst. The simulation was made with CINA [65].




Figure 7. Evolution of the sign of the time derivative of abundances on the nuclide chart (N - Z plane) for an energetic nova explosion at times before the peak temperature (a), nearly at peak (b), and after peak (c). The simulation was made with CINA [65].

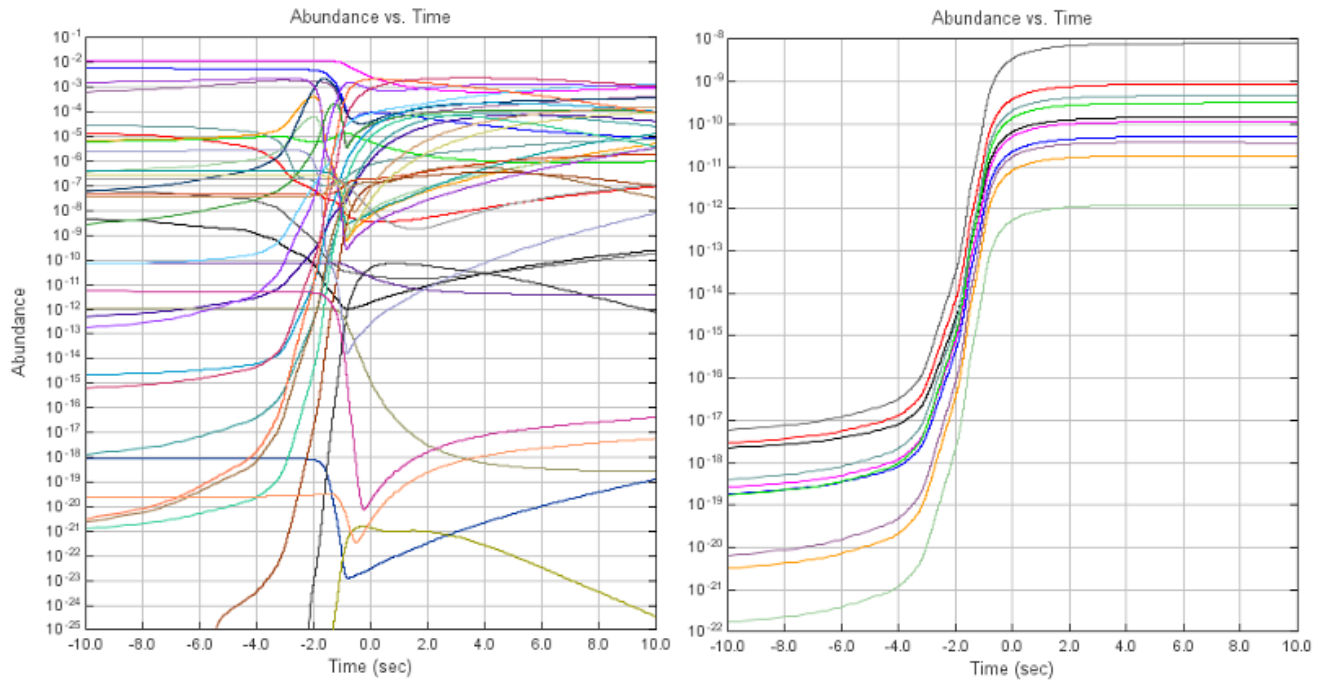


Figure 8. (Left) Abundance vs. time histories for many isotopes tracked in a nucleosynthesis simulation; (Right) subset of these isotopes exhibiting similar abundance vs. time behavior [85]. The zero time is set at the peak temperature of the event. The simulation was made with CINA [65].




Figure 9. Temperature profiles for some of the hydrodynamic zones of an energetic nova explosion [142]. The innermost (outermost) zones have the highest (lowest) peak temperatures. Some zones are not plotted for clarity. The zero time is set at the peak temperature of the event. Inset: The same profiles shown over a longer time scale.