ARL-TR-10491 ● Nov 2026

A Transformer Model to Evaluate Wearable Blast Overpressure Sensor Data

by P. Justin McKee

DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

ARL-TR-10491 ● Nov 2026





A Transformer Model to Evaluate Wearable Blast Overpressure Sensor Data

P. Justin McKee
DEVCOM Army Research Laboratory

DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

REPORT DOCUMENTATION PAGE

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
		START DATE	END DATE
November 2026	Technical Report	10/1/2025	8/1/2026

4. TITLE AND SUBTITLE
A Transformer Model to Evaluate Wearable Blast Overpressure Sensor Data

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**

6. AUTHOR(S)
P. Justin McKee

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
DEVCOM Army Research Laboratory ATTN: FCDD-RLA-TB Aberdeen Proving Ground, MD 21005	ARL-TR-10491

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
Blast overpressure can produce concussive-like symptoms in service members in training and deployed environments. The U.S. Army is developing body-worn sensors for use by the Joint Conventional Force to measure, document, and improve the response to this overpressure. Measurements from pressure sensors are heavily influenced by reflections that result from body positioning and do not directly provide consistent metrics that can be compared between exposures. An ML transformer model has been developed to read pressure history data from multiple sensors and extract information from that combined input. Models have been developed to estimate incident overpressure and the direction to the blast and make use of techniques to improve training efficiency by using low-fidelity simulations to supplement high-fidelity experiments.

15. SUBJECT TERMS
Terminal Effects, blast overpressure, machine learning, time series extrinsic regression, multi-fidelity

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT	b. ABSTRACT	C. THIS PAGE		
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UU	28

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
P. Justin McKee	(520) 691-4788

STANDARD FORM 298 (REV. 5/2020)
Prescribed by ANSI Std. Z39.18

Contents

List of Figures iv

List of Tables v

1. Introduction 1

2. Methods 2

2.1 Datasets 2

2.1.1 Low-Fidelity Data 2

2.1.1 High-Fidelity Data 4

2.2 Data Preparation 5

2.3 Model Architecture 5

2.3.1 Low-Fidelity Model 5

2.3.2 High-Fidelity Model 8

3. Results 9

3.1 Attention Mapping 9

3.2 Low-Fidelity Results 10

3.3 High-Fidelity Results 14

4. Discussion 16

5. Conclusion 18

6. References 19

List of Symbols, Abbreviations, and Acronyms 20

Distribution List 21

List of Figures

Figure 1. Example BlastX setup showing 8 boxes that represent mannequins with 20 sensors (blue dots) per box. Boxes are centered around a blast source (red sphere). A ground surface is simulated but hidden in the figure. 3

Figure 2. Example pressure plots generated from BlastX for sensors on three different faces of a box. 4

Figure 3. An example of pressure data from three sensors formatted for input into the model. Colors indicate different patches used as tokens to be processed by the transformer model. Dotted lines show the start and end of senor data and contain a special token to signify the start of new sensor data..... 6

Figure 4. Diagram of the multi-fidelity model that uses the low-fidelity model prediction and high-fidelity peak pressure data to predict high-fidelity IOP. 9

Figure 5. An example of normalized encoder layer attention weights for key and query pairs based on four pressure sensors. Circular plots show the input plot starting at the top and circling clockwise with lines connecting the top 2.5% of connections. The square plots show the full set of data in matrix format. 10

Figure 6. Comparison of actual simulated IOP to model predictions of IOP when running the model with two to five input sensors. Markers for experimental positions are low-fidelity simulations in the same configuration as available experimental data..... 11

Figure 7. A comparison of the target IOP to model results at locations where experimental data is available for different charge sizes. The target IOP is the low-fidelity-simulated IOP and not the experimentally measured IOP. 12

Figure 8. A comparison between simulated box positions and the position from the model prediction. Markers for experimental positions are low-fidelity simulations in the same configuration as available experimental data. 13

Figure 9. A comparison between low-fidelity model prediction of direction to the blast source when run with low-fidelity and high-fidelity data. Colors indicate the actual position, and the marker is plotted on the predicted position. Predictions from low-fidelity data are plotted on the outer ring with a solid line indicating the average result. Predictions from high-fidelity data are plotted on the inner ring with dashed lines to indicate the average result..... 14

Figure 10. Results of the multi-fidelity model for each charge size. Horizontal dotted lines show experimental IOP ± 1 StdDev. Gray dashes indicate a single model prediction. Vertical lines show average model IOP ± 1

StdDev. Markers are offset from their exact position degree to visually separate results for each charge size. 15

Figure 11. A visual representation of results for prediction of direction. Training data is plotted on the outer ring with a solid line to indicate the average value. Test data is plotted on the inner ring with a dashed line to indicate the average value. 16

List of Tables

Table 1. Hyperparameters used to define and train the model..... 7

Table 2. Results of the multi-fidelity IOP model compared with experiments for three charge sizes. Results for test results are in bold..... 14

Table 3. Results from fine-tuning the low-fidelity model on high-fidelity data to predict the direction to the blast. Test results are in bold. 15

1. Introduction

Understanding the effects of primary blast overpressure (BOP) for both short-term injury and long-term health is a unique challenge faced by the U.S. Army. Soldiers are exposed to low-level blasts during training and active duty while performing routine operation of weapon systems such as artillery, shoulder-fired weapons, and high-caliber rifles. The overpressure from operating these weapon systems is relatively low, and operating guidelines are in place to reduce exposure risk, but there is still a need to track and understand blast exposure. Creating a record of blast exposure helps to enforce training guidelines and validate that they are sufficient when comparing exposure history against long-term health outcomes.

Wearable BOP sensors (WBOPS) are instruments that are being developed with the goal of measuring and tracking exposure to BOP. However, accurate and consistent tracking remains a significant challenge because obscurations and reflections from surfaces heavily influence measurements by a sensor. The incident overpressure (IOP) is the overpressure with no influence from a reflecting surface. It can be measured most accurately with a sensor oriented to face perpendicular to the shock velocity. Pressure measurements increase as the sensor is rotated to face into the shock velocity because of reflection from the sensor surface. The value referred to as reflected pressure is measured when the sensor is facing directly toward the blast source, parallel with the shock velocity. IOP and reflected pressure values are measured in laboratory experiments with large instruments that must be carefully positioned and calibrated. WBOPS orientation changes with posture and movement of the wearer causing them to measure some intermediate value between the IOP and reflected pressure. Directly using pressure values from WBOPS is not ideal for tracking exposure because an identical blast event will produce different measurements based on the Soldier's posture during the event. Additional post processing is required to determine a consistent metric for blast exposure tracking. There is an additional need to also interpret the effects on the body under the armor that adds a second layer of complexity.

ML provides methods to process and use complex data such as the pressure history recorded by WBOPS. In this report, we develop an ML transformer model as well as a multi-fidelity model to process WBOPS data to achieve two goals. First, we develop methodology to process the whole pressure history with flexibility to handle variable length data from a variable number of sensors. Flexibility to read variable length data is important to capture blasts of different durations and allows the model to be run even in cases in which some sensors did not record the event. Second, we develop an approach to make use of data from multiple fidelity levels to improve training of the ML model. Blast experiments

are expensive and time consuming, which limits the quantity of high-fidelity experimental data. Simulations provide an alternative data source, but high-fidelity simulations are also limited by computation time and resources. Low-fidelity simulations provide a greater quantity of data, but methods must be developed to handle differences in structure and account for physics that is not captured by a low-fidelity simulation. Both a multi-fidelity model approach and fine-tuning are explored to demonstrate how a low-fidelity data source can improve efficiency of training when high-fidelity data is limited. IOP and the direction to the blast source are used as metrics to evaluate performance of the ML model.

2. Methods

The goal of this research is to develop methods to process pressure data from multiple WBOPS and make predictions of metrics that cannot easily be manually extracted. IOP and the direction to the blast source are values that are available from simulations and experiments that are used to demonstrate the methods developed here. IOP represents a single-pressure value to relate all sensor pressures measurements. The direction to the blast source is a more complex output and an example of using pressure to predict a metric other than pressure. Two separate ML models are developed for IOP and direction using data from low-fidelity simulations for initial training and high-fidelity experiments to adapt the low-fidelity model to better match realistic data.

2.1 Datasets

2.1.1 Low-Fidelity Data

A low-fidelity dataset is created by running simulations of primary blast in BlastX[1] software. BlastX is a fast-running simulation code that is intended to simulate primary blast with interactions and reflections from structures. A low-fidelity approximation of experiments is made by creating a set of box structures with dimensions that roughly match mannequins that are used in experiments. Sensors are placed on the surfaces of the box to represent WBOPS. Simulations are run with calculations for shock and gas at 101.3-kPa atmospheric pressure and 15° C. Boxes are set to the concrete wall type. These simulations lack important geometric features and realistic material models, but they provide an initial low-fidelity source of pressure data from sensors placed on different surfaces on an object.

An example of the BlastX simulation setup is shown in Figure 1. Each simulation consists of eight 0.4-m-wide × 0.24-m-thick × 0.8-m-tall boxes. Each box has six

sensors on the front and back face, three sensors on the sides, and two sensors on the top face. Sensors are arranged in a grid for half of the boxes and randomly distributed on the faces of the other half of boxes. The random sensors are restricted to maintain a spacing of at least 3 cm from the edge of the box. The motivation for the random sensor placement is to provide an extra source of uncertainty when training the model. Each box is evenly spaced around a central charge with 45° between each box. A ground surface is present in the model to allow ground reflections. A set of models are created to vary the standoff to box center between 1.25 and 12 m, height of box base between 0.4 and 1.9 m, charge height between 0.005 and 3 m, and charge mass between 3 and 200 g composition C-4. The radial position of the first box is varied in simulations between 0° and 45° so that any location between 0° and 360° is possible over the full eight-box set. A custom Python script is used to define the geometry of the boxes, sensor positions, charge position, and then run simulations in batches to generate a large dataset.

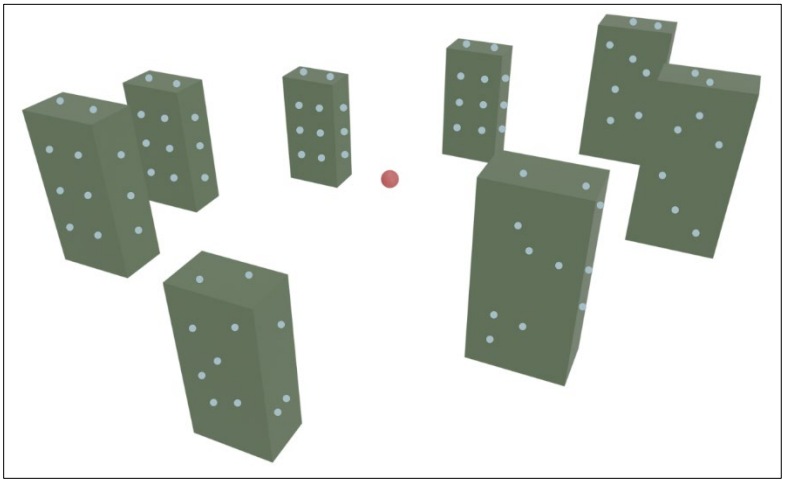

Figure 1. Example BlastX setup showing 8 boxes that represent mannequins with 20 sensors (blue dots) per box. Boxes are centered around a blast source (red sphere). A ground surface is simulated but hidden in the figure.

An example of BlastX pressure data for sensors on the front, right, and top of the same box is shown in Figure 2. Pressure data roughly captures the expected Friedlander waveform along with reflections from the ground and neighboring boxes. Simulation data is resampled to a frequency of 0.1 MHz, and zero padding values are inserted before the peak. Resampling is performed starting at the peak pressure value to ensure that it is not changed. The dataset is filtered to remove cases with pressure above 689.5 kPa that are outside of the range of interests and to remove errors where no pressure was recorded or the initial spike to peak pressure was not found. A total of 22,235 boxes with valid data make up the dataset that is used for training and evaluation. An additional 24 boxes that match experimental conditions were also simulated and added to an evaluation dataset.

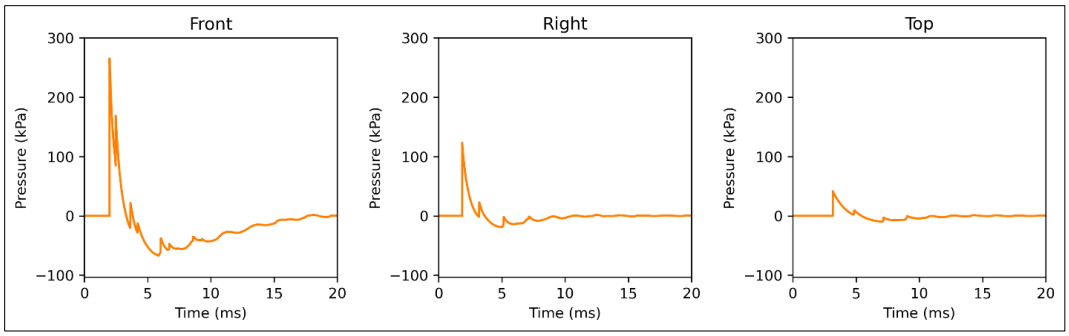

Figure 2. Example pressure plots generated from BlastX for sensors on three different faces of a box.

The target IOP used to train the low-fidelity model is also defined by BlastX simulations. Simulations are run covering the same range of charge sizes with a line of sensors in open air. The ground is not included in these simulations. Interpolation is used to obtain IOP for conditions between those that were directly simulated using a distance from the blast equal to the average distance of all sensors used in the prediction.

2.1.1 High-Fidelity Data

High-fidelity data comes from arena blast experiments described by Dean et al.[2] where mannequins with WBOPS are placed around an explosive charge. Data from these experiments include pressure history data from WBOPS and corresponding laboratory-quality pencil and plate sensors to measure IOP and reflected pressure, respectively. Details on the WBOPS and their data are not included in this report.

Experiments include mannequins placed from 0° to 180° with WBOPS on both sides of the body or on the center of the body. Experimental data is compared with simulations over the full 360° by flipping the left and right sensors. For example, a sensor on the left side of the body at 45° is equivalent to a sensor on the right side of the body at 315°.

Experimental data is augmented to increase the size and variability of training data. The average and standard deviation (StdDev) of IOP recorded from pencil gauges are used to create a normal distribution that is used to set target IOP values during training. Sensor data from different experiments with the same charge mass and position are mixed to create a set of input data for the model. This approach allows for a larger combination of inputs for the model under the assumption that results are reasonably consistent between repetitions of an experiment. High-fidelity data is resampled to 0.1 MHz to match the low-fidelity data. Resampling is performed

starting at the peak pressure to make sure the peak value remains in the resampled data.

2.2 Data Preparation

Pressure data requires additional processing before use in an ML model. First, the models developed in this report do not make use of the time component of the data. Instead, all input pressure must be resampled to a constant sampling rate for all data sources. Next, a log normalization method is used to increase sensitivity for low-pressure values. This normalization method is useful to make the model more sensitive to low-level blast that is of primary interest here while also still processing significantly higher overpressure data. Pressure data is first transformed using Equation 1, where p is the pressure, p^* is the transformed pressure, c is the pressure value that will be set to zero after being transformed, and s is the steepness of the transformed output. The transformed pressure values are then normalized to values between −1 and 1. Models shown in this report use c set to 0 kPa and s set to 0.7.

$$p^* = sign(p)\, log\big((s\,|p - c|) + 1\big) \tag{1}$$

Next, the window of data provided to the model is adjusted. The amount of data before the peak pressure is particularly important because the timing of events recorded by WBOPS may not be reliable enough to estimate the relative time of arrival between sensors. A random offset in the data window is used for training so the model will not be able to learn to make use of the arrival time.

2.3 Model Architecture

2.3.1 Low-Fidelity Model

A transformer architecture[3] is applied to the low-fidelity data source. Transformer models excel at processing sequential data by using an attention mechanism that allows the model to learn relations between any sections of the data regardless of their distance within the data. Transformer models are well known for natural language processing (NLP) and large language models, but they have also been applied to image data for object detection[4,5] and time series data.[6] A transformer model is chosen here for its ability to process and link long-range connections within the data such as the locations of peak pressure and positive duration between multiple sensors. The attention mechanism can also be used as a mask to allow the model to process data of variable lengths more effectively compared with other ML models with more strict constraints on data input size. This flexibility is useful in processing pressure data where it is not always easy to crop, resample, and resize data to fit model size constraints. This flexibility also allows the same model to be

used for a variable number of sensors, which could be helpful when fine tuning for different systems or allowing the model to run even if one of the sensors fails to record an event.

A transformer model processes data as a set of token embeddings. Pressure plot data is converted to tokens using an approach comparable to vision transformer models[4] where the data is split into patches of a constant size. The patches serve as a token, and a linear network with no activation function is used on the patch to create an embedding with a size to match the dimensions of the model. Position encoding values are created using methods from NLP[3] where a sinusoidal function creates values that are added to the embeddings to enable the model to learn relative positions.

Data from multiple pressure sensors are combined by concatenating them together into a single set of values. A special token inserted between the plots using a patch where all values in the patch are set to one after normalization so that the model can learn to separate data when multiple plots are provided. An example of pressure from three sensors is shown in Figure 3 where different colors indicate patches that will be converted to tokens. The narrow bands marked by vertical dotted lines between sensors contain the special token that separates plots with a pressure value that will become equal to one after the data is normalized. The model has a fixed context length, which is the maximum number of tokens that can be processed by the model. The masked data section in Figure 3 shows unused tokens where attention masking will cause the model to ignore them and focus on tokens that contain pressure data.

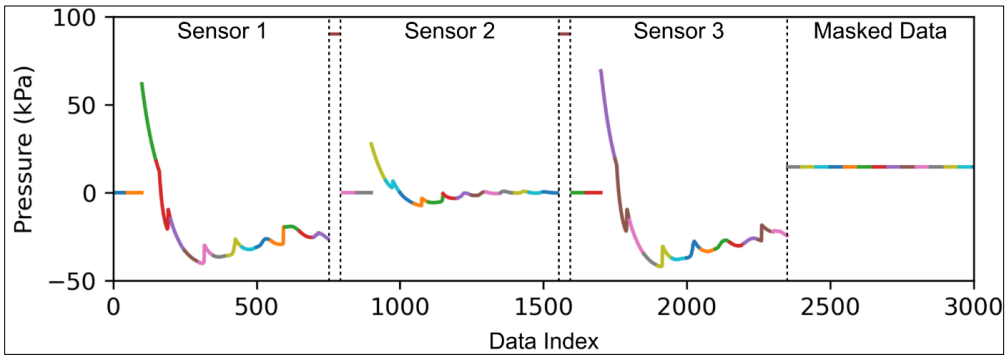

Figure 3. An example of pressure data from three sensors formatted for input into the model. Colors indicate different patches used as tokens to be processed by the transformer model. Dotted lines show the start and end of senor data and contain a special token to signify the start of new sensor data

The model is composed of a set of encoder and decoder layers as described by Vaswani et al.[3] To summarize, encoder layers are composed of a multi-head attention (MHA) block followed by a normalization layer, a feed forward network,

and a second normalization layer. The decoder layers make use of a similar structure with a second MHA layer that uses the output of the encoder as the values and keys and the output of the first decoder MHA layer as the query. This is followed by a feed-forward layer. Normalization is applied after each MHA and the feed-forward layer. The decoder model uses an initial input to generate a new token embedding. The initial input into the decoder model is an embedding with trained parameters based on methods used in vision transformers.[4] Finally, a multilayer perceptron (MLP) network is used as an output head to predict a specific metric from the decoder embedding output. Although it is possible for one model to output multiple different metrics, two models to predict IOP and direction to the blast are trained separately in this report.

The PyTorch package is used to construct and train the model. Hyperparameters used to define the current model are listed in Table 1. We found that a relatively small model dimension was able to converge more effectively on the available dataset. A cosine warm-up scheduler is used to linearly increase the learning rate up to the maximum level and then decrease it with a cosine function over the remainder of training. The AdamW[7] optimizer is used to optimize model parameters using gradient clipping and weight decay.

Table 1. Hyperparameters used to define and train the model.

Model		Data		Training	
Patch size	50	Number sensors	2–4 (IOP)	Max. learning rate	5×10^{-4}
Max. number tokens	83	Tokens per sensor	8–20	Clip gradient norm	0.1
Dimension model	128	Tokens before peak	1–3	Weight decay	1×10^{-6}
Dimension feed forward	512	Sample rate (MHz)	0.1		
Attention heads	8				
Encoder layers	6				
Decoder layers	6				

The IOP model is trained with two to four random sensors from unique faces of the box in a random order. These random selections make the problem more challenging so the trained model will be more robust to variability. This is important in real-world applications in which posture and proximity to other objects that can cause reflections that will greatly increase variability compared with simulations or controlled experiments. In contrast, the direction model is trained with a set of three sensors that are selected from the front, right, and top faces in the same order. The choice of sensor from each face is still random. This approach is used because the position of sensors is critical in estimating the direction to the blast. There is currently no position information encoded into the model input, so it must be inferred from the ordering of inputs. The output of the direction model is a 3D-unit vector that points toward the blast source. The direction is simplified

by setting the height component to zero and renormalizing it to convert to 2D and an angle for comparison to experimental mannequins' positions.

2.3.2 High-Fidelity Model

Several options are available to use a model trained on a low-fidelity data source to enhance development of a model for a high-fidelity data source. Typical approaches include fine tuning a model where all, or subsets, of a model's parameters are retrained using the new data source. Transfer learning freezes parameters of the original model and appends a new output head that is trained on the new data. Multi-fidelity approaches train a new model to predict a correlation between fidelity levels that can be used to transform results rather than directly predicting a high-fidelity output. This multi-fidelity approach is effective when the relationship between fidelity levels is easier to learn than directly learning to predict a metric from the high-fidelity data.

Attempts to fine-tune the model or use transfer learning for the IOP model were unsuccessful because the model very quickly overfits the experimental dataset. The experimental dataset used in this report only contains results for three charge sizes. This limitation can cause the problem to effectively become a classification problem in which the model learns to only output the few possible outputs from the training data and performed poorly on test data. The IOP model instead makes use of a multi-fidelity approach like the method described by Meng and Karniadakis.[8] A diagram of the model is shown in Figure 4. First, the pressure data is run through the low-fidelity model to produce an initial prediction of IOP. Next, linear and nonlinear models are trained to output parameters that will be summed with the low-fidelity IOP prediction. The linear and nonlinear models use the low-fidelity model's IOP prediction as well as the peak pressure of each high-fidelity sensor as input. The linear model is a single, fully connected layer with no activation function. The nonlinear model is an MLP model with hidden layers of size 512, 512, 256, and 64. Gaussian error linear unit activation is used for all except the last layer of the model that uses a hyperbolic tangent activation function.

Fine-tuning was found to be effective for the direction model. Fine-tuning was performed over 20 epochs of high-fidelity training data with a batch size of 256. A cosine warmup schedular was used to increase to a maximum learning rate of 0.0005 at approximately 5% of the training set before lowering it. Other training parameters were set to the same values used during the initial training. The multi-fidelity approach was also attempted on the direction model, but it produced poor results compared with fine-tuning.

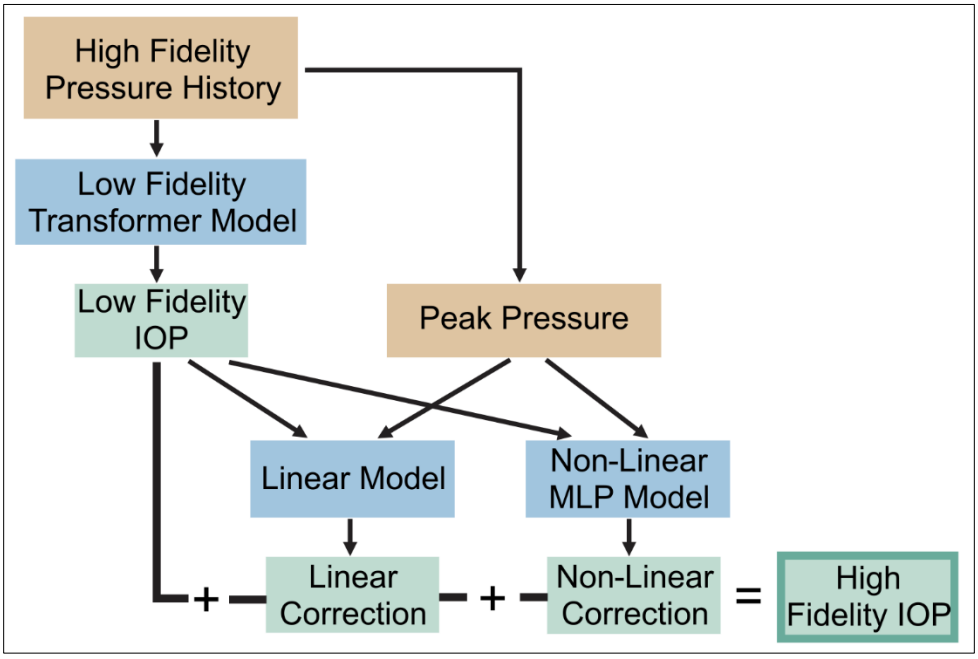

Figure 4. Diagram of the multi-fidelity model that uses the low-fidelity model prediction and high-fidelity peak pressure data to predict high-fidelity IOP.

3. Results

3.1 Attention Mapping

An example of the IOP model's attention weights from encoding an example pressure input with four sensors is shown in Figure 5. This example illustrates how attention is distributed to focus on different features of the input at each layer of the encoder. Each attention map shows the relationship between key and query pairs where the weight represents how important the key token is in understanding the query token. The attention weights are plotted on a circular diagram with lines that link the top 2.5% of connections as well as a square matrix that shows the same data for all pairs. The color map on all plots shows the attention weight normalized for that layer. There are several features of these attention maps that help to provide a more intuitive understanding of the relationships the model is using when making predictions of IOP. For example, layer 1 has its strongest connections at the bottom of the negative phase of the blast extending through the negative phase for the two largest peaks that have a larger negative phase. Layer 1 also shows attention at neighboring tokens throughout the plot. Layer 2 shows stronger connection between the positive duration of the pressure to tokens in the surrounding area. The two larger peaks have stronger connections to the full dataset compared to the smaller peaks. Layers 3 and 5 show connections with the special tokens that

indicate the start of a new plot. Layer 4 is focused on the peak pressure for each plot, whereas layer 6 is focused primarily on the peak pressure of the two largest peaks.

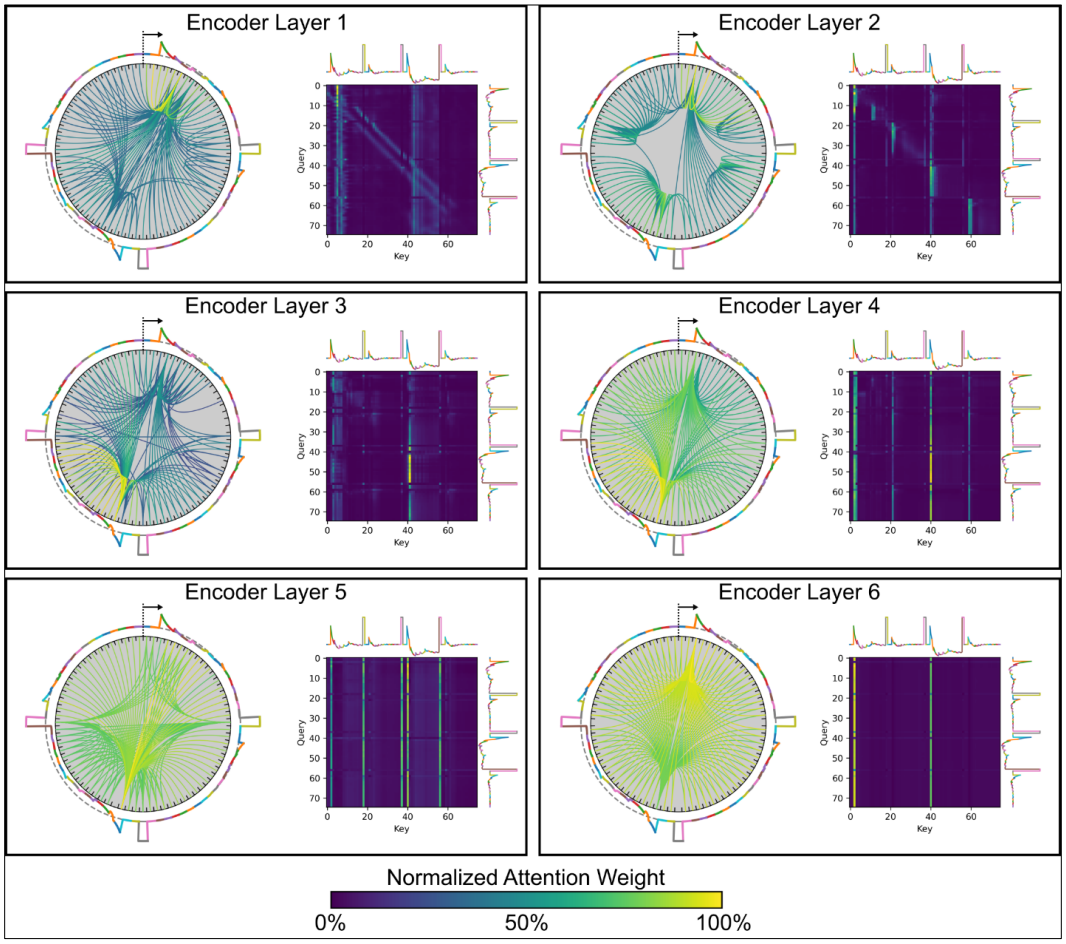

Figure 5. An example of normalized encoder layer attention weights for key and query pairs based on four pressure sensors. Circular plots show the input plot starting at the top and circling clockwise with lines connecting the top 2.5% of connections. The square plots show the full set of data in matrix format.

3.2 Low-Fidelity Results

Figure 6 compares the actual IOP and predicted IOP for test data on which the model was not trained. Simulations of boxes positioned to match available experimental data are blue, whereas general test data is orange (Figure 6). The model performance is compared when provided data for two to five sensors at a time. The same faces are used for consistency of results, but their order is randomized when generating predictions of IOP. The model has good average performance with average error below 10% even with data from only two sensors, but significant outliers exist when fewer sensors are used. The average error and

number of outliers decrease with more sensor data, including the case with five sensors for which the model was not trained.

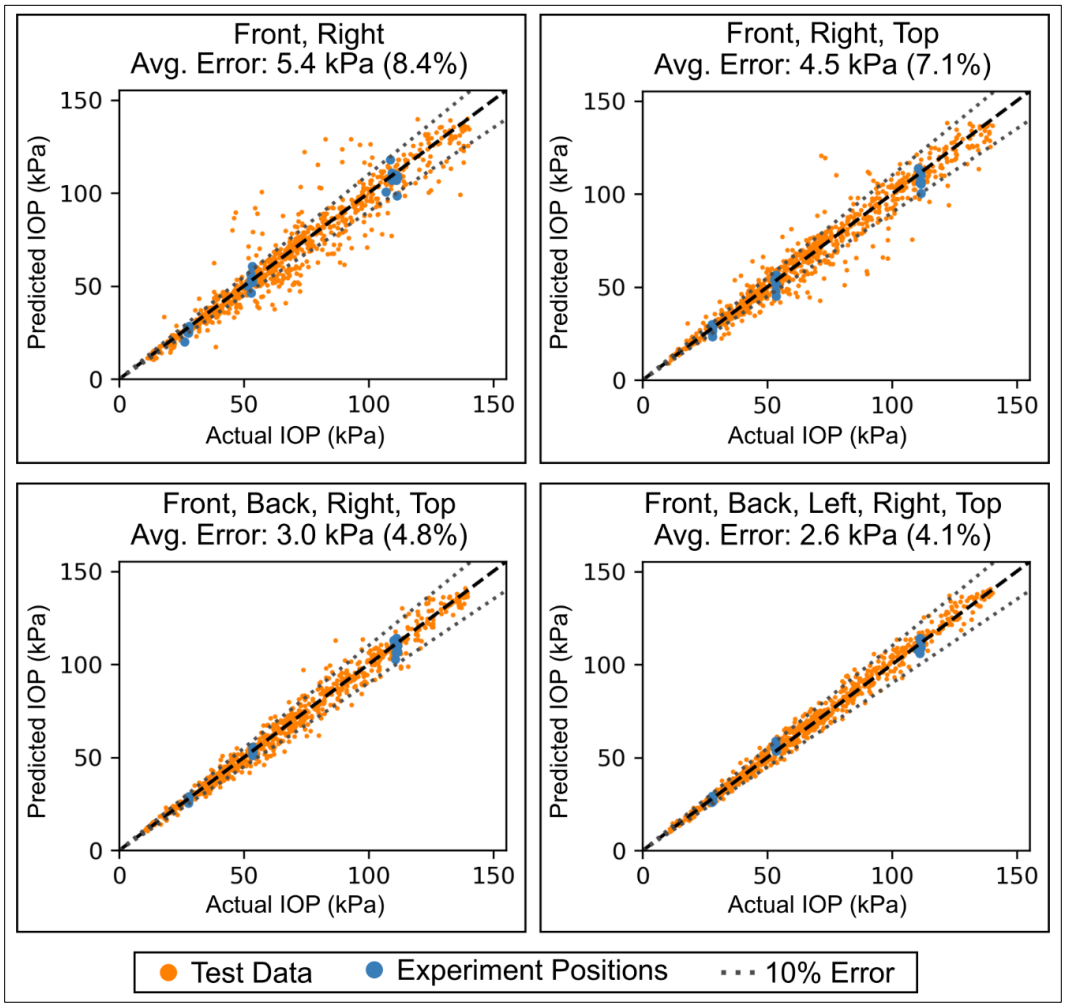

Figure 6. Comparison of actual simulated IOP to model predictions of IOP when running the model with two to five input sensors. Markers for experimental positions are low-fidelity simulations in the same configuration as available experimental data.

A model trained on low-fidelity data is not expected to be able to make accurate predictions when provided data from a high-fidelity source. However, it is important that it can still make consistent and sensible predictions given the differences between low- and high-fidelity data trends to enable a multi-fidelity approach. Figure 7 shows the target IOP from low-fidelity simulations compared to model predictions using low- and high-fidelity data for the three different charge sizes used in experiments. As expected, the model makes poor predictions when run with data from high-fidelity experimental sensors, but it does provide consistent results. Low-fidelity simulations tend to have higher peak pressure, likely due in part to the perfectly flat geometry on which simulated sensors are placed. Lower peak pressure measurements from experimental sensors result in predictions of

lower IOP from the model as expected. These results indicate that a multi-fidelity approach can be applied to make use of the trained low-fidelity model.

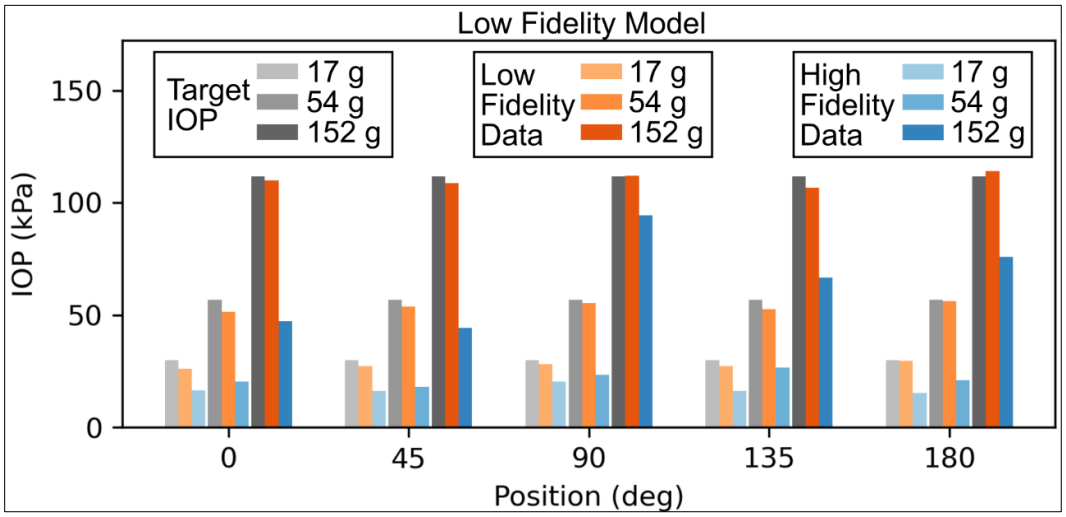

Figure 7. A comparison of the target IOP to model results at locations where experimental data is available for different charge sizes. The target IOP is the low-fidelity-simulated IOP and not the experimentally measured IOP.

Figure 8 shows the error when estimating direction to blast from test data using sensors on the front, right, and top of the box. General test data is shown as orange markers and low-fidelity simulations in positions to match experimental data are shown as blue markers (Figure 8). Error is higher when predicting direction to the blast compared with IOP. The average error was 9.0° for the test set, but outliers have significant errors. The amount of error changes as the box transitions to positions where the sensors are exposed to the blast or obscured by the box. The front sensor is exposed and facing into the blast source at 0°, the right sensor is exposed and facing into the blast at 270°, and the top sensor may be exposed or obscured based on the height of the box and charge in a simulation. The highest error occurs at positions where the front sensor is obscured by the box. Boxes positioned to match experimental positions did not have significant outliers.

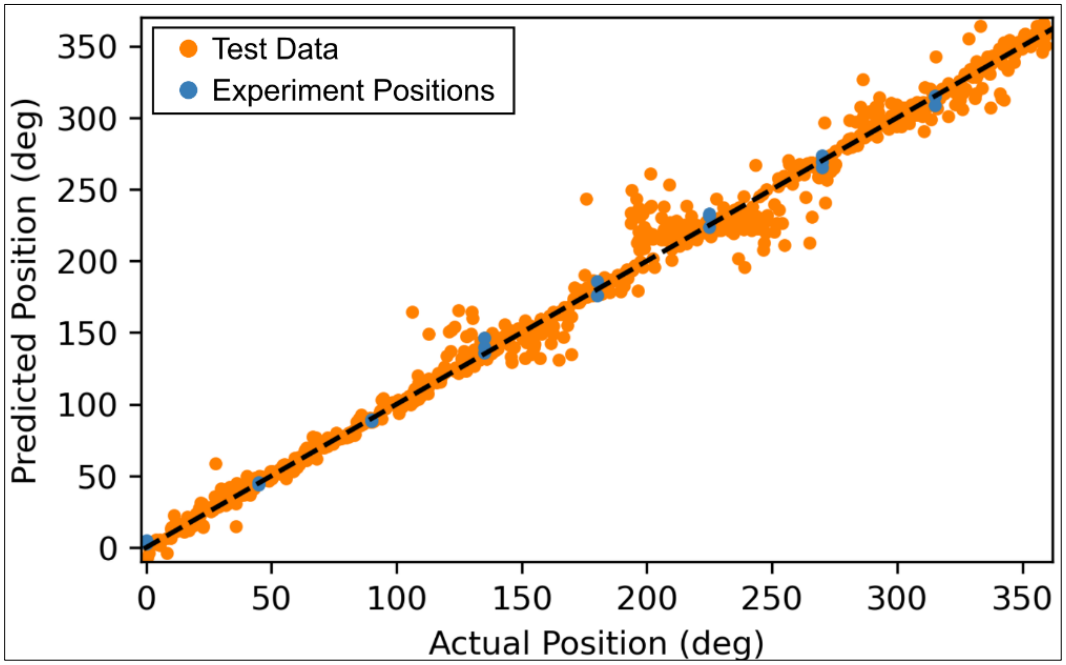

Figure 8. A comparison between simulated box positions and the position from the model prediction. Markers for experimental positions are low-fidelity simulations in the same configuration as available experimental data.

The low-fidelity direction model was also run with high-fidelity data to evaluate how it performs on the other dataset. Figure 9 shows markers on the outer ring of the circle for predictions when using low-fidelity data that matches experimental positions. The solid lines connecting to the center of the circle indicate the average prediction, which shows good results for this set of data. The inner ring has markers showing results when using high-fidelity data with the low-fidelity model. Dashed lines indicate the average result from high-fidelity data. Unlike the IOP model, predictions for direction using high-fidelity data show a large amount of scatter with less consistent results. Directions for the smaller 17-g charge in particular had poor results with predictions in all directions. It is in part because of this result that fine-tuning was chosen for training on high-fidelity data over a multi-fidelity model.

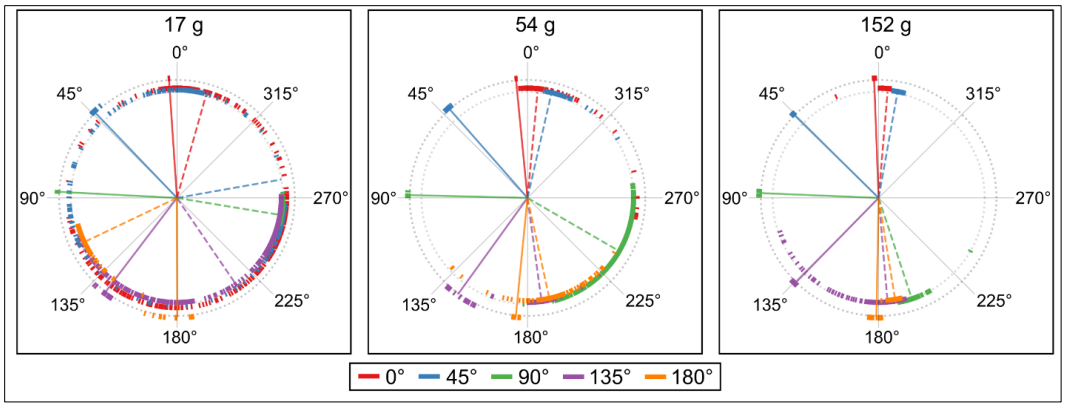

Figure 9. A comparison between low-fidelity model prediction of direction to the blast source when run with low-fidelity and high-fidelity data. Colors indicate the actual position, and the marker is plotted on the predicted position. Predictions from low-fidelity data are plotted on the outer ring with a solid line indicating the average result. Predictions from high-fidelity data are plotted on the inner ring with dashed lines to indicate the average result.

3.3 High-Fidelity Results

Results for training the IOP multi-fidelity model are shown in Table 2. These results are plotted in Figure 10 to compare both average results and plot individual predictions to visualize any outliers. Average test results typically fall within the bounds of ±1 StDev of experimental IOP. The 54-g charge at 90° is the only result that fell outside of this range, which had the highest average error at 4.6%.

Table 2. Results of the multi-fidelity IOP model compared with experiments for three charge sizes. Results for test results are in bold.

Position (deg)	17 g			54 g			152 g		
	Avg. error (kPa)	Avg. error (%)	StDev (kPa)	Avg. error (kPa)	Avg. error (%)	StDev (kPa)	Avg. error (kPa)	Avg. error (%)	StDev (kPa)
0	**0.2**	**0.6**	**0.2**	0.3	0.6	0.4	0.3	0.3	0.4
45	0.3	1.0	0.3	0.9	1.7	1.1	0.4	0.3	0.4
90	0.1	0.5	0.2	**2.5**	**4.6**	**0.6**	0.8	0.7	0.9
135	0.3	0.9	0.4	0.9	1.6	1.3	**2.5**	**2.4**	**3.4**
180	0.1	0.5	0.1	**1.9**	**3.4**	**2.4**	1.1	1.1	1.3
225	**0.4**	**1.5**	**0.4**	1.1	2.0	1.3	1.3	1.2	1.7
270	0.3	1.1	0.2	0.4	0.7	0.4	1.2	1.1	1.5
315	0.2	0.8	0.2	1.2	2.2	1.4	**2.2**	**2.1**	**1.8**

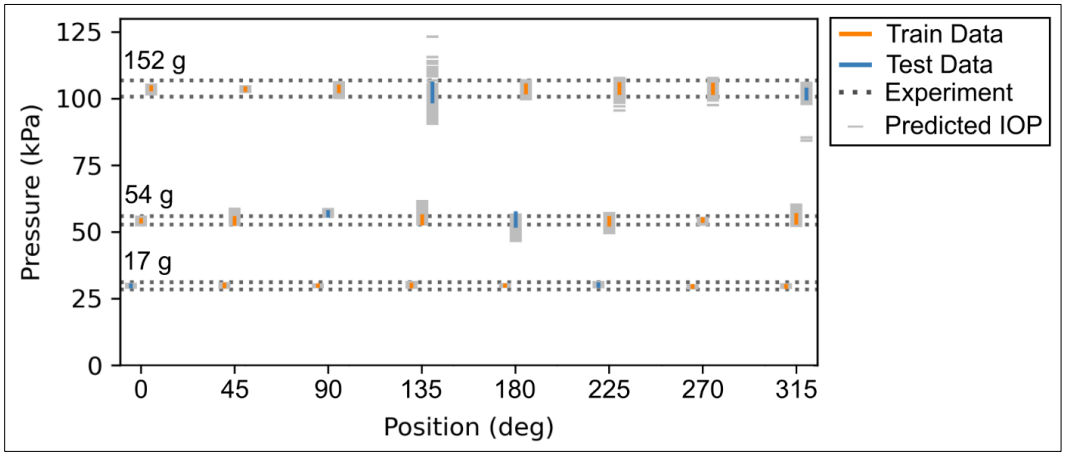

Figure 10. Results of the multi-fidelity model for each charge size. Horizontal dotted lines show experimental IOP ± 1 StdDev. Gray dashes indicate a single model prediction. Vertical lines show average model IOP ± 1 StdDev. Markers are offset from their exact position degree to visually separate results for each charge size.

The best results for predicting direction for high-fidelity data are achieved through fine-tuning the low-fidelity model. Table 3 lists the average error for each charge size and position. Figure 11 is a visual representation of the results where markers are plotted in positions based on the predicted direction to the blast. Training results are plotted on the outer ring with a solid line to indicate the average. Test results are plotted on the inner ring with a dashed line to indicate the average. Test results generally have low average error but higher StdDev compared with training results. The 54-g charge at 180° had the highest average test error at 16.4°.

Table 3. Results from fine-tuning the low-fidelity model on high-fidelity data to predict the direction to the blast. Test results are in bold.

Position (deg)	17 g		54 g		152 g	
	Avg (deg)	StdDev (deg)	Avg. (deg)	StdDev (deg)	Avg. (deg)	StdDev (deg)
0	**2.7**	**3.2**	0.1	0.1	0.1	0.1
45	0.4	0.3	0.3	0.2	0.8	0.6
90	0.7	0.5	**0.8**	**0.8**	0.6	0.4
135	2.1	1.1	1.9	0.8	**3.4**	**2.6**
180	1.1	0.7	**16.4**	**4.7**	0.9	0.7
225	**3.3**	**4.7**	0.3	0.3	0.3	0.2
270	1.8	0.4	2.1	0.4	**1.2**	**1.0**
315	0.6	0.4	0.6	0.4	0.8	0.6

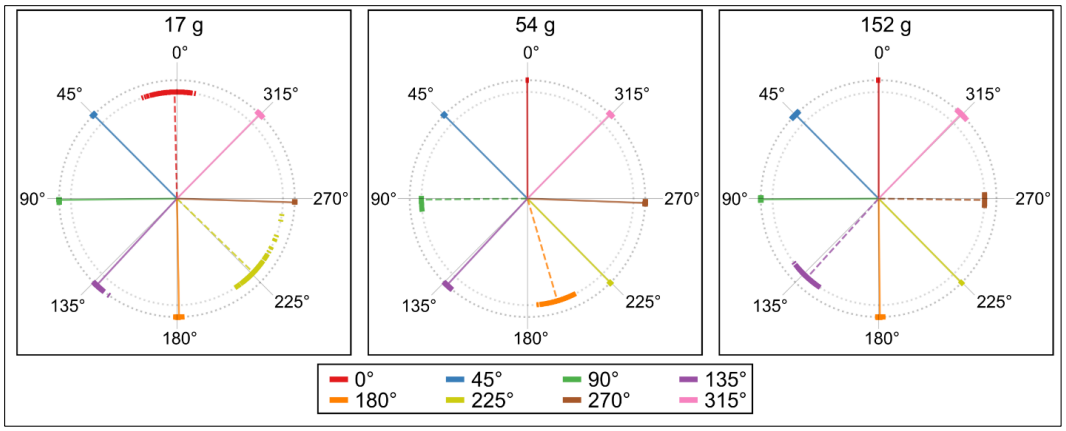

Figure 11. A visual representation of results for prediction of direction. Training data is plotted on the outer ring with a solid line to indicate the average value. Test data is plotted on the inner ring with a dashed line to indicate the average value.

4. Discussion

The primary goals of this report are to first develop and train an ML model to extract important features from pressure history data with a variable structure when there is an abundant quantity of data and then adapt that model to a smaller dataset from a different source. The advantage of the transformer model presented in this report is the ability to use the attention mechanism to enable greater flexibility for data input. Alternatives such as convolutional neural networks (CNNs) could also be applied, but they required more constraints for processing data to ensure inputs are a consistent size. These constraints could make a CNN more challenging for adapting to systems with different numbers of WBOPS. Although the transformer model did lose accuracy when making predictions of IOP with data from only two sensors, it is still able to make a reasonable estimate that could be useful in cases where one sensor fails to record a blast event. Similarly, the model continues to improve performance when more sensor data is input.

Several challenges still exist for maintaining input flexibility. First, fine-tuning the low-fidelity transformer model will maintain the same flexibility in input data, but the multi-fidelity model used for IOP has constraints on input size. However, the simple architecture of the multi-fidelity model allows for fast training that makes training unique models for each possible number of input sensors reasonable. This multi-fidelity model also only uses peak pressures as input to maintain simplicity, so the length of pressure data from different blast sources does not present a concern. Next, the current direction model has more restrictions than the IOP model. The low-fidelity direction model was not effective unless it was provided sensors from the same faces of the box in the same order. This model was only trained with a fixed number of sensors for this reason, but future work is focused

on resolving this constraint. Training the model where the input contains positions for all possible input sensors and then using an attention mask to ignore unused sensors could enable the model to be run and trained with a variable number of sensors.

The difference between methods used for the IOP and direction models illustrate how differences in data affect how models perform. Predicting IOP was insensitive to the location and order of sensors used in the low-fidelity model, which allowed the model to be easily trained to make predictions when picking any set of sensors. In contrast, the direction model was more limiting and required using sensors in a fixed order. This is an important distinction when applied to real-world applications because the ideal sensor design may change based on what output metric is considered the most important. The differences between use of a multi-fidelity model and fine-tuning also present important considerations for data structure. Attempts to fine-tune a model for IOP caused the model to quickly overfit because there were limited possible IOP values in the experimental dataset. The simpler model used for multi-fidelity IOP predictions was less susceptible to overfitting and performed better. In contrast, fine-tuning the direction model was more effective compared with the multi-fidelity approach. More directions were possible in experimental data, and the need to output a vector rather than a scalar value may have helped to reduce overfitting. However, tuning this model did still require a careful selection of hyperparameters to prevent overfitting from excessive tuning. These differences illustrate how the optimal method of training a model may change based on the available training data and goals of the model.

The need for different methods between the IOP and direction model may also relate to the simplified geometry of the low-fidelity model. The ability to make an accurate prediction of the direction to the blast is very sensitive to the orientation of the sensor face, as seen by the need to use fixed sensor location. The box geometry cannot replicate sensor orientations seen in experiments. Even a sensor placed on the chest will have a difference in orientation depending on if it is placed in the center or to either side. This resulted in less consistent results when inputting high-fidelity data into the low-fidelity model as seen by comparing Figure 9 with Figure 7 where prediction of IOP had low accuracy but showed more predictable results. Fine-tuning the model has a greater capacity to adjust the model parameters to make the greater change necessary for a high-fidelity direction model.

Both the IOP and direction high-fidelity model showed the ability to make reasonable predictions on test data on which the model was not trained. The high-fidelity experiments have relatively large gaps when test data is withheld from training, unlike the low-fidelity data, which emphasizes the ability of the model to interpolate between the inputs on which it was trained. However, experimental data

with higher variability or regions that exist at the boundaries of the input space where the model needs to extrapolate present a greater challenge. Withholding certain groups of data for training caused significantly larger test errors compared to the results displayed. For example, the 152-g charge at 315° had more variability in experimental results compared with other positions, which caused significant errors when it was used as a test set for the direction model. The IOP model was less sensitive to this variability, but some outliers can still be seen in Figure 9. The 54-g, 135° charge had a higher error when used as a test set for the IOP model. In this report, one of the experimental sensors had relatively high pressure, which caused difficulty in interpolating between neighboring positions of the 54-g charge or between the 17- and 152-g charges at 135°. The multidimensional nature of this problem can make interpreting and anticipating regions that may have higher error difficult. For future work, it is important for researchers to consider where a model may perform poorly and where additional data would be useful to allow the best interpolation.

5. Conclusion

ML models have been developed to process data from WBOPS and make predictions from pressure history data. The model is demonstrated using IOP and the direction to the blast as output metrics. The transformer architecture allows the model to read in data of variable length from a variable number of sensors. Finally, methods were shown to use low-fidelity simulations to assist with training the model more efficiently when high-fidelity data is more limited. Models for IOP and direction to the blast source were developed here to make use of available low-fidelity simulations and experimental data, but these methods also provide an approach to make predictions on additional metrics for Soldier safety. Research is currently underway to perform experiments and develop high-fidelity models to better understand behind armor loading and its impact on the body that will provide additional data sources that can be used to train ML models for additional metrics.

6. References

1. Bessette G. BlastX user manual for software versions 9.0 and higher. Army Engineer Research and Development Center (US); 2023 Apr. Report No.: ERDC/GSL SR-25-2.

2. Dean SW et al. Comparative evaluation of wearable blast overpressure sensor designs with personal protection armor systems, year 1: test methodology and reference data. DEVCOM Army Research Laboratory (US); 2025 Sep. Report No.: ARL-MR-1138.

3. Vaswani A et al. Attention is all you need. Adv Neural Inf Process Syst. 2017;30.

4. Carion N et al. End-to-end object detection with transformers. European Conference on Computer Vision; 2020 Aug 23. p. 213229.

5. Zhu X, et al. Deformable DETR: deformable transformers for end-to-end object detection. arXiv. 2020 Oct 8. arXiv:2010.04159.

6. Wen Q et al. Transformers in time series: a survey. arXiv. 2022 Feb 15. arXiv:2202.07125.

7. Loshchilov I, Hutter F. Decoupled weight decay regularization. arXiv. 2017 Nov 14. arXiv:1711.05101.

8. Meng X, Karniadakis GE. A composite neural network that learns from multi-fidelity data: application to function approximation and inverse PDE problems. J Comput Phys. 2020;401:109020.

List of Symbols, Abbreviations, and Acronyms

2D/3D	Two-/Three-Dimensional
BOP	Blast Overpressure
CNN	Convolutional Neural Network
IOP	Incident Overpressure
MHA	Multi-Head Attention
ML	Machine Learning
MLP	Multilayer Perceptron
NLP	Natural Language Processing
WBOPS	Wearable BOP Sensors

1 (PDF)	DEFENSE TECHNICAL INFORMATION CTR DTIC OCA
1 (PDF)	DEVCOM ARL FCDD RLB CB TECH LIB