MODELLING THE DEPTH OF PENETRATION OF CERAMIC TARGETS USING ARTIFICIAL NEURAL NETWORKS

C.C. Renahan[1], W.S. Andrews[1] and K.M. Jaansalu[2]

[1]*Department of Chemistry and Chemical Engineering, Royal Military College of Canada, Kingston ON Canada K7K 7B4*
[2]*Department of Metallurgical and Materials Engineering, Montana Tech of the University of Montana, Butte Montana USA 59701-8997*

Data collected from depth of penetration (DOP) testing has been used to develop an artificial neural network (ANN) model that was able to predict the ballistic performance of the ceramic armours Al_2O_3, SiC, B_4C, and TiB_2. The model provided good predictions of DOP and interpolated well over the input space of the specific experimental conditions and ranges. The ANN employed the back propagation paradigm with a single hidden layer architecture and, considering up to 39 input variables and parameters, was able to establish a strong correlation between predicted DOP and impact velocity, ceramic type, ceramic density, ceramic thickness, projectile length, projectile diameter, and projectile mass.

INTRODUCTION

The increased range of threats facing combat vehicles has led to developments in protective measures, including the use of ceramics, with the earliest and most widely used ceramic for armour being alumina (Al_2O_3). Other common ceramic armours now include the more expensive silicon carbide (SiC), titanium diboride (TiB_2), and the most expensive, boron carbide (B_4C).

Depth of penetration (DOP) testing has been used to investigate the effectiveness of ceramics since approximately 1968, and is the most common methodology for characterizing the performance of ceramic materials as armour components. It is, however, an experimental procedure, that is best suited to classifying or ranking the performance of various ceramics.

The fracture of ceramics has been studied in depth, but an accurate model of their failure has yet to be developed. The high velocity of the impact and dynamic loading, along with the brittleness of the ceramic, make predictions difficult. The mechanisms involved in this impact are also not completely understood and it is not clear how to relate a ceramic's mechanical properties to its ballistic efficiency.

Artificial neural networks (ANNs) have become one of the fastest-growing fields of artificial intelligence. They can be trained to identify non-linear patterns between input and output values, and can learn complex and non-linear relationships, involving noisy or incomplete data. The back propagation algorithm used in this study (the most common of ANN types) employs gradient descent to minimize the error between the predicted and target values over the output space.

MODEL DEVELOPMENT

A significant amount of time and effort went into the selection and development of the data file, which had to be complete, balanced, and representative of the range of ceramics and factors in question. The majority of the data used in the development of the network were collected from experimental trials of ceramic plates, beginning in 1990. The sources of data used had varying degrees of completeness. In some cases, data had to be discarded due to incomplete information on what were felt to be critical areas, such as the thickness of the ceramic plate in question.

The test facility allowed for projectiles of up to 14.5 mm in calibre to be fired into target plates, that were backed by witness blocks of Al 6061-T6 aluminum cylinders. A magnetic detector at the muzzle was used to trigger four X-ray cameras that took images that were used to measure the projectile’s attitude and velocity at impact.

The data consisted of the results of individual test firings, with each having associated with it up to 63 individual characteristics or measurements. Over 300 data vectors were assembled and evaluated for their completeness and usefulness. Overall, 40 of the 63 characteristics or measured values were determined to be sufficiently complete in a majority of the data vectors available, and were felt to be relevant and important to the performance of the ceramics (see Table 1). The resulting data file had 267 data points (results of individual tests) that were deemed to be valid for the development of the network model. In order to ensure that there would be no inadvertent biases introduced into the ANN, the data vectors were manipulated somewhat (some variable values were transformed and some vectors were repeated) to ensure to the greatest degree possible that there were uniform distributions throughout the output and input spaces.

The distribution and relative number of the 39 input variables required more manipulation in order to have them more evenly distributed. This balancing of data resulted in a data file of 1238 vectors, that was further divided into two, a training and a test set, with the latter containing 250 vectors (selected randomly), approximately 20% of the training set. A large number of networks were trained using various architectures (numbers of hidden layers and hidden nodes in these layers) and network parameters, such as epoch size (the number of data vectors trained on before updating the connecting weights). The most successful ANN model developed consisted of one hidden layer containing 87 hidden nodes, having an RMS error of 0.074 and a correlation coefficient of 0.94 for the test set.

Table 1. Experimental data values used in model development.

Ceramic Properties	Projectile Properties
Tile type (Al_2O_3, SiC, B_4C, TiB_2)	Mass
Average grain size	Diameter
Void fraction	Length
Longitudinal velocity	Material (W or steel)
Shear velocity	Hardness
Bulk velocity	Nose shape (ogive or cone)
Young's modulus E	Projectile core L/D
Shear modulus G	Distance from muzzle to target
Bulk modulus K	Impact velocity
Poisson's ratio	Impact yaw angle (°)
Compressive strength	Adjusted residual DOP
HEL (GPa)	
Tile thickness	
Tile surface area	
Tile density	
Areal density	
Confinement material (Al or steel)	
Confinement type (semi or full)	

In order to create a more manageable model, an effort was made to reduce the number of inputs required. This removal of inputs, or pruning, took place by removing one input at a time and assessing the impact of the resultant lack of data on the performance of the network. The end result was a more general model, requiring fewer inputs to produce the same quality of output. The process followed was to identify the input parameter that least affected the output of the network, remove it, and retrain the network. This process was repeated until the network showed a marked decline in its ability to provide adequate predicted values. In fact, network performance remained relatively consistent as input parameters were pruned, but dropped off significantly with less than ten input parameters.

It must be noted, however, that these results are a function of the specific data studied and that there may be other reasons why some factors are seen to have minimal effect on the output of the network. Inputs that were underrepresented in the training data, or else were very limited in range, may have been interpreted as being less important by the network during training. If their relative importance or overall effect on the output were underestimated during network development, they would be removed quickly. Therefore the ranking of these factors must be evaluated in the context of this specific data set (for example, the ranges of impact velocities were probably not great enough for the ANN to be sensitive to changes in these values), and broader generalizations or conclusions on a ceramic's performance on the whole should be made with caution. That said, the factors determined to have the most significance (in decreasing order) are shown in Table 2.

Table 2. Factors determined through a sensitivity analysis to have a significant effect on model predictions (listed in order of decreasing significance).

Parameters in Decreasing Order of Significance
Projectile diameter
Ceramic tile density
Ceramic tile thickness
Projectile core nose shape
Ceramic tile shape
Ceramic tile average grain size
Confinement material
Ceramic tile areal density
Ceramic tile void fraction

The result, when using the reduced, or pruned, networks of only 10 inputs, is shown in Figure 1 below.

MODEL PREDICTIONS

The performance of the selected network was evaluated through the use of artificial data vectors. For each of the 39 data values, the trained network was tested using a new data file consisting of synthetic or artificial data vectors, where, in most cases, a single data vector was considered in which one field was varied to produce a range of input vectors. This resulted in a data file where one of the data characteristics was varied, while the remaining attributes remained constant. When the network processed this data file, the resulting output (predicted DOP) changed as a function of that single variable's value in the data file. That is, in the case of impact velocity, the data file contained a range of artificial velocity data as a part of a series of otherwise constant input vectors and, when processed by the network, produced a series of predicted DOP value outputs that were the function of the change in impact velocity.

To create these artificial vectors, four vectors for each of the four ceramic types were chosen, with each being representative of the full range of data characteristics for each of the four ceramics. Figure 2 includes the predicted value of DOP, indicated by lines, with the actual measured or real value of the particular variable under consideration also indicated on the graph. The range of data included in the development of the network is indicated below the predicted curves, be it for a specific ceramic or else the overall range of the variable in question. An example of the ANN model's quality of predictions is shown for varying impact velocities for alumina. For all four Al_2O_3 cases, as impact velocity increased, DOP increased. This is as expected, and is illustrated in Figure 2.

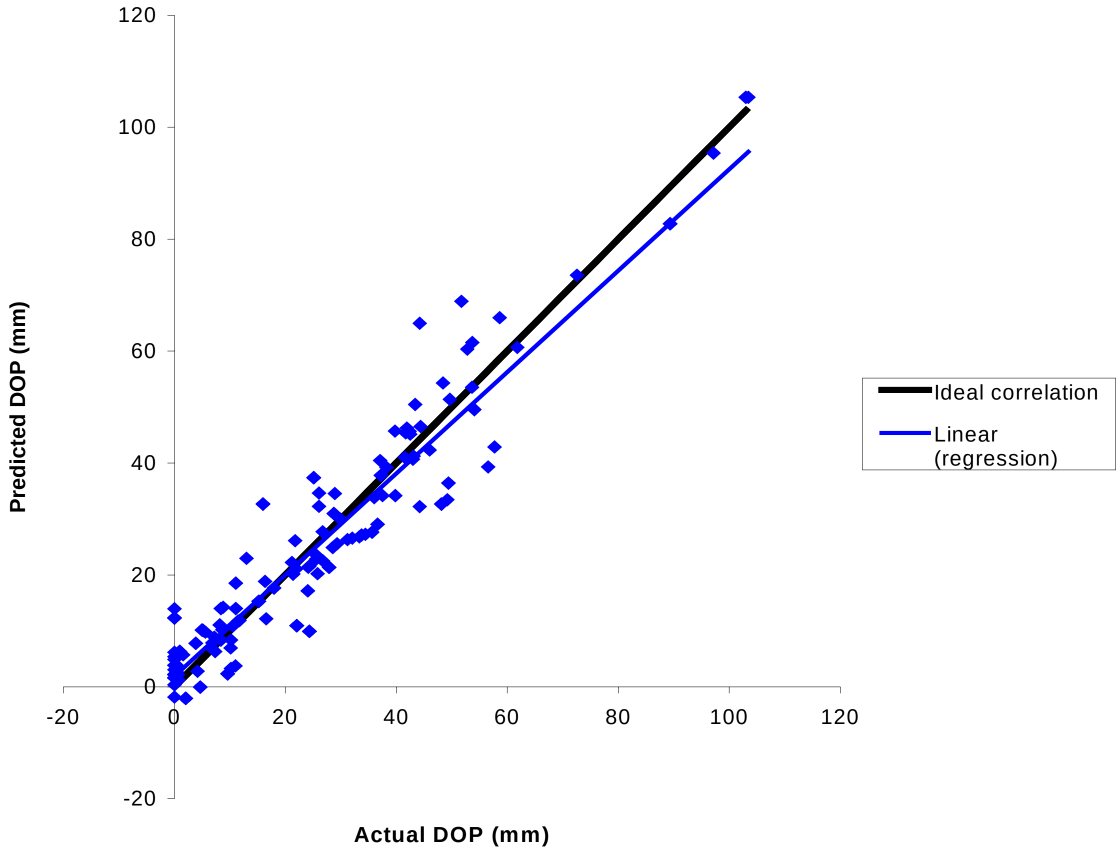

Figure 1. Predicted DOP versus actual DOP when using only 10 remaininginputs after pruning.

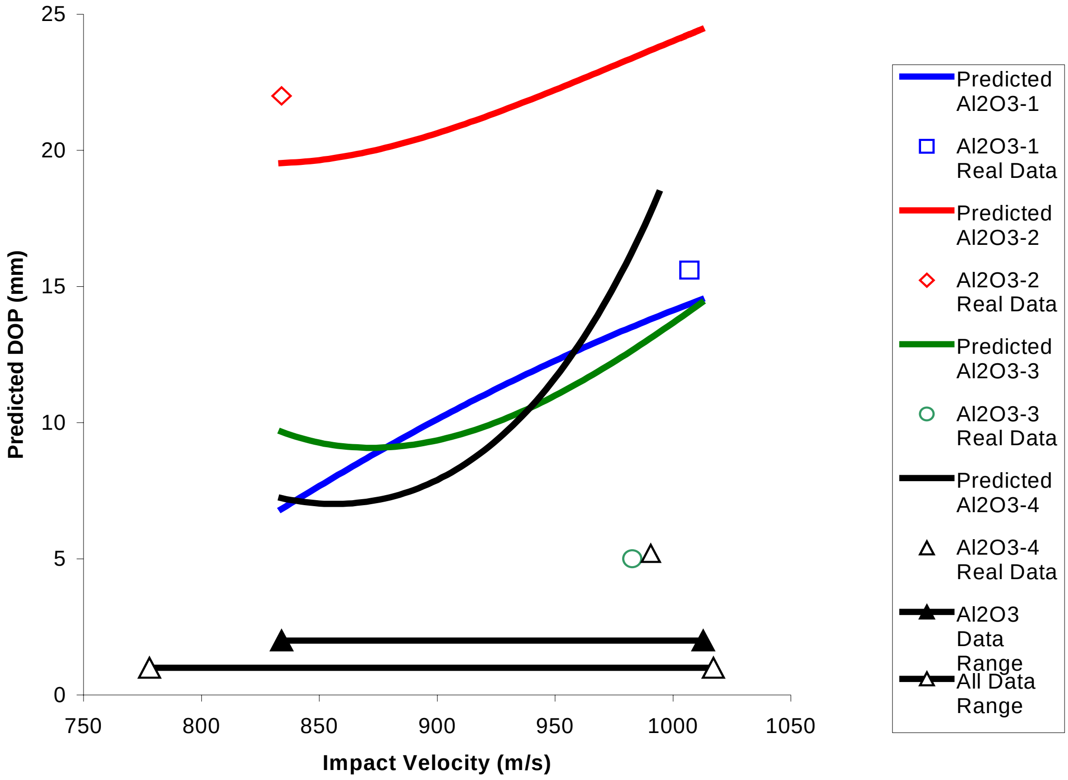

Figure 2. Predicted DOP as a function of impact velocity for Al_2O_3.

A number of the predictions show non-linear behaviour, clearly evident in cases Al_2O_3-3 and Al_2O_3-4 above, with the increasing DOP at lower impact velocity. Al_2O_3-1's (A1) curve is significantly different from the other three, as it is concave downwards as opposed to the others (concave upwards). One difference that sets this case apart is that the yaw of the projectile was 3.2°, which was higher than the value of 3° that was initially deemed as the limit for valid experimental results. This is the Al_2O_3 curve that is felt to best demonstrate the expected v_{50} curve shape within the velocity range in question.

Ceramic A2 was about one quarter the thickness of the other three, which accounts for the DOP being much higher for the same impact velocity. A much smaller projectile was used, although the impact velocity was approximately the same. A2 was also the only one of the four with slightly different ceramic material properties, having slightly lower longitudinal velocity, Young's modulus, shear modulus, Poisson's ratio and compressive strength. In contrast, the values of average grain size, void fraction, shear velocity, bulk velocity, bulk modulus and HEL were constant for all the Al_2O_3 tiles.

The impact velocity of the projectile for the firings used for case A1 varied only between 1001-1009 m/s, with those of A2 being 834-844 m/s, and for A3 and A4 being 974-1013 m/s. The discreteness of these ranges for each projectile type may explain the curves' uncharacteristic performances outside of each projectile's or ceramic's given range of data. In each specific range band, the DOP versus impact velocity predictions were as expected.

When experimental uncertainty was applied, two of the curves (A1 and A3) overlapped for 85% of the range in question, with A4 overlapping again at the lower end of the curve. This shows that some of the curves in certain areas of the data space are not as different as they may appear. Some seemingly significant differences, such as the concave shape of A1 and the convex shape of A3, are much less of a concern when it is seen that the curves generally overlap within the applied error bars. The reason for predictions shown for case A2 falling outside of the general range covered by the others, may be due to differences in ceramic composition or dimensions alluded to above.

Similar analyses were done for the other ceramic materials under consideration and for many of the parameters involved (where there was sufficient range in the values available to warrant such an examination). The results were comparable. Overall, the model generated good predictions of the performance of various ceramic tiles within the range of data on which it was trained. In this sense, it became less a generalized model than an intelligent look-up table, albeit performing in multi-dimensional space.

Using DOP test data, then, an ANN was successfully trained and evaluated to determine its suitability in predicting the ballistic performance of various types of ceramic armour. An overall uncertainty of 5% was determined to apply to the predicted DOP output of the network. The model developed was found to be tuned to the data ranges with which it was trained. It should also be noted that for significant gaps in a data range, the ANN also had difficulty in generating sound predictions in areas of sparse training data. Where there existed sufficient amounts of data for network development, the networks provided predictions that were consistent with expected trends.

The network was found to perform well in the representation of the relationships between predicted DOP and impact velocity, ceramic density, ceramic thickness, projectile length, projectile diameter, and projectile mass. For impact velocity, in the majority of the cases considered, the behaviour of the predicted DOP versus velocity curves was as expected, with DOP increasing as the impact velocity increased. The curves also corresponded to the expected profile of a v_{50} curve.

For density, the curves were the most consistent of all the predicted DOP relationships, both within each ceramic type, and among all four types of ceramic. While the predicted values outside the data ranges in question may have been suspect, the uniformity of the results, and the consistency with the expected trend, are strong indications that this is a valid result and prediction.

For ceramic thickness, the curves for three of the four ceramics were quite consistent, and the fourth followed the expected trend. The curves, however, did not exhibit a very significant decrease in DOP for increased ceramic thickness, indicative of the little change in DOP for changes in thickness of the ceramic within this range of data, due to factors such as sonic velocity and comminution of the ceramic. While the ceramic is, in all likelihood, performing as expected, a wider range of experimental data would be required for validation.

Varying projectile length, and the resultant change in the projectile's mass and therefore kinetic energy, provided a consistent general trend in prediction. The same can be said for diameter, although the added factor of increasing the surface area of impact as the diameter of the projectile increases is also a factor. The L/D relationship was not felt to be significant, as the projectile was not being consumed in the impact.

The relationship between predicted DOP and projectile mass was as expected, with DOP increasing with the increased mass, and therefore kinetic energy, of the projectile. The consistency of these results is indicative of the significance of mass in the relationship, and the network's sensitivity to this factor. The significant range of data points in most cases is also an indicator that this relationship was well accounted for in the development of the network.

Predicted DOP values were found to be relatively insensitive to the yaw of the projectile at impact.

Limited data, be it a very small range or with only very few discrete values available, made it difficult to assess the network's ability to correlate or predict DOP as a function of many of the properties included in the data set, including projectile core hardness, ceramic average grain size, void fraction, longitudinal velocity, shear velocity, bulk velocity, Young's modulus, shear modulus, bulk modulus, Poisson's ratio, compressive strength, and HEL. That being said, there were several cases where the model was able to accurately rank the ceramics' effectiveness in relation to one another, in keeping with the characteristics of the DOP methodology.

CONCLUSION

In summary, then, a neural network model has been trained and tested that provides good DOP predictions within the ranges of the data used in its development for four ceramic materials, Al_2O_3, SiC, B_4C, and TiB_2.