

AI-Driven Casting Simulation for Faster Design Developments

Milan P. Raval
Altair Engineering, Troy, Michigan, USA

Copyright 2025 American Foundry Society

ABSTRACT

Casting simulation was a ground-breaking innovation for casting experts. When it was introduced in the early '70s, the industry felt relief as they were facing challenges like cost, trial and error methodology, and time consumption. Simulation technology solved most of these problems, but additional challenges arose such as the need for design engineers to perform multiple casting simulations for decision-making, and simulations that could take days to run on large models such as seen with megacastings. In both cases computational time became a bigger challenge. As the traditional trial and error method was replaced by casting simulation, here we will be talking about leveling up with the fundamentals of artificial intelligence (AI). Casting simulations that could take days to compute, AI integration can now predict in seconds or minutes.

Keywords: casting simulation, design, manufacturing, artificial intelligence, AI

INTRODUCTION

Comprehensive solutions can help organizations develop reliable and innovative casting components. Blending industry knowledge and innovative technologies, artificial intelligence (AI), and a generative design approach, validated by advanced simulation, organizations can deliver better products through multidisciplinary automotive design exploration. There are certain challenges in today's industries. These challenges are weight, performance, manufacturability, and sustainability. A product's manufacturability is an essential part of the design evaluation process. While high-pressure die casting (HPDC) is well-suited for high-quality parts, the manufacturability challenges posed by the complexity of new designs and finite-element analysis (FEA) with finer elements on intriguing locations are entirely new. Here, the feasibility of such a design component addresses the overall design and optimization process parameters to assure the quality of the component.

A solution for this is an AI-powered generative design workflow. This solution maximizes product efficiency while achieving performance targets and minimizing components' carbon footprint. For lightweight design, we can use topology optimization to generate the most efficient design alternatives and a response surface

modeling (RSM) optimization approach that uses AI and machine learning to classify and cluster results of a large multidisciplinary optimization study. By considering performance and manufacturability together early in the design process, we can assure the part's manufacturability, reduce casting defects, and analyze different product and process alternatives.

Artificial intelligence (AI) and machine learning help to make sense of the enormous amount of data generated by multi-objective optimization. This enables designers to understand the substantial number of simulations performed to optimize the design performance and its manufacturability while evaluating different design options very quickly. Thanks to machine learning and AI-driven expert emulation, designers can quickly identify the cluster of designs that offer the optimal balance of weight, safe space, energy absorption, and manufacturability.¹

CASTING MANUFACTURING

Casting is one of the most crucial manufacturing processes used by many industries around the globe. Its origin can be traced back to 4000 B.C. The definition of this process can be described as pouring or filling molten liquid metal into an empty cavity to obtain a desired shape upon solidification. Automotive, aerospace, machinery, defense, and household appliances are some of the industries that leverage this process. There are various manufacturing processes available to obtain desired products with different methods. Technical knowledge helps industry experts to understand the benefits and limitations of their chosen manufacturing process. Understanding the material property, process requirements, and the application of the product can become vital in defining a specific manufacturing process to work with. In other words, these manufacturing processes can transform low-efficacy products into high-utility products by altering some of their material properties.

A casting manufacturing process can be characterized further by gravity casting, high pressure, low pressure, permanent or semi-permanent mold casting. After part design, four major factors that allow industry experts to determine the correct process are geometry, material, surface finish, and cost. Designs with complex geometry, tight dimensional tolerances, and thin sections tend to require more costly processes. Designs with tight tolerances and minimal drafts may require special

molding processes.² In this case study, the author used real-life casting component for a casting simulation and an AI predictive study. The model assembly of an actuator housing was used as shown in Figure 1.




Figure 1. The model assembly of the actuator housing used in this study.³

ARTIFICIAL INTELLIGENCE

Artificial intelligence terminology can be easily understood by visualizing a 3-layered stack (Figure 2). In the first layer, all the fundamentals of data collection, gathering, analysis, and neural networks play crucial roles. The second layer is the outer shell of deep learning, also known as machine learning, which improves the performance based on exposed data. A program that leverages layers 2 and 3 to react and adapt for better outcomes is known as artificial intelligence.

Recent advancements in AI have made it possible to discover complex relationships from data. These advancements can generally be attributed to three factors: 1) algorithmic improvements to neural networks, 2) an increased availability of data, and 3) improvements to deep learning hardware, namely graphics processing units (GPUs). While neural network architectures for learning on images, video, audio, and text are ubiquitous, more recent architectures can learn on 3D shapes, sparking a new subfield known as Geometric Deep Learning.⁴

The ability to learn on 3D shapes has contributed to several fields, including geometry processing, computer vision, and robotics. Pertinent to this work, geometric deep learning has been applied to engineering simulation. The seminal works of Sanches-Gonzales et al.⁵ and Pfaff et al.⁶ demonstrated the use of geometric deep learning to map the relationship between domain geometry, boundary conditions, and physical behavior, essentially learning the solutions to partial differential equations (PDEs). Since then, the application of deep learning to PDE approximation has grown substantially.




Figure 2. The 3 Layers of Artificial Intelligence.⁷

SIMULATION SOFTWARE

Modern simulation software not only allows assessment of a part's performance in its operating environment but also of the manufacturing process employed (Figure 3). Historically, commercially available finite element analysis software was used by specialists to model manufacturing processes. Typically, these have been employed late in the design process to solve issues. More recently, process-specific software that is suitable for the non-specialist has become widely available. This allows simulation to be applied much earlier in the design cycle to not only refine a design for a given process but when combined with relevant cost models can also be used to select a preferred process.

Simulating part performance and the manufacturing process iteratively and early in the design cycle allows intelligent decisions and trade-offs to be made between a perfectly optimized part and a part that can be manufactured with the chosen process.² In this study, the author used a casting simulation tool for the high-pressure die casting process to simulate various designs and create datasets.




Figure 3. Prediction of casting manufacturing through simulation software.³

AI TECHNOLOGY

Today, a portfolio of AI tools is available in every sector. Using the right AI tool based on the requirements is a crucial step, for example, whether it is a research, manufacturing, or process-based prediction model. A commercial implementation of a geometric deep learning model was used to learn the relationship between mesh geometry and the porosity resulting from casting. The underlying neural network has tens of thousands of learnable parameters which are iteratively optimized during the training process until the prediction error. This is known as the loss, which is sufficiently small. Once trained, the network can predict the porosity of new designs based solely on their geometry.

This work demonstrates the use of geometric deep learning for approximating the solutions to casting simulations to accelerate the design process. The overall objective of the study was to re-create the fundamentals of casting manufacturing simulation and utilize the simulation datasets for geometric deep learning.

OBJECTIVES

1. Perform casting simulation for data generation;
2. Utilization of simulation datasets for pre-processing; and
3. Use of historical simulation datasets to predict new designs faster.

METHODOLOGY

A computer-generated case study was considered to generate, evaluate, train, and predict new designs for casting manufacturing. A practical use case of the high-pressure die-casting process for an actuator housing model was used to perform a design prediction and the porosity defect as a result type was considered to be the anchor point to carry out this study. This case study does not discuss replacing traditional casting simulation with AI tools but about leveraging them to create a workflow where new design decisions can be made with higher efficiency and speed. The methodology of this study focuses on these four areas:

1. Model preparation, simulation, and analysis;
2. Collection of historical simulation datasets;
3. Training machine-learning algorithms based on simulation data; and
4. Prediction of new designs faster with a machine-learning model.

Altair simulation tools Inspire™ Cast and physicsAI™ within the HyperWorks® platform were used to carry out this predictive analysis.⁹

MODEL PREPARATION, SIMULATION, AND ANALYSIS

The HPDC process involved a shot weight of 6.8 kilograms, casting alloy AlSi10Mg, casting temperature of 680C (1256F), mold material Steel H-11, and an initial mold temperature of 210C (410F). The first phase's velocity was 0.2 meters/s, and the second phase was 4.7 meters/s. For this case, cyclic simulations were performed to obtain an appropriate temperature distribution in the mold, followed by filling and then solidification simulations. The simulation model consisted of around 11 million tetrahedral elements, and an Intel® Core™ i9-9900K CPU @ 3.60GHz with 16 cores was deployed for the simulation.

The simulation results were analyzed in detail and compared with X-ray images from casting experiments (Figure 4). The simulations provided valuable insights into the melt flow, air entrapments, and potential shrinkage-related porosity – all of which were effectively addressed. The project compared simulation results with actual experimental data, and this comparative analysis deepened the team's understanding of the filling and solidification processes, enabling them to identify and mitigate defects. Casting simulation's advanced algorithms also enabled visualization of casting defects and optimized process parameters to achieve the desired product quality.³




Figure 4. Comparison between shrinkage porosity and X-rays.³

A single simulation to perform the casting analysis took around 20 hours for meshing, filling and solidification analysis. For any design engineers to perform such simulation on different design sets can lead to a time-consuming approach in the manufacturing cycle. Hence, this study using AI methodology was used to overcome the hurdle related to time consumption. Utilization of data analytics in manufacturing depends on the type and quality of data. However, this data is not only related to the designs and process of manufacturing. The dataset is also generated when manufacturing simulation is performed for early-stage analysis.

Preparing Various Designs for Data Generation:

As this study focuses on design changes and predictions, the process parameters were kept constant for all designs. Design variables are considered to be the change in thickness of two gates in the casting model, thickness of the center of the hub, and rib thickness of the actuator core. Design variable functionality in the casting simulation tool helped in creating different designs faster. Included were three major parameters: the thickness of gate 1, the thickness of gate 2, and hub section dimensions. It created seven variables and by adjusting these variable values, fifteen different design sets were generated as given in Figure 5:




Figure 5. Seven design variables of the actuator model.

To generate these datasets, each casting simulation was performed on a local desktop. The 5-step process of performing the casting simulation is described below:

1. Model/Assembly Import
2. Model Setup
3. Defining Casting Components
4. Process Parameters
5. Meshing Criteria and Simulation

Model/Assembly Import

The assembly of the actuator housing along with the runner system was imported into the user interface of the software. Simulation software allows users to import various types of extension files like computer-aided design (CAD) native formats, Parasolid, Iges, Step, etc. In this study, the Parasolid file format was used to import into the software environment.

Model Setup

Once the assembly was imported, step two required designating the casting part, runner system, and gate location. In this step, casting material and temperature were chosen. Gravity direction was also appointed correctly as it is crucial to define the right gravity direction based on a specific casting simulation method.

Defining Casting Components

The final step of model setup includes creating or designating any casting components within the software. For our actuator housing, a steel mold with a temperature of 210C (410F) was created. A shot sleeve representing the actual casting method was created.

Process Parameters

Step 4 is crucial in terms of specific process criteria. The software supports a variety of casting simulation processes like gravity casting, low-pressure die casting, and high-pressure die casting. In this study, high-pressure die casting was selected for the actuator housing similar to real-world manufacturing of the part (Figure 6). All crucial parameters of shot actuation during filling were defined.

Meshing Criteria and Simulation

After model setup, casting components designation, and process parameter selection, the last step of the simulation run is mesh creation and stage selection. The simulation software used for the study is based on finite element methodology. It creates a triangular mesh on the surface and a tetrahedral for volume. Approximately 11 million elements were created for the entire assembly. The simulation software used can mesh automatically by simply selecting the element size, making the workflow seamless for any design engineer without the need for any manual meshing.




Figure 6. High-pressure die casting simulation setup of actuator housing.

COLLECTION OF HISTORICAL SIMULATION DATASETS

The simulation results library consists of a unique file that stores information about design, meshing, setup, and results. This file can be used by AI tools to prepare and validate data for training machine-learning algorithms. As mentioned earlier, 15 datasets were created by running casting simulations that created necessary information.

These fifteen datasets can be used directly in the AI tool, or the user can perform pre-check and cleanup before importing in the AI user interface. This cleanup includes removing unnecessary casting components (like mold and runner system) from the file to only focus on the casting part results. Based on the target results type for the AI prediction, only the final time steps can be considered through this cleanup approach for results like filling, solidification, air entrapment, porosity, etc. For our case study, we used the last time step of the solidification to train porosity results.

TRAINING MACHINE-LEARNING ALGORITHMS BASED ON SIMULATION DATA

The author conducted the following procedure to load, organize, train, and validate the historical datasets generated through casting simulations on various designs. As mentioned earlier, relevant pre-processing and cleaning were performed before organizing this data into the AI tool. This 4-step procedure (Figure 7) is as follows:

1. Create and Load Project
2. Create and Manage Datasets
3. Training the Machine-Learning Model
4. Testing the Trained Model

Create and Load Project

The very first step in the AI prediction was to create a project session and save it to the desired location. This allowed the author to easily access the project multiple times with convenience.




Figure 7. The 4-step workflow from left to right for project creation, training, and testing.

Create and Manage Datasets

Datasets are collections of one or more simulation results used for training or scoring a model. The AI tool is compatible with many CAE file formats. It also supports all readable CAE element types including 0D, 1D, 2D, and 3D elements, and both first and second-order elements.

Both static and transient simulations are supported by the tool. If the results are transient, all samples must have the same number of time steps. Numerous types of physics simulation are supported, including structural, computational fluid dynamics (CFD), thermal, manufacturing, and so on. The tool expects that results are defined on all nodes/elements. If the result is undefined for some parts, they should first be removed during pre-processing before creating the dataset in the AI tool. The tool predicts only the von Mises component for tensor results like stress and strain. If another component is desired, like Tresca stress, a scalar field can be created from this component as a preprocessing step.

When shell elements have multiple layers, the tool predicts the maximum value across all layers. If the solid elements exist in the computer-aided engineering (CAE) data, selecting the Extract solid faces checkbox means the prediction will be made only on the outer surfaces. The Extract solid faces checkbox is selected by default because it ensures faster prediction time. If the Extract solid faces checkbox is cleared, predictions will be made for all nodes/elements throughout the interior of the solid body as well.⁸

With a single click on the plus icon over Datasets, historical simulation results were accessed. As shown in Figure 8, an interactive dialogue box pops up which helps to import and organize data as needed.

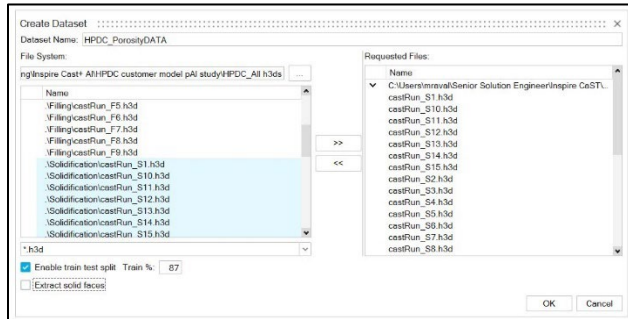


Figure 8. Simulation results import and organization settings.

The focus of this study was on porosity result type and solidification results for training. Two types of datasets can be created here. One is considered for training the machine learning model, and the other for testing the validation of the trained model. Users can manually divide the number of results files needed for training and testing or let the AI tool automatically divide them based on selected percentages. Considering 87% out of 15 datasets, 13 results were used to create training datasets, and 2 results were used as test datasets (Figure 9). Our goal was to focus on the porosity result type which is based on 3D tetrahedral volume mesh. Datasets were created disabling the Extract solid faces option.




Figure 9. Casting simulation of 15 datasets separated by test and train datasets.

Training the Machine-Learning Model

The AI tool allows users to train the datasets on the local machine, remote high performance computing (HPC), and GPU processing. Training times vary significantly depending on various factors, such as the number of samples, number of elements, time steps per sample, model size, hyperparameters, number of training epochs, and Hardware (processor speed, random-access memory (RAM), access to GPU). The AI software is based on geometric deep learning. Hence, it trains the model while considering the design and results irrelevant to mesh size, physics, and design variables.

In this particular case study, the author used a local machine with and without GPU processing to train the datasets as the number of samples was not too large. With a simple click on the model training plus icon, the tool opens up access to define training parameters as shown in Figure 10a.

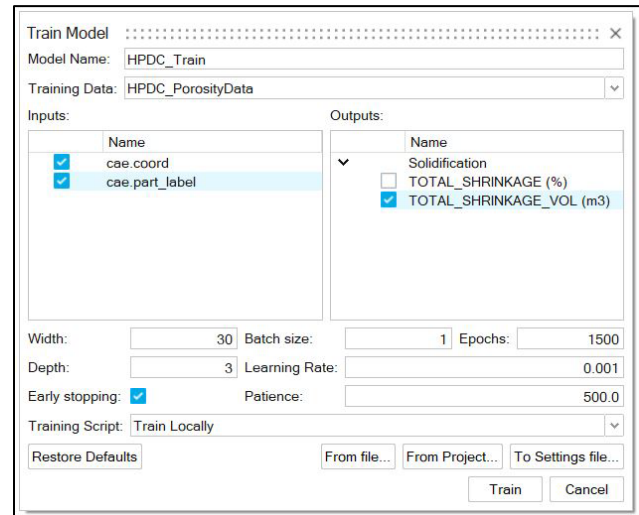


Figure 10a. Training the machine-learning model for porosity results.

Width

Width is the number of neurons per layer of the neural network. Larger widths result in neural networks with more learnable parameters and the ability to model more complex relationships.

Depth

Depth is the number of layers in the neural network. Larger depths result in neural networks with more learnable parameters and the ability to model more complex relationships.

Epochs

Epochs are the number of iterations for which training takes place. It is considered best practice to train until the training loss has converged.

Patience

Terminates the training early if the loss has not improved for this number of epochs. Early stopping saves time by ending the training early if the loss has converged before the total number of requested epochs.

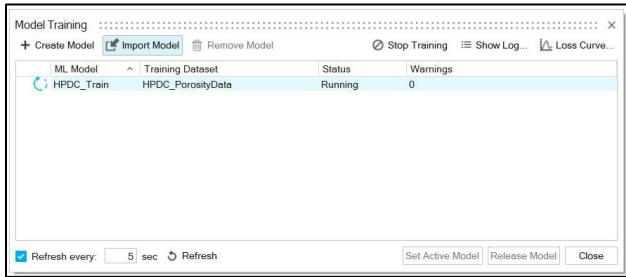


Figure 10b. Training the machine-learning model for porosity results.

Training configuration was run with default parameters suggested by the tool. Once the training is initiated, another dialogue box pops up showing the information on the ongoing training module as shown in Figure 10b. It also allowed us to access the loss curve to understand the model behavior. The loss curves are useful to visualize the progress of the training process.

In a well-fit model, the training and validation losses become nearly identical. If validation never approaches the training loss, this is indicative of underfitting; increased training time can leave files to improved model performance. A validation loss that approaches the training loss but diverges higher likely indicates overfitting; the point of low validation loss is the ideal model to avoid loss of generalization. Figure 11 shows an example of loss curve behavior.




Figure 11. Loss curve indicating training behavior over epochs.

Testing the Trained Model

Once the training on 13 results samples is done. It was validated against the 2-test sample to check the accuracy of the trained model. Here, 2 test data kept aside were not used in the training of the machine-learning model in the AI tool. The idea is to check and validate if the trained model shows higher accuracy in the prediction of the porosity result on test data.

The comparison was made between the AI tool predicting the porosity of 2 test samples against the simulation results (porosity) of these 2 test samples (Figure 12).




Figure 12. Testing designs 6 and 12 against the machine learning model.

Analyzing the results generated by the machine learning model on both designs 6 and 12 shows significant comparison. The porosity distribution is predicted correctly on both designs against their actual casting simulation. To verify the machine learning model's accuracy, the author investigated the MAE (Mean Absolute Error) factor, which appeared low in value, indicating a higher prediction accuracy. If the number of datasets and variables in designs is less, this accuracy can be achieved quickly. However, if accuracy is not achieved, then it means that either more data or a finer training configuration is required. To fine-tune training configuration, width, epochs, depth, and patience criteria can be manipulated as needed.

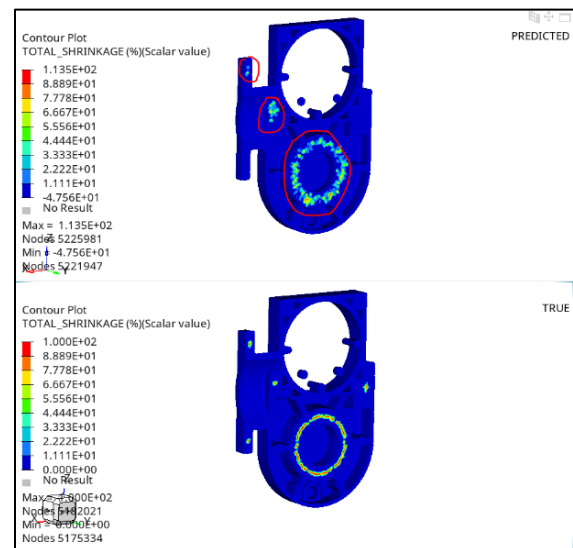


Figure 13a. Comparison of machine learning model prediction against actual simulation result for design 6.

Figure 13a shows the comparison of porosity results generated by the machine learning model and casting simulation results for design 6. The top half of the image shows the predicted porosity result which was not used for training the model. The bottom half of the image shows the real casting simulation result for porosity generated earlier as a part of historical datasets. Analyzing the porosity behavior and percentage of the porosity against each other shows significant similarities.

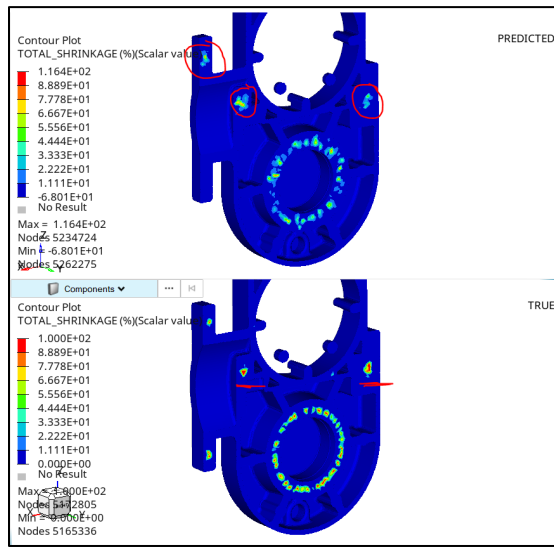


Figure 13b. Comparison of machine learning model prediction against actual simulation result for design 12.

Figure 13b shows the comparison of porosity results generated by the machine learning model and casting simulation results for design 12 as well. The top half of the image shows the predicted porosity result which was not used for training the model. The bottom half of the image shows the real casting simulation result for porosity generated earlier as a part of historical datasets. Analyzing the porosity behavior and percentage of the porosity for design 12 also shows major resemblances between them.

The author performed the testing not only with the default training configuration but also by adding more epochs iterations to improve the results further. There were two result types predicted for test data and validation. One was the porosity as a percentage, and the other one was the porosity as volumetric shrinkage. Total shrinkage by percentage shows the qualitative results, while total shrinkage volume shows quantitative results, indicating the amount of shrinkage in the part. This helps to compare the direct values of shrinkage amount between casting simulation data and results predicted by the machine learning model.

Once these outcomes satisfied the need for the case study, the author decided to predict the new design directly through the AI tool instead of running the casting simulation.

PREDICTION OF NEW DESIGNS FASTER WITH A MACHINE-LEARNING MODEL

Initially, the simulation datasets were created. Using the AI tool, data was imported and organized. The machine learning model was trained based on historical datasets. Validation of the machine learning model was performed for test datasets on design 6 and design 12. Now, it was time to consider the new design 16 (Figure 14), which was much different than the other design datasets for running the porosity analysis directly using the AI tool. Once the machine learning model is trained correctly based on historical datasets, there is no need to perform traditional casting simulations for analysis. With this AI prediction, porosity on design 16 will be performed in minutes to seconds. The author activated the machine-learning model before importing design 16. The new design can either be imported as CAD or as mesh files, such as meshed Finite Element Mesh (FEM) files to be predicted directly.




Figure 14. Design 16 for predicting porosity using the AI tool.

Design 16 consists of changes in the thickness of the gates, the thickness of the ribs, and the actuator hub diameter, considering that these variables in the new design are not similar to any of the previous 15 samples. As mentioned earlier, the author created two types of training models. One for porosity as a percentage and another for porosity in terms of shrinkage value. On the new design, we can predict porosity shrinkage volume directly by activating the relevant model in the AI tool. Once the new design as CAD or mesh was imported into the tool, a simple click on the predict icon quickly predicted the result in no time. Figure 15 shows the total shrinkage volume in the new design 16.

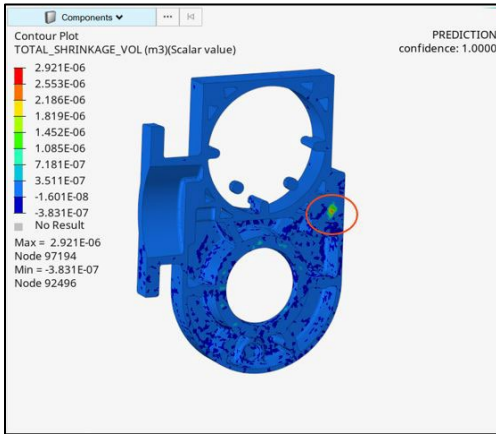


Figure 15. Porosity as a total shrinkage volume prediction on new design through AI tool.

Newly predicted results showed around $\sim 2.9 \times 10^{-6} \text{ m}^3$ volumetric shrinkage for design 16 to the certain location shown in Figure 15. Even though the training model was validated with a high confidence level and the least mean absolute error, the author was curious to check what type of porosity would be predicted if the traditional casting simulation was performed on the new design 16. (Figure 16) Hence, another casting simulation was executed. After analyzing porosity results in the casting simulation of design 16, the author was convinced of the accuracy of the AI prediction for new designs. Total shrinkage volume was predicted by the casting simulation in design 16 on the same location, similar to AI tool prediction. Also, the amount of volumetric shrinkage was $\sim 2.81 \times 10^{-6} \text{ m}^3$. Although the additional casting simulation was not necessary for design 16, it was still performed to confirm the accuracy of the machine-learning model in the AI tool. The difference in prediction between the casting simulation and AI was less than 4%. This difference can be reduced further by adding more datasets or optimizing the training criteria of the model.




Figure 16. Porosity as a total shrinkage volume prediction on new design through casting simulation.

COMPARISON AND DISCUSSION

This case study demonstrates convincing results by using AI technology for casting simulation. Differences between the predictions of results are negligible and could be reduced even further. While a traditional casting simulation takes specific time for model setup and computation, AI predictions are faster and can be performed with the speed of design. Once the machine-learning model is adequately trained, any new designs, later on, can be predicted quickly without the efforts of model setup and avoiding computational time. In this study, a single-casting simulation took around 15 hours for model setup and computation. Considering the generation of 15 designs, the cumulative time consumption for all iterations was approximately ~ 230 hours in total. Once the historical dataset was generated, the model setup, validation, and training time for the machine-learning model were ~ 17 hours on the local desktop through CPU and ~ 2 hours on the local desktop with GPU processing. After the model training, prediction on any new design took less than 1 minute (~ 39 seconds in this case).

ADVANTAGES OF AI TECHNOLOGY WITH CASTING SIMULATION

This method does not replace the need for casting simulation, but it leverages it through historical casting simulation performed on previous designs. It becomes more beneficial when the design changes are frequent and more complex, the size of the model is large with finer mesh, and there are multiple casting components. Large models developed using megacasting consist of millions of mesh elements that require multiple days to run casting simulations on local machines, or hours on high-performance computing. Contrary to that, AI tools will be able to predict new designs in a much faster way.

LIMITATIONS AND RECOMMENDATIONS

The author found that the AI tool used in this study is based on design and geometry only. It is independent of the process parameters for casting manufacturing. Optimization or changes in design along with process conditions cannot be considered with the AI tool used in this study. For process optimization or predicting the behavior of casting quality through process condition changes, alternative AI tools should be applied based on the requirement. As part of the implementation and feature improvement, process parameter changes would be considered for future research studies once such capability is available in the current AI tool.

ACKNOWLEDGMENTS

The authors thank Ms. Sarah Gauthier, and Mr. Eamon Whalen for their support as reviewers for this paper.

REFERENCES

1. “Guide To Megacasting: AI-Powered Generative Design,” Altair. <https://altair.com/resource/guide-to-megacasting-ai-powered-generative-design> (Link last accessed 02-11-25.)
2. M. Raval & A. Gray & J. Bangal, “Manufacturing Process Selection Based On Simulation-Driven Design” *AFS Metalcasting Conference* (2022).
3. Sydow Druckguss Gmbh, “Practical Use Case: High-Pressure Die Casting Of An Actuator Housing.” <https://altair.com/resource/practical-use-case-high-pressure-die-casting-of-an-actuator-housing> (Link last accessed 02-11-25.)
4. Bronstein, Michael M., et al., "Geometric Deep Learning: Going Beyond Euclidean Data," *IEEE Signal Processing Magazine*, 34.4 (2017): 18-42.
5. Sanchez-Gonzalez, Alvaro, et al. "Learning to Simulate Complex Physics with Graph Networks," *International Conference on Machine Learning*, PMLR (2020).
6. Pfaff, Tobias, et al. "Learning Mesh-based Simulation with Graph Networks," arXiv preprint, arXiv:2010.03409 (2020). <https://arxiv.org/abs/2010.03409> (Link last accessed 02-11-25.)
7. “3 Layers of Artificial Intelligence,” Shutterstock.
8. “Physics AI,” [Altair Online Training Materials](#). (Link last accessed 02-11-25.)
9. “Altair PhysicsAI Geometric Deep Learning,” <https://altair.com/physicsai> (Link last accessed 02-11-25.)