

Effectiveness of Neural Networks for Research on Novel Thermoelectric Materials. A proof of concept.

Filippo Remonato^{1,*}, Ole M. Løvvik², and Espen Flage-Larsen²

¹ SINTEF Digital, department of Applied Mathematics and Cybernetics, Oslo, Norway *corresponding author: `filippo.remonato@sintef.no`

² SINTEF Industri, department of Materials Physics, Oslo, Norway
`olemartin.lovvik@sintef.no`, `espen.flage-larsen@sintef.no`

Abstract. This paper describes the application of neural network approaches to the discovery of new materials exhibiting thermoelectric properties. Thermoelectricity is the ability of a material to convert energy from heat to electricity. At present, only few materials are known to have this property to a degree which is interesting for use in industrial applications like, for example, large-scale energy harvesting[3,8]. We employ a standard neural network architecture with supervised learning on a training dataset representing materials and later predict the properties on a disjoint test set. At this proof of concept stage, both sets are synthetically generated with plausible values of the features. A substantial increase in performance is seen when utilising available physical knowledge in the machine learning model. The results show that this approach is feasible and ready for future tests with experimental laboratory data.

Keywords: Neural Networks · Thermoelectric Materials · Physics-oriented Machine Learning.

1 Introduction

About 70% of the energy globally produced is released in the atmosphere as heat [2]. Given the predicted increase in energy demands, coupled with the pressure on our planet's resources and climate, it is clear we cannot afford to lose such amounts of energy, and is therefore important to be able to recover the produced heat and use it in further processes. The most common way to recover waste heat is for heating buildings or smaller environments like a car. This produces already good results: An average fossil-fuel power plant is able to turn into electricity only about 30% of the energy contained in the fuel, but when the excess heat is recovered, the total efficiency rises to 60-80% [1]. However, heat is sub-optimal compared to other energy sources: It is difficult to transport, it is easily lost, and in general our technology to exploit heat is not particularly advanced. The reason why waste heat is mostly reused simply in its raw form is that heat has high entropy, it is difficult to transform it back into more valuable energy types. Thermoelectric materials accomplish just that: They are capable of

generating an electrical current from a temperature difference, and are therefore highly valuable ingredients in heat recovery processes and development of new technologies.

This property of converting heat to electricity has sparked wide attention in several research communities and application areas [9,6,4]. Enabling the recovery and use of what was previously almost lost thermal energy, they are seen as an important step toward more sustainable, efficient, and circular industrial processes. Despite numerous ongoing researches, the amount of available data on thermoelectric materials is unfortunately still somewhat limited; this because most current state of the art techniques for material analysis make use of the density functional theory (DFT) [5], which is computationally relatively expensive [7]. In principle, DFT allows us to study in detail and precisely a variety of physical properties of a crystalline material, including whether the material in question presents good thermoelectric properties or not, but doing this for every possible material is clearly not yet feasible.

One quantity indicating how “thermoelectric” a material is, is the so called *figure of merit*, ZT , given by the formula:

$$ZT = \frac{T\sigma\alpha^2}{\kappa_E + \kappa_L}, \quad (1)$$

where T is the temperature, σ is the electrical conductivity, α is the Seebeck coefficient, κ_E is the electronic part of the thermal conductivity, and κ_L is the phonon part (lattice vibrations) of the thermal conductivity. A brief description of these and other physical quantities is given below in Section 2.1.

Machine learning (ML) is a collection of techniques, and the study of such techniques, to produce algorithms capable of inferring the solution to a given problem by analysing sets of (related) data. Machine learning algorithms have the distinctive property of being able to extrapolate from the provided data the often very complex underlying patterns linking input to output, in a “cause and effect” fashion. In this work we focus on using such techniques to predict the ZT value given some physical descriptors of the materials.

We see this as a step toward a practical methodology where ML can be used to scan large parts of the input domain, i.e. the different materials compositions, in order to search for candidate thermoelectric materials. This will allow a laboratory to focus their research and computational efforts to few selected promising materials, thus accelerating the development of new technologies.

2 Methodology

The problem is handled with a classical regression/supervised learning approach: We first train our chosen neural network (NN) architecture using pairs (x_i, ZT_i) , where x_i is a database entry vector collecting a material’s information and ZT_i is its associated ZT value. Then we perform predictions on a test set containing datapoints which have not been used before.

2.1 The database

The dataset we employed consisted of roughly 250000 synthetically-generated datapoints. The descriptors (features) included in the dataset are:

- *num bands*: The number of electronic bands included in the simulation, gives an idea of the complexity of the band structure around the chemical potential;
- *effmass arithmetic*: The effective mass of the charge carriers using the arithmetic mean, in units of the bare electron mass;
- *effmass geometric*: The effective mass of the charge carriers using the geometric mean, in units of the bare electron mass;
- *shift arithmetic*: The average energy shift from zero (location of the chemical potential on the energy scale where 0 eV is the energy reference), of all the electronic bands using the arithmetic mean, units of eV;
- *shift min*: The minimum energy shift (see definition of shift arithmetic);
- *shift max*: The maximum energy shift (see definition of shift arithmetic);
- *T*: The temperature in Kelvin;
- *chempot*: The chemical potential in eV;
- *n*: The free carrier concentration in 10^{21} cm^{-3} ;
- *sigma*: The electrical conductivity, σ , in S/m;
- *seebeck*: The Seebeck coefficient, α , in V/ μK ;
- *kappae*: The electronic part of the thermal conductivity carried by the electrons, κ_E , in units of W/(m · K);
- *ZT_01*: The value of ZT when $\kappa_L = 0.1 \text{ W}/(\text{m} \cdot \text{K})$;
- *ZT_1*: The value of ZT when $\kappa_L = 1 \text{ W}/(\text{m} \cdot \text{K})$;
- *ZT_10*: The value of ZT when $\kappa_L = 10 \text{ W}/(\text{m} \cdot \text{K})$;

The descriptors' values were chosen in a plausible range, so that the tests, while not being connected to any application, could still be instructive regarding the feasibility and performance of this approach when enough real-world data is available. For a similar study performed on a smaller but real-world dataset, we reference the interested reader to [10]. The full dataset of 250000 samples has been randomly shuffled and split in a train set (0.6 fraction), validation set (0.2 fraction), and test set (0.2 fraction).

2.2 Features selection

Not all of the descriptors listed above have been used in our ML model. To choose which descriptors to use as training features we employed a simple correlation analysis, selecting those that displayed high correlation with the target output and discarding those that are highly correlated with each other. The logic behind the former criteria is obvious: In order to exploit the approximating power of neural networks, the features used as input have to be meaningfully related to the target output. The reason for the latter criteria is that one wants to avoid over-representing some characteristics of the data. For example, it's clear that the descriptors *effmass arithmetic* and *effmass geometric* both represent the same

information just in two different ways, therefore only one should be included in the training data.

We note that, since the dataset has been synthetically generated, one could think a correlation analysis approach is not necessary. However, we thought it meaningful to treat the dataset like if it contained unknown real-world data, and therefore applied the same strategies we would have applied in that case.

With the above in consideration, the descriptors that were selected as training features are: *num bands*, *effmass geometric*, *shift arithmetic*, *shift max*, *chempot*, and *n*. Note, in particular, that we did not select the values of σ , α , and κ_E as features since those quantities are the ones that are computed with the expensive DFT and the whole reason to use a NN to predict ZT values is to avoid that step. Indeed, if the values of σ , α , and κ_E were available, one should employ the formula in (1) directly.

2.3 Neural network architecture

Since the purpose of this investigation was to check the feasibility of the approach, we didn't delve too much in fine-tuning the network's parameters and instead wanted to test how an almost out-of-the-box approach performed. For this reason we employed a very standard NN consisting of the input layer for the features, three hidden layers with 100 neurons each with Re-LU activation function, and output nodes, so that the results contained in this work form a baseline for further research and comparison. All layers are fully connected. Both the amount of neurons and number of hidden layers has been decided a priori, before doing any tests, and no regularisation, dropout, or other fine tuned features have been used, so that the results obtained here constitute indeed a lower bound for the quality and accuracy one can expect from this approach. The only tuning we performed was to prevent strongly overfitting the train data. This, staying true to the philosophy of keeping the network's parameters at a minimum, was achieved with early stopping based on the error on the validation set. The model has been constructed in Keras and trained using the Adams optimizer with default parameters values. As loss function we employed the standard mean square error between predictions and target values.

We tried two different choices of activation function for the output layer: In Test 1 no special function was employed, so that the last layer produced a simple linear output. Moreover, since the maximum values of the three ZT s differed at most by one order of magnitude, we did not normalise the target values. In Test 2, on the other hand, we wanted to use the physical knowledge at our disposal, namely that ZT values are positive, to increase the performance of the NN approximation, and therefore enforced output values to be positive through the application of a sigmoid activation function on the last layer. Since the range of the sigmoid is $(0, 1)$, this also naturally called for the normalisation of the target values.

The two simple changes above considerably improved the performance of the model, as shown in the results below.

3 Results

3.1 Test 1

Our first test was conducted with a neural network consisting of an input layer with 6 input nodes, three hidden layers with 100 neurons each and Re-Lu activation function, and a *linear activation* output layer with 3 nodes, one for each value of ZT . The output data has not been normalised.

Figure 1a shows the training history, i.e. the loss function value, for both the train and validation set over the 400 training epochs.

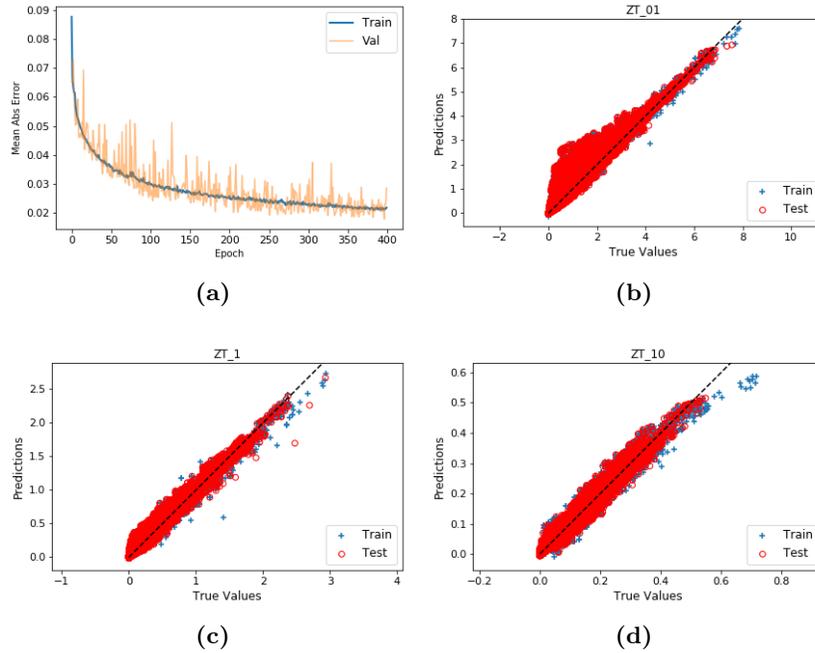


Fig. 1: Test 1. (a) The error function on the train and validation set during the training epochs. (b) The predictions for $ZT_{0.1}$ on train and test set. (c) The predictions for ZT_1 on train and test set. (d) The predictions for ZT_{10} on train and test set. The black dashed line is the bisector of the first quadrant.

Probably thanks to the fact that our dataset is synthetically generated, and therefore we have ample quantity of data and no skewed classes, no appreciable overfitting in the model is observed. The decision to early-stop training at 400 epochs was taken because that is the point at which the validation error reaches a plateau. We also note that although the general trend of the ZT values is correctly identified, the accuracy of the prediction is somewhat coarse. This becomes even more apparent when plotting the same data as in Figures 1b –

1d on a logarithmic scale, to better see the differences in orders of magnitude. Figure 2 shows that the network fails for small values, with the predictions being several orders of magnitude off.

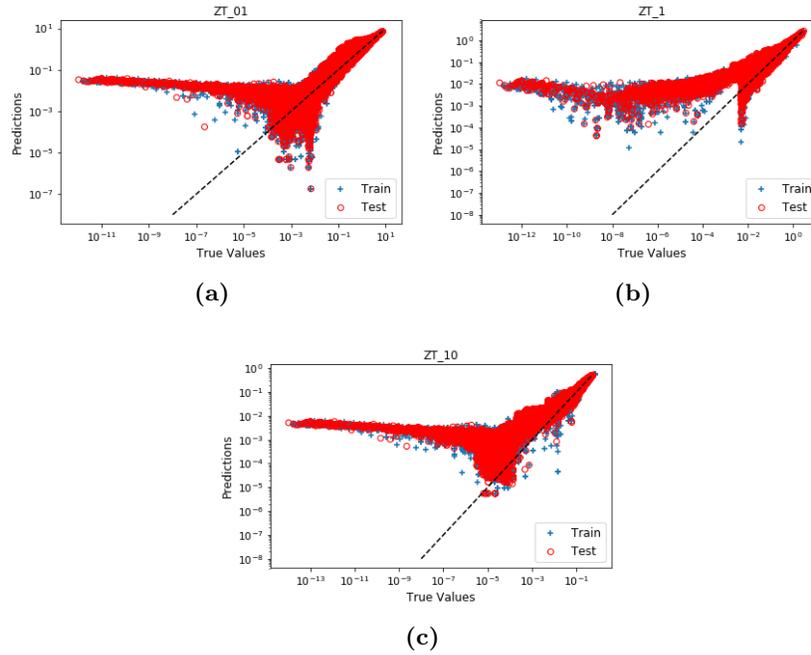


Fig. 2: Test 1. (a) Logarithmic plot of the predictions for $ZT_{0,1}$ on train and test set. (b) Logarithmic plot of the predictions for ZT_1 on train and test set. (c) Logarithmic plot of the predictions for ZT_{10} on train and test set. The black dashed line is the bisector of the first quadrant.

It has to be remarked that the accuracy in the predictions for large values of ZT , which is the area of interest, is acceptable, but the strange behaviour exemplified in Figure 2 made us wonder whether we could improve the performance by adding just a little of physical knowledge to the model; this led us to Test 2.

3.2 Test 2

In the second test we incorporated some physical knowledge into the model. In particular we enforced the output of only positive values by applying a sigmoid activation function to the last layer in place of the normal linear activation. Since the sigmoid range is in $(0, 1)$, the target output has also been normalised before training. Figure 3 shows the results in this case.

As before, no appreciable overfitting is found, but the accuracy of the predictions is sensibly improved as the bands around the diagonal line are narrower.

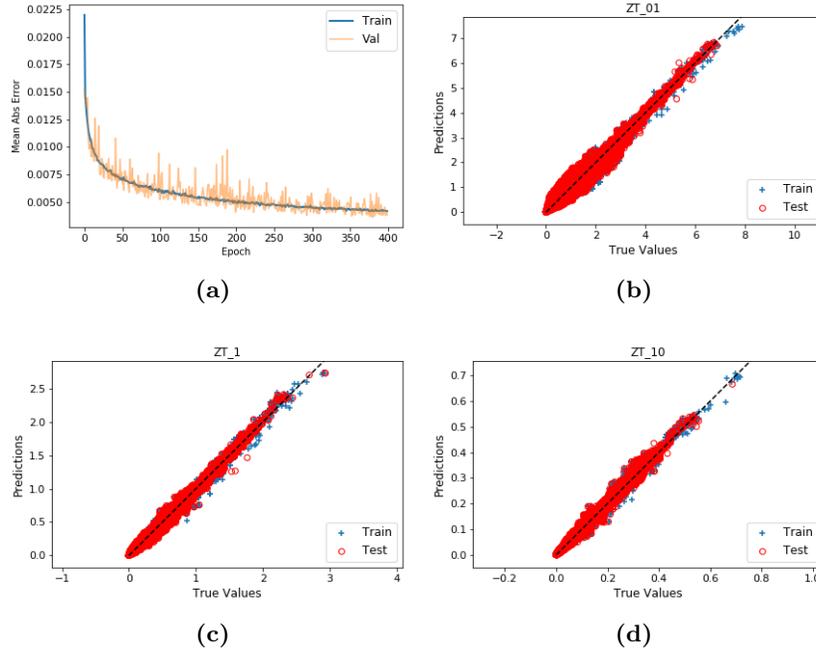


Fig. 3: Test 2. (a) The error function on the train and validation set during the training epochs. (b) The predictions for $ZT_{0.1}$ on train and test set. (c) The predictions for ZT_1 on train and test set. (d) The predictions for ZT_{10} on train and test set. The black dashed line is the bisector of the first quadrant.

This is visible also from the error recorded in the training history, Figure 3a, which after 400 epochs is about four times lower than that in Test 1.

The change in behaviour of the model transpires clearly also in the logarithmic plots for the predictions, displayed in Figure 4. Now the network predicts accurately across the entire range of values, with errors well within acceptable range.

The introduction of the sigmoid function in the last layer, forcing the NN to learn only positive values, has thus substantially improved the overall performance of the ML model.

4 Conclusions

We applied a very standard neural network model to the problem of predicting the thermoelectric figure of merit, the ZT value, of a set of synthetic band structures corresponding to idealised unknown test materials. The dataset has been synthetically generated in order to have plenty of data, which allowed investigating the feasibility of such an approach in scenarios where large quantities of real-world data are available. The results have been very promising.

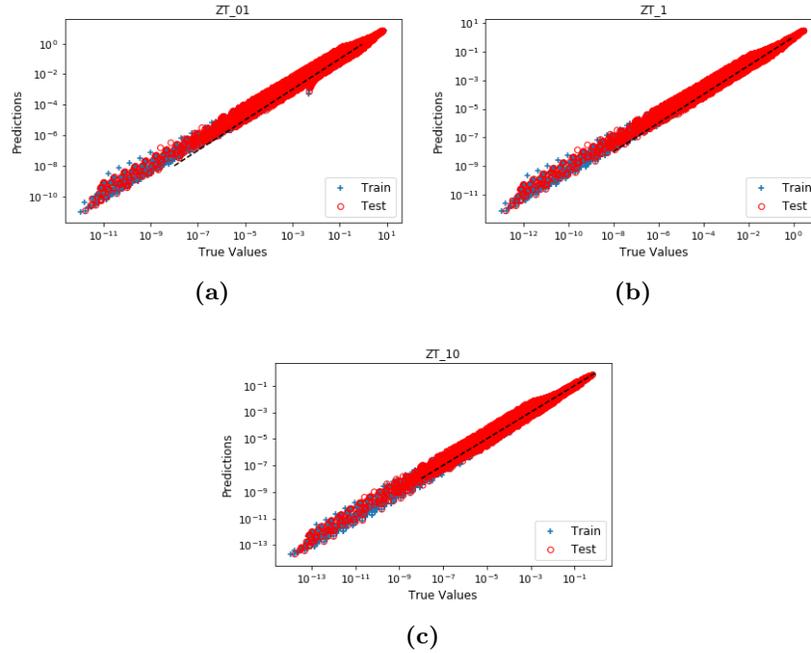


Fig. 4: Test 2. (a) Logarithmic plot of the predictions for $ZT_{0.1}$ on train and test set. (b) Logarithmic plot of the predictions for ZT_1 on train and test set. (c) Logarithmic plot of the predictions for ZT_{10} on train and test set. The black dashed line is the bisector of the first quadrant.

In our first test we predicted the values of ZT using a completely out-of-the-box approach. While the general trend of the figure of merit was correctly identified, the quality of the predictions was somewhat coarse. This applied especially to the smaller values which, while being the least interesting quantities for the application in which this work is concerned, could still be important for other aspects.

In our second test we specialised the neural network to output only positive values. The reason behind this was the physical knowledge that the ZT is a positive number. This change impacted the performance of the model very visibly, gaining accuracy in the predictions across the whole range of materials used for testing.

Our results show that a neural network approach can be a powerful tool in predicting physical properties of materials which would otherwise be very expensive to compute in classical ways, for example through a DFT approach. The application of machine learning techniques can therefore boost the research of novel energy materials by analysing large sectors of the parameters space, i.e. large quantities of chemical compounds, and providing an indication about which materials are promising candidates exhibiting high thermoelectric properties.

One can then perform the expensive, but accurate, DFT laboratory tests only on those promising candidates. Machine learning approaches have therefore a very good potential to greatly increase the efficiency in the use of scientific resources. This simple but significant experiment further shows that it can be beneficial when physical prior knowledge is built-in in a machine learning model.

References

1. United states environmental protection agency. <https://www.epa.gov/chp/chp-benefits>, accessed: 04/04/2019
2. Forman, C., Muritala, I., Pardemann, R., Meyer, B.: Estimating the global waste heat potential. *Renewable and Sustainable Energy Reviews* **57**, 1568–1579 (2016)
3. Gayner, C., Kar, K.: Recent advances in thermoelectric materials. *Progress in Materials Science* **83**, 330–382 (2016)
4. Løvvik, O., Berland, K.: Predicting the thermoelectric figure-of-merit from first principles. *Materials Today: Proceedings* **5**, 10227–10234 (2018)
5. Parr, R.: Density functional theory of atoms and molecules. In: *Horizons of Quantum Chemistry*, pp. 5–15. Springer (1980)
6. Petsagkourakis, I., Tybrandt, K., Crispin, X., Ohkubo, I., Satoh, N., Mori, T.: Thermoelectric materials and applications for energy harvesting power generation. *Science and Technology of Advanced Materials* **19**(1), 836–862 (2018)
7. Schuch, N., Verstraete, F.: Computational complexity of interacting electrons and fundamental limitations of density functional theory. *Nature Physics* **5**(10), 732 (2009)
8. Shi, X., Chen, L., Uher, C.: Recent advances in high-performance bulk thermoelectric materials. *International Materials Reviews* **61**(6), 379–415 (2016)
9. Sirusi, A., Ross, J.: Chapter three - recent nmr studies of thermoelectric materials. *Annual Reports on NMR Spectroscopy*, vol. 92, pp. 137–198. Academic Press (2017)
10. Tabib, M.V., Løvvik, O.M., Johannessen, K., Rasheed, A., Sagvolden, E., Rustad, A.M.: Discovering thermoelectric materials using machine learning: Insights and challenges. In: *Artificial Neural Networks and Machine Learning – ICANN 2018*. pp. 392 – 401 (2018)