

Transfer learning to model inertial confinement fusion experiments

K. D. Humbird*[†], J. L. Peterson*, B. K. Spears*, R. G. McClarren[‡]

*Lawrence Livermore National Laboratory, 7000 East Ave, Livermore, CA 94550

[†]Department of Nuclear Engineering, Texas A & M University, 3133 TAMU, College Station, TX 77843

[‡]Department of Aerospace and Mechanical Engineering, University of Notre Dame, 365 Fitzpatrick Hall, Notre Dame, IN 46556

Abstract—Inertial confinement fusion (ICF) experiments are designed using computer simulations that are approximations of reality, and therefore must be calibrated to accurately predict experimental observations. In this work, we propose a novel technique for calibrating from simulations to experiments, or from low fidelity simulations to high fidelity simulations, via “transfer learning”. Transfer learning is a commonly used technique in the machine learning community, in which models trained on one task are partially retrained to solve a separate, but related task, for which there is a limited quantity of data. We introduce the idea of hierarchical transfer learning, in which neural networks trained on low fidelity models are calibrated to high fidelity models, then to experimental data. This technique essentially bootstraps the calibration process, enabling the creation of models which predict high fidelity simulations or experiments with minimal computational cost. We apply this technique to a database of ICF simulations and experiments carried out at the Omega laser facility. Transfer learning with deep neural networks enables the creation of models that are more predictive of Omega experiments than simulations alone. The calibrated models accurately predict future Omega experiments, and are used to search for new, optimal implosion designs.

I. INTRODUCTION

Many physical systems and experiments are designed using models – analytical theories or computer simulations that attempt to take into account the various components of the system to determine the efficiency, performance, and reliability of a design. In applications where the physics of the system is well known, the models can be accurate depictions of reality, however when dealing with systems at extreme conditions (such as extremely high temperatures, pressures, and densities), the physics is not as well understood and the models are not always validated with experimental data. This is often the case in inertial confinement fusion (ICF), in which lasers are used to compress a small fuel capsules filled with deuterium and tritium to high density, temperature, and pressure in order to create conditions that are favorable for nuclear fusion reactions [1], [2]. The computer simulations that model ICF experiments are complex and involve a wide variety of physics models: radiation hydrodynamics, atomic physics, nuclear burn physics, laser and plasma physics, magnetic field effects, and more [3]. These codes are validated in certain regimes, but acquiring data at the extreme conditions reached in ICF experiments is challenging and expensive, thus the accuracy of the models away from the validation data is not well known. Furthermore, fully-integrated simulations

which model everything from the laser beam propagation to the particle transport within the fuel capsule are extremely expensive to run in 3D with high resolution. Researchers often need to make many approximations, such as running the simulation in 2D with axi-symmetric constraints, in order to efficiently search the parameter space for promising experimental designs. The use of surrogate models trained on large databases of simulations have enabled rapid exploration of design spaces [4], [5], but these models are only as good of a representation of reality as the simulations upon which they are trained.

When the simulations used to design new ICF experiments contain simplifying assumptions, it is expected that there will be discrepancies between the simulation prediction and what is observed in the experiment. A common statistical approach to correcting an inaccurate simulator is model calibration – using experimental data to “calibrate” an inaccurate model to produce one which is more consistent with reality. Model calibration is a broadly-researched topic [6]–[9], with one of the most popular techniques developed by Kennedy and O’Hagan [10]. In this approach, the true model is assumed to be an additive combination of a simulator, Gaussian distributed error due to measurement uncertainty, and an unknown discrepancy term which is learned using experimental data. The form of this discrepancy term is often specified by the user (for example, the user might chose a Gaussian process discrepancy with a particular kernel function) and the complexity is limited by the amount of experimental data that is available. Researchers have explored Bayesian calibration with discrepancy terms for ICF data [11]; in this work, we propose an alternative calibration technique borrowed from the machine learning community, referred to as “transfer learning”, for calibrating ICF models.

Traditional machine learning models gain knowledge by observing large quantities of labeled data. For example, if the task is to classify photos of animals, a model will need to be exposed to millions of labeled images of all the animals it is expected to classify, in a variety of different scenarios, colors, perspectives, etc, in order to classify the animals correctly. Supervised learning tasks are straightforward to solve when large quantities of data are available, however many of these techniques fail when limited to small sets of labeled data.

Transfer learning is an alternative learning technique that can help overcome the challenge of training on small data

sets. Transfer learning is a method for using knowledge gained while solving one task, and applying it to a different, but related task. This approach is most commonly used for image classification [12]–[14], for which there are many large databases of labeled images [15]–[17] and several pre-trained neural network models that are available for download, such as AlexNet [18] and Inception [19]. These neural networks have been studied extensively, and they appear to learn how to recognize images in a logical series of steps as you traverse hidden layers in the models. First, the networks often search for edges in the images, then for simple geometric patterns, and eventually begin to recognize specific characteristics, such as eyes, arms, ears, etc [19]. In general, the deeper in the network you go, the finer the details the network appears to focus on. It therefore seems logical to expect that for a neural network trained on any image dataset, the first several hidden layers are essentially the same – they learn about features common to all images. One can thus take a pre-trained neural network, freeze the first several layers of weights, and focus on re-training only the last few layers on a new image dataset to learn how to appropriately classify this new set of images. The old frozen layers are where the network learns to “see”; the last few are where the network learns to “recognize” specific types of images. Transfer learning is the process of taking a network trained on a large dataset, freezing several layers of the network, then retraining the unfrozen layers on a different, often smaller dataset. As an example relevant to ICF, researchers at the National Ignition Facility (NIF) [20] have used transfer learning to classify images of different types of damage that occur on the optics at the NIF. There are not enough labeled optics images to train a network from scratch, but transfer learning with a network pre-trained on ImageNet [15] produces models which classify optics damage with over 98% accuracy. This methodology has enabled the group to automate their damage inspection by letting the network process the images of optical components, rather than having an optics expert inspect each image manually [21].

In ICF it is often the case that we do not have enough experimental measurements to train a machine learning model on the experimental data alone. Image classification is not the primary task of interest in ICF, so traditional transfer learning using a pre-trained open-source model is not appropriate. However, it is possible to create massive databases of ICF simulations for neural network training [5], [22]. The simulations are suspected to be a good reflection of reality, but require tuning to be more consistent with experiments. We therefore propose the use of transfer learning as a non-parametric approach for calibrating a simulation-based neural network to experiments. The general concept is illustrated in Figure 1: Train a feed-forward neural network on simulation data to relate simulation inputs to observable outputs. Next, freeze many of the layers of the neural network, but retrain a small number of layers of the network using the sparse set of experimental data for which the inputs and output observables are known.

The feasibility of using transfer learning to produce neural networks that accurately predict the outcomes of ICF experiments is tested on data generated at the Omega laser facility [23] at the Laboratory for Laser Energetics (LLE). In




Fig. 1. To transfer learn from simulations to experiments, the first three layers of the simulation-based network are frozen, and the remaining two layers are available for retraining with the experimental data.

section II we introduce a proposed hierarchical approach to transfer learning for numerical simulation and experimental data. In section III, we compare standard and hierarchical transfer learning for the Omega dataset, and use the transfer learned models to study the discrepancies between the Omega simulations and experiments in section IV.

II. HIERARCHICAL TRANSFER LEARNING

Computer simulations of complex physical systems are often modeled at varying levels of fidelity. Fast, low fidelity models are used to explore vast design spaces for optimal settings, and expensive high fidelity models might be used in interesting regions of design space to compute predictions of planned experiments. The high fidelity simulations are often more accurate and reliable than the fast, approximate models, but the expense of running the simulation often prevents their use in large parameter scans. It might be possible to create models that emulate high fidelity simulations with reduced computational cost with transfer learning. For example, a model trained on a dense set of 1D simulations could be calibrated to a sparse set of 2D simulations that fill the same design space. Rather than running a dense set of 2D simulations to train a 2D surrogate model, an equally accurate surrogate might be obtained by transfer learning from 1D to 2D with a relatively small number of 2D simulations, saving substantial computational resources. Furthermore, this 2D-calibrated model can be subsequently calibrated to experimental data. If the 2D model is a better reflection of reality than the 1D model, the transfer learning step between 2D and the experiment should be easier than jumping from 1D directly to the experimental data. We refer to this technique of transfer learning from low to subsequently higher fidelity simulations to the experimental data as “hierarchical transfer learning”.

To demonstrate the utility of hierarchical transfer learning, consider the simple function:

$$f(a, x) = xe^{ax}, \quad (1)$$

where x and a are random variables; x between $[-1,1]$ and a between $[0,1]$. This expression will be the “experiment” or true function. We also have a low and high fidelity approximation of the experiment:

$$f_{low}(a, x) = x, \quad (2)$$

$$f_{high}(a, x) = x + ax^2, \quad (3)$$

which are the first (low fidelity) and second (high fidelity) order Taylor expansions for the true function, respectively. This problem is used to study the benefits of hierarchical transfer learning; specifically to see if stepping through the hierarchy results in better models than calibrating directly from low fidelity to experimental data.

This study uses fully-connected neural networks designed with the algorithm “deep jointly-informed neural networks” (DJINN) [24]. DJINN is chosen for its efficiency in choosing an appropriate neural network architecture for arbitrary data sets, and eliminates the need for extensive hyper-parameter optimization. The algorithm uses a decision tree model trained on the data, and maps the decision tree structure into a neural network architecture with initialized weights. The networks have rectified linear unit activation functions [25], and minimize the L2 loss with weight regularization using the Adam optimizer [26] in Tensorflow [27]. DJINN performs particularly well for inertial confinement fusion data [28], [29], which is the focus of sections III and IV. For further details on the algorithm, the reader is referred to the publication or publicly-available software for DJINN [24].

For this comparison, DJINN models with 3 hidden layers are trained to map from (x, a) to $f(x, a)$ or one of the approximations in Equations 2-3. First, the average explained variance score (averaged over 5 random training/testing data splits of 80/20%) for DJINN models trained on experiments alone is computed; this is the baseline to which various transfer learning techniques is compared, as the transfer learned models are not expected to exceed the accuracy of a DJINN model trained purely on experimental data. Next, transfer learning from high fidelity simulations to experiments, then from low fidelity simulations to experiments, is tested. Finally, transfer learning from low to high fidelity, and then to experiments is performed to evaluate the benefits of hierarchical transfer learning. To compare the models, the mean and standard deviation of the explained variance score are computed. A Student’s t-test [30] is performed to determine if the transfer learned models are statistically significantly different in performance than a model trained exclusively on experiments. The explained variance scores and p-values from the Student’s t-test are summarized in Table II;]. The neural network hyper-parameters are noted in Table I and are kept the same for all of the models; these parameters are manually tuned to optimize performance. The choice to retrain the final layers of the model is motivated by applications of transfer learning to image recognition models; finer details are often the focus of

TABLE I
HYPER-PARAMETERS FOR ORIGINAL MODEL AND TRANSFER LEARNING FOR THE TAYLOR EXPANSION EXAMPLE.

Original Model Parameters	
Number of models	5
Hidden layer widths	4-8-14; 4-7-15, 4-8-14; 4-7-15; 4-7-14
Learning rate	0.004
Batch size	50
Epochs	300
Transfer Learning Parameters	
Retrained layers	Final 2
Learning rate	0.0001
Batch size	1
Epochs	300

the final layers, while the first layers in the networks focus on large scale features [12]–[14]. Similarly for regression tasks, modifying the final layers of the network allow for simple modifications to the network’s predictions, while modifying early layers could lead to substantial nonlinear changes to the overall response surface. This is a design choice that is used throughout the remainder of the manuscript, however it is not a necessary requirement; the layers chosen for retraining can also be optimized for a specific task.

For this simple example, there is not a statistically significant difference between a model trained exclusively on a large dataset of experiments, models that are trained on high fidelity simulations and calibrated to experiments, and models that are hierarchically calibrated. However, these three models are statistically significantly better than the model which is calibrated directly from low fidelity to the experiments, illustrating that there is an advantage of informing the model of high fidelity data prior to experimental calibration. In order to make an accurate emulator of the experiments, the cost of each of these routes should be the determining factor in which method to choose; however it is expected in most cases that the hierarchical method, which requires the least number of experiments and/or high fidelity simulations, is the least expensive approach. For this simple example the computational cost is negligible, but for applications in which complex multi-physics systems are being modeled, the computational cost difference between low and high fidelity simulations can be substantial; for such systems experiments are also often costly and limited in number.

In Table II only a single size dataset is considered for the hierarchical model. To determine the minimum number of high fidelity simulations and experiments that is needed to create a model that is not significantly different than the baseline (experiment only) surrogate, the mean explained variance score is computed as the dataset sizes are varied. First, the minimum number of high-fidelity points that are required to produce a transfer-learned model that is of similar performance to a model trained exclusively on 100 high fidelity simulations is determined. Next, the necessary number of experiments needed to re-calibrate this model to be on par with the experiment-only baseline model is determined. As in the models from Table II, the number of points in the dataset includes training and testing data, which are split into 80/20%

TABLE II
COMPARISON OF HIERARCHICAL AND ONE-STEP TRANSFER LEARNING (TL) TO DIRECT MODELING OF EXPERIMENTS. LOW/HIGH FI. INDICATES DATA PRODUCED WITH THE LOW/HIGH FIDELITY SIMULATIONS, RESPECTIVELY. EXP. INDICATES “EXPERIMENTAL” DATA PRODUCED WITH THE ANALYTIC EXPRESSION IN EQ. 1.

Model	Mean \pm SD Explained Variance	p-value
Train with 100 exp.	0.994 ± 0.004	-
Train with 100 high fi.; TL with 25 exp.	0.994 ± 0.006	0.957
Train with 100 low fi.; TL with 25 exp.	0.954 ± 0.041	0.016
Train with 100 low fi.; TL with 50 high fi. + TL with 25 exp.	0.981 ± 0.025	0.279




Fig. 2. “High fidelity” prediction quality as the number of high fidelity data points used for transfer learning from low fidelity data is increased. Transfer learned models require approximately 30 high fidelity data points to achieve a model accuracy that is comparable to a model trained exclusively on 100 high fidelity simulations.




Fig. 3. Experimental prediction accuracy as the number of experimental data points is increased in hierarchical transfer learning. Models are first trained on 100 low fidelity simulations, calibrated to high fidelity with 30 high fidelity data points, then subsequently calibrated to the experiments. The model quality converges with about 25 experiments, and is comparable to the performance of a model trained on 100 experiments alone.

sets. The transfer learning parameters are held constant as recorded in Table I, and each model starts with the same low fidelity surrogate trained on 80 data points and validated on the remaining 20 points.

Figure 2 illustrates how the transfer learning quality improves as the number of high fidelity data points is increased; the error bars reflect the variation in performance when the testing/training datasets are shuffled. Transfer learning with about 30 high fidelity samples produces models that perform similarly to models trained exclusively on high fidelity simulations. Thus the first step in the hierarchical calibration is performed with 30 high fidelity simulations. Next, the minimum number of experimental data points needed to calibrate this model to the experiments is determined. Figure 3 illustrates the quality of the transfer learned model predictions as the number of experiments is varied. For this example, performance comparable to the baseline is achieved with 25 experiments; beyond this the improvement in model quality is minimal.

Whether it is beneficial to perform hierarchical transfer

learning depends on the relative expense of the varying levels of fidelity in simulations and the experiments. If 100 low fidelity and 30 high fidelity simulations are less expensive than 100 high fidelity simulations, it is worth taking the hierarchical approach for subsequently calibrating to 25 experiments; however there may be situations in which running a high fidelity database is preferred. This example suggests that hierarchical transfer learning does offer improvements over transfer learning from low fidelity data directly to the experiments. The hierarchical calibration approach could also be improved by optimizing where the high fidelity simulations and experiments are placed in the design space; future work will explore sampling strategies for efficient hierarchical transfer learning.

The previous results are for a very simple, low dimensional function that is fast to evaluate. In the next sections, the same technique is applied to a real-world application: ICF experiments from the Omega laser facility.

III. TRANSFER LEARNING FOR ICF MODEL CALIBRATION

The performance of hierarchical transfer learning is next tested on a real-world dataset containing 23 experiments from the Omega laser facility that lie within a 9D design space. A database of 30k Latin hypercube sampled [31], [32] LILAC [33] simulations span the 9D space encompassing the experiments. The nine input parameters varied in the databases include laser pulse parameters: the average drive, drive rise time, energy on the target, the first picket power, foot power, foot width, and foot picket width, and capsule geometry parameters: the ice thickness and the outer radius of the capsule.

The 30k simulations are low fidelity; they are 1D, do not account of laser-plasma interactions (LPI) such as cross-beam energy transfer (CBET) [34]–[36], and use flux-limited thermal diffusion models [37]. Each simulation requires about ten wall-clock minutes to run on a single CPU. Each of the 23 experiments is accompanied by a high fidelity simulation, which is a 1D LILAC simulation with CBET, more accurate equations of state [38], and non-local transport; these simulations require approximately ten wall-clock hours to run. Both the low and high fidelity simulations produce a large number of scalar outputs; 19 of which are included in the following analysis. There are only 5 observables available for all 23 experiments that are common to the simulation database that will be used to test transfer learning with experimental data.

The low fidelity simulation-based neural networks are trained using DJINN, these models are referred to as “low fidelity DJINN” models. In section III-A, transfer learning is used to calibrate from low fidelity to high fidelity simulations, and then to experiments.

A. Hierarchical transfer learning with the Omega dataset

The hierarchical approach, in which the low fidelity model is calibrated first to high fidelity simulation data, then to the experimental data, is compared to transfer learning directly from low fidelity simulations to experiments for the Omega database. The models transfer learned to high fidelity simulations are referred to as “high fidelity DJINN” models, and those that are subsequently transfer learned to experiments as “experiment DJINN” models. If the high fidelity simulations are more accurate depictions of reality than the low fidelity simulations, priming the DJINN model with high fidelity information prior to calibrating to the experiments could improve the ability of the model to adapt to the experimental data.

An ensemble of five low-fidelity DJINN models is trained on the database of 30k low fidelity simulations. The models predict all 19 observables simultaneously and are individually cross-validated. The variance between DJINN models, each of which have been trained on a different random 80% subset of the data, will provide uncertainty estimates on the model predictions. The hyper-parameters used to train the networks are summarized in Table III. The mean explained variance score of the low fidelity models for each output is given in Table IV.

The same ensemble of low fidelity DJINN models are used to compare hierarchical calibration to calibration di-

TABLE III
HYPER-PARAMETERS FOR ORIGINAL MODEL AND TRANSFER LEARNING FOR THE OMEGA DATASET.

Original Model Parameters	
Number of models	5
Hidden layer widths	11-13-22-16; 11-14-19-26; 11-14-24-16; 11-15-29-30; 11-14-22-29
Learning rate	0.004
Batch size	1500
Epochs	400
Transfer Learning Parameters	
Retrained layers	Final 2
Learning rate	0.0003
Batch size	1
Epochs	2300

TABLE IV
MEAN EXPLAINED VARIANCE SCORES ON THE TEST DATASETS FOR FIVE DJINN MODELS TRAINED ON RANDOM 80% SUBSETS OF THE 30K LOW-FIDELITY LILAC SIMULATIONS.

Observable: Mean Explained Variance			
Absorption Fraction	0.991	R0	0.962
Adiabatic	0.886	RhoR	0.968
Burnwidth (BW)	0.869	Shock Mass	0.856
BangTime	0.990	RhoMaxBT	0.962
Convergence Inner	0.964	Tion	0.956
Convergence Outer	0.962	Tion_DD	0.959
In-flight aspect ratio (IFAR)	0.885	Vi	0.971
Peak Kinetic Energy	0.981	Yield	0.944
Pressure	0.958	Yield_DD	0.949
RhoNAve	0.967		

rectly from low fidelity simulations to experiments. The low fidelity DJINN models are calibrated independently, each on a different random subset of the high fidelity or experimental data. For each of the models, the first three layers of weights are frozen, and the remaining two layers are available for retraining, as shown in the cartoon in Figure 1. Note that the architecture of the networks is not reflected in this cartoon; the true architectures are given in Table III for the ensemble of five DJINN models. The last two layers of weights are retrained to convergence for 2000 epochs with a batch size of one and a learning rate of 0.0003 in the Adam optimizer (these parameters are tuned to optimize predictive performance of the models). Each model is trained on a random 90% of the experimental data (20 points) and tested on the remaining 10% (3 points). The low fidelity and post-shot simulations have 19 outputs, but the experiments only have 5 available observables. To calibrate to the experiments, the cost function, which is the MSE of the 19 scaled outputs, is modified such that the missing 14 outputs are not weighted. More explicitly, the cost becomes a weighted MSE where the weights are zero for outputs not measured in the experiment, and unity for those that are observed in the experiment. The predictions for the remaining 14 observables may change in non-physically motivated ways, thus we focus only on the 5 available observables. An equivalent approach to the weighted cost function would be to train the low fidelity and high fidelity DJINN models with only the 5 outputs available in the experiment. We choose to retain all 19 outputs so the accuracy of the high fidelity calibration for all 19 observables can be evaluated. Note that all input and output data is scaled [0,1], using the

parameter ranges set by the database of 30k simulations, prior to training. This prevents the cost function from being biased toward outputs that are larger in magnitude due to the choice of units.

First we consider transfer learning from the low fidelity simulations to the high fidelity simulations. Figure 4 illustrates the prediction error (calculated on training and testing data combined), computed as:

$$\text{Error} = \frac{(\text{Prediction}) - (\text{High fidelity truth})}{(\text{High fidelity truth})}, \quad (4)$$

for all nineteen available outputs for the uncalibrated (low fidelity) and calibrated (high fidelity) DJINN models. The error bars reflect the standard deviation in prediction errors from the ensemble of DJINN models; the points on Fig. 4 illustrate the mean error.

The low fidelity and high fidelity simulations differ significantly in their predictions of the nineteen observables, shown by the error in the blue points of Figure 4. The low error in the red points indicates that the networks are able to successfully learn the high fidelity outputs via transfer learning.

The high fidelity calibrated models are next calibrated to the experimental data, again by transfer learning the last two layers of the network – the same layers that were modified to calibrate to the high fidelity data. The mean and standard deviation of final prediction error is now computed using the experiment as the ground truth:

$$\text{Error} = \frac{(\text{Prediction}) - (\text{Experiment truth})}{(\text{Experiment truth})}, \quad (5)$$

The prediction quality is illustrated in Fig. 5 for the low fidelity, high fidelity, and experiment DJINN models.

Figure 5 illustrates that the high fidelity simulations are not necessarily more predictive of reality than the low fidelity simulations, however transfer learning to experiments is still successful. The largest error for transfer learning is in the yield, due to the model needing to adjust its predictions by over an order of magnitude for most experiments, however the mean prediction error for all observables is close to zero.

Since the post-shots are no closer to the experiments than the low fidelity models, hierarchical modeling does not offer significant benefits for this dataset; comparable results are achieved by transfer learning directly from low fidelity simulations to the experiments. Table V records the average explained variance ratio for each the experimental observables for the hierarchical models, and those that are calibrated directly from low fidelity to experiments. The explained variance ratios are computed on the test dataset, and are averaged for the five models.

The previous analyses involve randomly choosing the training and testing data for model calibration. To illustrate how these models can be used to predict the outcomes of future experiments, a model calibrated to the high fidelity simulations is calibrated to the experiments using only the oldest 19 experiments in the dataset. The model is tested on the 4 most recent experiments; the predictions are shown in Figure 6. Training on the old data and predicting the four most

TABLE V
EXPLAINED VARIANCE SCORES FOR MODELS CALIBRATED FROM LOW FIDELITY TO HIGH FIDELITY SIMULATIONS TO EXPERIMENTAL DATA, AND FOR MODELS CALIBRATED DIRECTLY FROM LOW FIDELITY SIMULATIONS TO EXPERIMENTS. THE HIGH FIDELITY SIMULATIONS ARE NOT AN ACCURATE PICTURE OF REALITY, AND THUS THERE ARE NO SIGNIFICANT BENEFITS OF FIRST CALIBRATING TO THE HIGH FIDELITY DATA FOR THIS PARTICULAR DATASET.

Observable	Mean±SD Explained Variance		
	Hierarchical	Low fi. - Exp.	p-value
Burnwidth	0.975 ± 0.023	0.889 ± 0.079	0.139
Bangtime	0.987 ± 0.015	0.942 ± 0.092	0.364
ρR	0.874 ± 0.097	0.835 ± 0.179	0.712
T_{ion}	0.988 ± 0.009	0.924 ± 0.094	0.211
Yield	0.818 ± 0.143	0.956 ± 0.034	0.096

recent experiments requires the model to extrapolate in input space, away from the old experimental data. The model is able to successfully predict the outcomes of the newest four experiments, demonstrating it does have the ability to successfully extrapolate away from the experimental data, but within the bounds of the simulations. It should be noted that the same hyper-parameters given in Table III are used when training on the old data and predicting the most recent experiments. Since these hyper-parameters were chosen by doing random selections of training/testing data, the more recent experiments could have implicitly influenced the choice of the hyper-parameters. A more fair comparison would have held the four most recent experiments from the choice of hyper-parameters as well, although it should be noted that the performance of the models is robust to minor perturbations in the learning rate and number of training epochs and it is unlikely that optimization of these parameters without the four most recent experiments would have resulted in significantly different model performance.

The low generalization error of the transfer learned models is a result of hyperparameter tuning that prevents over-fitting of “catastrophic forgetting” [39]. Over-fitting is when a machine learning model essentially memorizes the training data and cannot accurately predict data that is not included in the training set. Catastrophic forgetting is a related challenge that occurs in transfer learned models, in which the model is over-trained on the new task, and thus forgets all knowledge of the original task for which it was trained. In the case of the Omega dataset, the neural network would “forget” all knowledge of the simulation data and would simply learn how to fit the experimental data, and thus its extrapolation properties would not be informed by the simulation data, as desired. The risk of over-training to the experimental data can be reduced by limiting the learning rate, the number of training epochs, and the number of retrained weights in the transfer learning step. Limiting these hyper-parameters prevents the neural network from changing significantly, and thus it retains memory of the simulation data on which it was originally trained. High cross-validation scores and the ability to extrapolate away from the experimental data used in the training step indicates that the model is unlikely to be over-fitting to the experimental data. Although caution should still be taken when extrapolating away from the experiments, it appears that small extrapolations are tolerated, and the models can be updated after each




Fig. 4. Prediction error (with the high fidelity simulation as the ground truth) for calibrated and uncalibrated DJINN models. The low fidelity model predicts high fidelity observables with significant error, as the models contain different physics. The models calibrated to high fidelity data are able to predict all 19 high fidelity observables with low error.

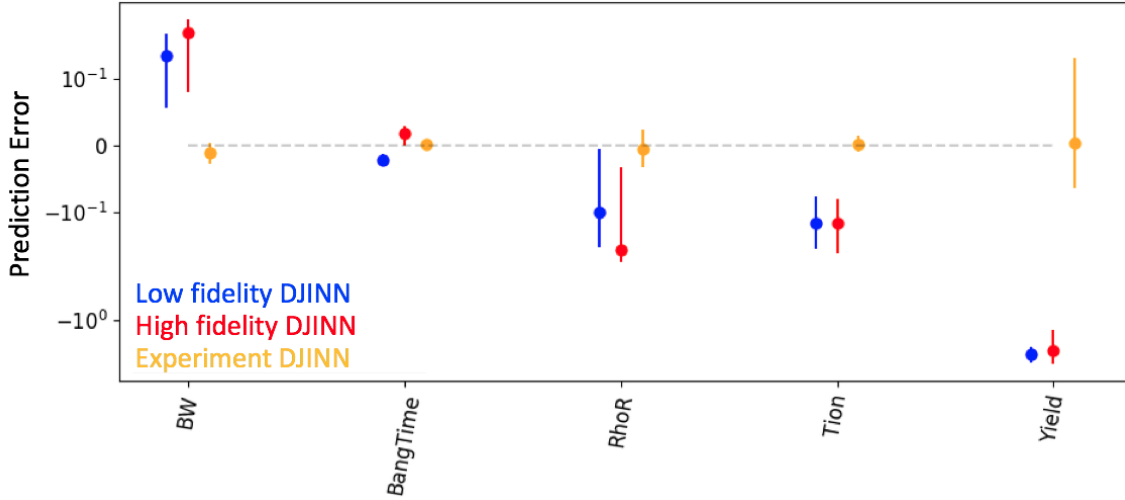


Fig. 5. Prediction error (with the experiment as the ground truth) for low fidelity DJINN models, DJINN models calibrated to the high fidelity data, and DJINN models calibrated to the experimental data. The experimentally calibrated models predict the five experimental observables with low error.

experiment to enable further exploration away from the current data.

IV. EXPLORING DISCREPANCIES BETWEEN SIMULATIONS AND OMEGA EXPERIMENTS WITH TRANSFER LEARNING

A result of hierarchical transfer learning is that it produces three sets of models which emulate low fidelity simulations, high fidelity simulations, and experiments. Studying the discrepancies between these models gives insight into the different physics included in the simulations and the experiment.

A primary use of ICF implosion simulations is to find optimal design settings for experiments. An interesting application of the three models is thus to search for “optimal” designs using each fidelity surrogate and determine if the simulation-based models suggest a similar “optimal” design as the experiment-informed model. For this exercise, we define an

optimal design as one that maximizes the experimental ignition threshold factor (ITFX) [40]:

$$\text{ITFX} \propto \text{Yield} \cdot (\rho R)^2. \quad (6)$$

where ρR is the areal density. The resulting optimal designs are illustrated in Fig. 7.

There are several important differences between the three optimal designs. First, consider the differences between the low fidelity (blue) and high fidelity (red) designs. The low fidelity design prefers high compression of a thick capsule, and achieves this by driving the capsule with a very high power. This differs from the high fidelity design, which prefers a lower power and thinner shell in order to achieve a similar implosion velocity and yield. The high fidelity design includes CBET effects, so it is reasonable for this design to lower




Fig. 6. Predictions of new experiments (yellow circled in bold black) after transfer learning using previous experimental data. The model is able to accurately predict new experiments, which are small extrapolations in input space from the old experimental data. The yellow data points indicate the training data predictions, the points with black outlines are predictions for the test data. The blue and red points indicate the low and high fidelity DJINN model predictions of the experiments, respectively.




Fig. 7. Designs which optimize ITFX according to the low fidelity, post-shot, and experiment DJINN models. The three designs are distinct due to the lack of accurate physics models, asymmetries, and other experimental sources of performance degradation not included in the simulations.

the peak power to avoid performance degradation from laser-plasma interactions [41].

Next consider the experimental design: unlike the high fidelity design, this design adjusts the picket and the foot of the pulse to guard against hydrodynamic instabilities that are not modeled in 1D simulations. Lowering the foot of the pulse (which occurs around 0.75 ns) lowers the adiabat inside the shell, allowing for higher compression and therefore higher areal density. To mitigate instabilities associated with higher compression, the picket of the pulse (occurring around 0.2 ns) is increased to increase the adiabat on the outer surface of the shell. The higher outer adiabat reduces hydrodynamic instabilities by increasing the ablation velocity [42]–[44]. The experimental capsule is thicker and is driven at an even lower power than the post shot for a longer period of time; perhaps due to underestimation of the CBET effects by the high fidelity simulations.

It is interesting to compare the predictions for each fidelity of DJINN model for the three optimal designs, given in Table VI. The low fidelity model optimal design is expected to perform well according to simulations, but the experiment-calibrated model expects it to perform 24% worse than the experimental optimal. Thus, relying on simulations leads to incorrect conclusions about how to best maximize implosion performance.

The maximum ITFX design according to the experiment-calibrated model is consistent with other analyses of this database [45]–[47]. The researchers at Omega are training power law-based models to relate simulation outputs and experimental measurements; through this process they found that to optimize yield, they should increase the thickness of the ice in the capsule; the experimentally-calibrated DJINN model makes a similar suggestions. They confirm their model predictions with a series of experiments, each time making

TABLE VI
 PREDICTIONS OF THE LOW FIDELITY (LOW FI.), HIGH FIDELITY (HIGH FI.), AND EXPERIMENT (EXP.) DJINN MODELS AT THE POINT OF OPTIMAL YIELD(ρR)² ACCORDING TO EACH OF THE THREE MODELS.

log10(Yield(ρR) ²) at:	Low Fi. Optimal	High Fi. Optimal	Exp. Optimal
Low Fi. DJINN	19.764	19.586	19.570
High Fi. DJINN	19.421	19.598	19.273
Exp. DJINN	18.951	19.019	19.070

small extrapolations in shell thickness and updating the models with the new experimental data before predicting the outcome of the next experiment. After maximizing the yield, the researchers performed a set of experiments to independently optimize the areal density by modifying the picket and foot of the pulse. This approach is largely physics-guided, and treats the pulse and capsule independently to optimize yield and areal density; it does not explicitly account for interactions between the capsule geometry and the laser pulse. In contrast to the power law fits to the data, the neural network models can consider nonlinear interactions between the inputs and outputs, tuning the pulse and capsule simultaneously to maximize ITEX. However, the neural networks might be inaccurate far from the experimental data, thus caution should be taken to make small extrapolations from the data with this technique as well.

The fact that two distinct methods for creating data-driven models suggest similar design choices for optimizing performance is encouraging, and illustrates the powerful role transfer learning can play in designing ICF experiments. Although a purely machine learning motivated ICF experiment has not yet been carried out, machine learning has been used to optimize performance of other large scale experiments with promising results [48], [49]. In future work, we intend to apply machine learning to guide the optimization of ICF experiments, however we feel it is still necessary to analyze such designs to ensure they do not encounter sources of performance degradation or physical limitations not included in the training data. For example, facility limits or safety requirements might be considered, or one might want to limit how far from previous experiments the model can extrapolate.

V. CONCLUSIONS

Transfer learning with deep jointly-informed neural networks has enabled the creation of surrogate models which emulate more expensive simulations and experiments, without requiring massive quantities of expensive data. Transfer learning uses low fidelity simulations to learn the approximate responses surfaces, then uses a sparse collection of expensive high fidelity simulations or experimental data to modify a limited number of weights in the network. The resulting networks accurately emulate the expensive data, without requiring a large database of high fidelity simulations or experiments to train the neural network from scratch. Hierarchical transfer learning, the process of calibrating from low to high fidelity simulations to experiments (or between levels of fidelity of simulations) enables the creation of accurate emulators for the

highest fidelity simulations or experiments, with lower computational cost than creating a neural network on the highest fidelity data alone. Transfer learning with DJINN enables the creation of neural network models that are predictive of direct drive ICF experiments at the Omega laser facility, which are used to design optimal implosions for future experiments.

VI. ACKNOWLEDGMENTS

The authors would like to thank Riccardo Betti, Varchas Gopalaswamy, and Jim Knauer at the Laboratory for Laser Energetics for providing the LILAC simulations and the Omega experimental databases, and for candid conversation. Thank you to Bogdan Kustowski for conversations on transfer learning.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Released as LLNL-CONF-764063. This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

REFERENCES

- [1] S. Atzeni, *The Physics of Inertial Fusion*. Oxford, UK: Oxford University Press, 2009.
- [2] J. Lindl, *Inertial Confinement Fusion: The Quest for Ignition and Energy Gain Using Indirect Drive*. College Park, MD: AIP-Press, 1998.
- [3] M. Marinak, G. Kerbel, N. Gentile, O. Jones, D. Munro, S. Pollaine, T. Dittrich, and S. Haan, "Three-dimensional HYDRA simulations of national ignition facility targets," *Physics of Plasmas*, vol. 8, no. 5, pp. 2275–2280, 2001.
- [4] R. Nora, J. L. Peterson, B. K. Spears, J. E. Field, and S. Brandon, "Ensemble simulations of inertial confinement fusion implosions," *Statistical Analysis and Data Mining*, 2016.

- [5] J. L. Peterson, K. D. Humbird, J. E. Field, S. T. Brandon, S. H. Langer, R. C. Nora, B. K. Spears, and P. T. Springer, "Zonal flow generation in inertial confinement fusion implosions," *Physics of Plasmas*, vol. 24, no. 3, p. 032702, 2017.
- [6] M. Van Oijen, J. Rougier, and R. Smith, "Bayesian calibration of process-based forest models: bridging the gap between models and data," *Tree Physiology*, vol. 25, no. 7, pp. 915–927, 2005.
- [7] H. Stripling, R. McClarren, C. Kuranz, M. Grosskopf, E. Rutter, and B. Torralva, "A calibration and data assimilation method using the Bayesian MARS emulator," *Annals of Nuclear Energy*, vol. 52, pp. 103–112, 2013.
- [8] A. Booth, R. Choudhary, and D. Spiegelhalter, "A hierarchical Bayesian framework for calibrating micro-level models with macro-level data," *Journal of Building Performance Simulation*, vol. 6, no. 4, pp. 293–318, 2013.
- [9] R. G. McClarren, "Uncertainty quantification and predictive computational science."
- [10] M. C. Kennedy and A. O'Hagan, "Bayesian calibration of computer models," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 63, no. 3, pp. 425–464, 2001.
- [11] J. A. Gaffney, "Making ICF Models More Predictive: Combining Simulations, Experiments and Expert Knowledge using Machine Learning and Bayesian Statistics," 2018.
- [12] S. Hoo-Chang, H. R. Roth, M. Gao, L. Lu, Z. Xu, I. Noguez, J. Yao, D. Mollura, and R. M. Summers, "Deep convolutional neural networks for computer-aided detection: Cnn architectures, dataset characteristics and transfer learning," *IEEE transactions on medical imaging*, vol. 35, no. 5, p. 1285, 2016.
- [13] M. Oquab, L. Bottou, I. Laptev, and J. Sivic, "Learning and transferring mid-level image representations using convolutional neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014, pp. 1717–1724.
- [14] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *Advances in neural information processing systems*, 2014, pp. 3320–3328.
- [15] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A Large-Scale Hierarchical Image Database," in *CVPR09*, 2009.
- [16] M. Everingham, S. A. Eslami, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes challenge: A retrospective," *International journal of computer vision*, vol. 111, no. 1, pp. 98–136, 2015.
- [17] G. Griffin, A. Holub, and P. Perona, "Caltech-256 object category dataset," 2007.
- [18] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [19] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1–9.
- [20] G. Miller, E. Moses, and C. Wuest, "The national ignition facility: enabling fusion ignition for the 21st century," *Nuclear Fusion*, vol. 44, no. 12, 2004.
- [21] L. Kegelmeyer, "Image analysis and machine learning for NIF optics inspection," Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States), Tech. Rep., 2018.
- [22] S. H. Langer, B. K. Spears, J. L. Peterson, J. E. Field, R. Nora, and S. Brandon, "A HYDRA UQ workflow for NIF ignition experiments," in *Proceedings of ISAV 2016: Second Workshop on In Situ Infrastructures for Enabling Extreme-scale Analysis and Visualization*. IEEE Computer Society, 2016.
- [23] J. Kelly, L. Waxer, V. Bagnoud, I. Begishev, J. Bromage, B. Kruschwitz, T. Kessler, S. Loucks, D. Maywar, R. McCrory *et al.*, "OMEGA EP: High-energy petawatt capability for the OMEGA laser facility," in *Journal de Physique IV (Proceedings)*, vol. 133. EDP sciences, 2006, pp. 75–80.
- [24] K. D. Humbird, J. L. Peterson, and R. G. McClarren, "Deep neural network initialization with decision trees," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 30, no. 5, pp. 1286–1295, May 2019.
- [25] V. Nair and G. E. Hinton, "Rectified linear units improve restricted Boltzmann machines," in *Proceedings of the 27th Intl. Conf. on Machine Learning (ICML-10)*, T. Joachims, Ed. Omnipress, 2010, pp. 807–814.
- [26] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," *ArXiv e-prints*, 2014.
- [27] M. Abadi and *et al.*, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015.
- [28] K. D. Humbird, J. L. Peterson, and R. G. McClarren, "Bayesian analysis with deep jointly-informed neural networks," in *submitted to Conference Proceedings for the Conference on Data Analytics*, 2018.
- [29] B. K. Spears and *et al.*, "Deep learning: A guide for practitioners in the physical sciences," *Physics of Plasmas*, vol. 25, no. 8, p. 080901, 2018.
- [30] M. O'Mahony, *Sensory evaluation of food: statistical methods and procedures*. Routledge, 2017.
- [31] R. C. W. McKay, M.D.; Beckman, *A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code*. Springer, 1979, vol. 21, no. 2.
- [32] T. J. Santner, B. J. Williams, and W. I. Notz, *The design and analysis of computer experiments*. Springer Science & Business Media, 2013.
- [33] J. Delettrez, R. Epstein, M. Richardson, P. Jaanimagi, and B. Henke, "Effect of laser illumination nonuniformity on the analysis of time-resolved x-ray measurements in uv spherical transport experiments," *Physical Review A*, vol. 36, no. 8, p. 3926, 1987.
- [34] I. Igumenshchev, D. Edgell, V. Goncharov, J. Delettrez, A. Maximov, J. Myatt, W. Seka, A. Shvydky, S. Skupsky, and C. Stoeckl, "Crossed-beam energy transfer in implosion experiments on omega," *Physics of Plasmas*, vol. 17, no. 12, p. 122708, 2010.
- [35] P. Michel, L. Divol, E. Williams, S. Weber, C. Thomas, D. Callahan, S. Haan, J. Salmonson, S. Dixit, D. Hinkel *et al.*, "Tuning the implosion symmetry of icf targets via controlled crossed-beam energy transfer," *Physical review letters*, vol. 102, no. 2, p. 025004, 2009.
- [36] D. Froula, T. Kessler, I. Igumenshchev, R. Betti, V. Goncharov, H. Huang, S. Hu, E. Hill, J. Kelly, D. Meyerhofer *et al.*, "Mitigation of cross-beam energy transfer: Implication of two-state focal zooming on omega," *Physics of Plasmas*, vol. 20, no. 8, p. 082704, 2013.
- [37] C. Levermore and G. Pomraning, "A flux-limited diffusion theory," *The Astrophysical Journal*, vol. 248, pp. 321–334, 1981.
- [38] H. Risken, "Fokker-planck equation," in *The Fokker-Planck Equation*. Springer, 1996, pp. 63–95.
- [39] I. J. Goodfellow, M. Mirza, D. Xiao, A. Courville, and Y. Bengio, "An empirical investigation of catastrophic forgetting in gradient-based neural networks," *arXiv preprint arXiv:1312.6211*, 2013.
- [40] P. Chang, R. Betti, B. Spears, K. Anderson, J. Edwards, M. Fatenejad, J. Lindl, R. McCrory, R. Nora, and D. Shvarts, "Generalized measurable ignition criterion for inertial confinement fusion," *Physical review letters*, vol. 104, no. 13, p. 135002, 2010.
- [41] B. J. MacGowan, B. Afeyan, C. Back, R. Berger, G. Bonnaud, M. Casanova, B. Cohen, D. Desenne, D. DuBois, A. Dulieu *et al.*, "Laser-plasma interactions in ignition-scale hohlraum plasmas," *Physics of Plasmas*, vol. 3, no. 5, pp. 2029–2040, 1996.
- [42] V. N. Goncharov, J. P. Knauer, P. W. McKenty, P. B. Radha, T. C. Sangster, S. Skupsky, R. Betti, R. L. McCrory, and D. D. Meyerhofer, "Improved performance of direct-drive inertial confinement fusion target designs with adiabat shaping using an intensity picket," *Physics of Plasmas*, vol. 10, no. 5, pp. 1906–1918, 2003.
- [43] J. P. Knauer, K. Anderson, R. Betti, T. J. B. Collins, V. N. Goncharov, P. W. McKenty, D. D. Meyerhofer, P. B. Radha, S. P. Regan, T. C. Sangster, V. A. Smalyuk, J. A. Frenje, C. K. Li, R. D. Petrasso, and F. H. Séguin, "Improved target stability using picket pulses to increase and shape the ablator adiabat," *Physics of Plasmas*, vol. 12, no. 5, pp. –, 2005.
- [44] K. Anderson and R. Betti, "Laser-induced adiabat shaping by relaxation in inertial fusion implosions," *Physics of Plasmas*, vol. 11, no. 1, pp. 5–8, 2004.
- [45] V. Gopalaswamy, R. Betti, J. Knauer, N. Luciani, D. Patel, K. Woo, A. Bose, I. Igumenshchev, E. Campbell, K. Anderson *et al.*, "Tripled yield in direct-drive laser fusion through statistical modelling," *Nature*, vol. 565, no. 7741, p. 581, 2019.
- [46] V. Gopalaswamy, "Optimization of direct-drive inertial fusion implosions through predictive statistical modeling," in *American Physical Society Division of Plasma Physics Meeting*, 2018.
- [47] R. Betti, "The one-dimensional cryogenic implosion campaign on OMEGA: Modeling, experiments, and a statistical approach to predict and understand direct-drive implosions," in *American Physical Society Division of Plasma Physics Meeting*. APS, 2017.
- [48] E. Baltz, E. Trask, M. Binderbauer, M. Dikovskiy, H. Gota, R. Mendoza, J. Platt, and P. Riley, "Achievement of sustained net plasma heating in a fusion experiment with the optometrist algorithm," *Scientific reports*, vol. 7, no. 1, p. 6425, 2017.
- [49] M. Streeter, S. Dann, J. Scott, C. Baird, C. Murphy, S. Eardley, R. Smith, S. Rozario, J.-N. Gruse, S. Mangles *et al.*, "Temporal feedback control of high-intensity laser pulses to optimize ultrafast heating of atomic clusters," *Applied Physics Letters*, vol. 112, no. 24, p. 244101, 2018.