

A Neural Modelling Tool for Non-Linear Influence Analyses and Perspectives of Applications in Medical Research

Antonello Pasini ^{1,*}  and Stefano Amendola ²

¹ National Research Council, Institute of Atmospheric Pollution Research, Monterotondo Stazione, 00015 Rome, Italy

² Italian Air Force Mountain Centre, 41029 Sestola, Italy; stefano.amendola@aeronautica.difesa.it

* Correspondence: pasini@iia.cnr.it

Abstract: Neural network models are often used to analyse non-linear systems; here, in cases of small datasets, we review our complementary approach to deep learning with the purpose of highlighting the importance and roles (linear, non-linear or threshold) of certain variables (assumed as causal) in determining the behaviour of a target variable; this also allows us to make predictions for future scenarios of these causal variables. We present a neural tool endowed with an ensemble strategy and its applications to influence analyses in terms of pruning, attribution and future predictions (free code issued). We describe some case studies on climatic applications which show reliable results and the potentialities of our method for medical studies. The discovery of the importance and role (linear, non-linear or threshold) of causal variables and the possibility of applying the relationships found to future scenarios could lead to very interesting applications in medical research and the study and treatment of cancer, which are proposed in this paper.

Keywords: neural networks; non-linearities; causality analyses; climate; medical applications



Citation: Pasini, A.; Amendola, S. A Neural Modelling Tool for Non-Linear Influence Analyses and Perspectives of Applications in Medical Research. *Appl. Sci.* **2024**, *14*, 2148. <https://doi.org/10.3390/app14052148>

Academic Editors:
Giovanni Costantini and
Daniele Giansanti

Received: 9 February 2024
Revised: 28 February 2024
Accepted: 2 March 2024
Published: 4 March 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

When dealing with a complex system—for example, the human body in interaction with its environment—it is often difficult to understand what causes certain changes in the system itself. This is due to the very characteristics of complex systems, which do not have linear causal relationships among their ‘components’ but often have non-linear, threshold or even circular relationships, i.e., with feedback.

For complex systems studied by the so-called ‘hard sciences’, one often finds non-linear equations that link different variables dynamically and thus manage to algorithmically reduce the complexity. In the environmental sciences, doing this is already difficult, and in the medical sciences, it is probably even more so. In many of the latter cases, algorithmic reduction of this complexity must be achieved using data-driven methods.

In any case, once possible causal factors of a change in the system have been identified, these data-driven methods need to be able to capture the non-linearities and feedback that exist in the relationships between these variables. These non-linearities are obviously not ‘seen’ by a linear multiple regression but can be captured, at least implicitly, by non-linear multiple regressions, such as those that can be carried out by neural network (NN) models.

Given this context, much emphasis is now placed on so-called ‘deep learning’, whereby automated machine learning systems use huge available datasets to find regularities in the data and possible causal relationships [1]. This activity has even led to speculation of an ‘end of theory’ [2] and has also been largely applied in medical research (see, for instance [3,4] and references therein). In general, the first sector to see applications of deep learning was clinical image analysis, obviously due to the large amount of data available. Today, deep learning is still widely used with success for this sector in particular, along with the analysis of biosignals, i.e., the set of all the electrical, mechanical and thermal

signals produced by a human body. Moreover, deep learning has recently been applied in biomedicine, the set of disciplines that study molecules involved in biological processes, and in the analysis of so-called 'electronic health records', the standard digital version of the patients' health records, although here with more difficulty due to the heterogeneity of the data. Obviously, by now, all areas of medical science where a large amount of data is available are involved in these deep learning analyses.

However, at least two problems arise here: Firstly, these huge systems (with many hidden layers) are certainly useful for discovering unsuspected relationships between apparently unconnected variables, but, usually, the stability of the systems requires quasi-linear transfer functions. Therefore, the relationships that are found are also quasi-linear (which can hide the real non-linear behaviour). Secondly, in environmental and medical sciences, there are often insufficient amounts of data available to make use of deep learning.

For all these reasons, a few years ago, one of us (A.P.) developed an NN modelling tool for small datasets [5], with an approach that can be considered complementary to that of deep learning: small NNs with just one hidden layer, highly non-linear transfer functions, and the possibility of highlighting non-linear links well. Furthermore, an ensemble strategy has given this tool the characteristics of reliability and robustness in its results.

In this paper, we present some methods of applying this tool (pruning, attribution and prediction), which have already been tried and tested in the climatic and environmental fields (see, e.g., [6–8]), and which make it possible to highlight the importance and roles (linear, non-linear or threshold) of certain causal variables in determining the behaviour of a target variable. This permits predictions of the target variable to be made for future scenarios of the causal variables, too. The tool has already been used in the biomedical field (see the many citations of paper [5], for instance, refs. [9–15]), but not with these specific applications, which can be very useful.

In short, the scope of this paper is to provide and discuss the application methods of a tool that can be used in all those fields in which our theoretical knowledge of the complex system of interest is insufficient for a detailed dynamical study of the causal relationships between different system variables. Obviously, the occurrence of diseases, including cancers, and their relationship to genetic, environmental or lifestyle causal variables is a topic to which this tool can certainly be applied and where it can yield interesting results.

The codes of the tool's applications are freely downloadable: see the Data Availability Statement in the Footnotes.

This paper is structured as follows: Section 2 briefly presents the NN tool in its basic features; in Section 3, the application methods for discovering the importance (linear, non-linear or threshold) of causal variables on climatic and environmental examples are explored; in Sections 4 and 5, respectively, application possibilities in the medical field are discussed and brief conclusions are drawn.

2. The NN Tool

When trying to achieve algorithmic reduction with a statistical model and small datasets, we often run the risk of finding unrealistic and unreliable relationships between causal variables and a variable representing the effect being studied. This is also the case with neural models, where we must avoid overfitting to obtain realistic relationships, and we must minimise the role of variability (due to the random choice of initial weights) to obtain reliable results. The tool presented here seeks precisely to avoid these problems. In the present paper, we describe NNs and this tool conceptually; for more technical details, see [5].

The NN models used in the tool are fairly standard, i.e., feedforward networks with backpropagation training; see [16,17] for the fundamentals of these kinds of networks (multi-layer perceptrons: MLPs). These networks possess just one hidden layer, as it is sufficient to approximate any continuous function [18,19]. The hidden neurons are endowed with hyperbolic tangent transfer functions, while the single output neuron is endowed with a linear one. In the latest applications of the tool, the training technique has been modified

with respect to the original one described in [5]: a further improvement was made for the analysis of small datasets. We now use the Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm [20], which shows better performance than the standard backpropagation. The usual way to obtain a non-linear relationship between inputs and targets (to be approximated by outputs), which is generally valid, is to divide the total dataset into three subsets and apply a training–validation–testing procedure. The connection weights (i.e., the free parameters of the NNs) are fixed on the training set by stopping the training phase when the error on the validation set begins to increase; then, the model performance is measured on a third set, unknown to the NNs—the so-called test set. In our tool, developed for analysing small datasets, a generalised leave-one-out procedure is adopted for the phases of training, validation and testing.

This procedure is depicted in Figure 1. In our complete dataset, formed by a number of input–target pairs (patterns), we extract a single input–target pattern and consider it as the test set. Then, a validation set is randomly chosen, and the remaining patterns constitute the training set. At the end of all the epochs (iterations) of an NN run, the model is able to fix the connection weights, and a transfer function from inputs to output is obtained. But this result can be influenced by the specific random choices of the initial weights and of the members of the validation set. Thus, the NN tool performs multiple runs in an ensemble approach by choosing different random values for weights and members of the validation set. This allows us to calculate the ensemble mean of the outputs and ‘average away’ the intrinsic variability of the NN results.




Figure 1. Sketch of the generalised leave-one-out procedure (from [7]).

After this reconstruction of one value of the target, the procedure is followed again for all the other patterns; each of them becomes, sequentially, the test set. As a final result, at the end of this generalised leave-one-out procedure, we are able to achieve the estimation of all output values and ensemble means.

Inside this tool, standard multilinear regressions are also performed to form a benchmark, which permits us to evaluate the specific value added by the use of NNs in the analyses. In these linear analyses, we adopt the same approach to training: for each single

pattern, the coefficients of the linear regression are fixed on the other data (the union of the training and validation sets used in the NN method). Thus, the linear model can fix its coefficients on more data (this is an advantage with respect to the NN model); as a result, and as we show in the following, the better performance of the NNs is even more notable. Obviously, in the linear model, no ensemble strategy is required.

3. Environmental Applications of the Tool and Results

As mentioned earlier, neural network models are now widely used in the environmental sciences, and our tool in particular has been mainly applied in the climate field. Only three examples are given here to illustrate some of the application methods. These make it possible to identify the importance/role (linear, non-linear or threshold) of the causal variables in determining the behaviour of the target variable and to obtain predictions of the value of the target variable for future scenarios of the causal variables.

3.1. Importance and Roles of the Causal Variables (Pruning)

The first example we present follows the discussion in [7]. There, the problem was to understand the importance and role of climatic influences on human migration from the Sahel Belt to Italy in the period 1995–2009.

The data available for this analysis in our dataset were the annual flows of migrants from the ten Sahel countries to Italy (our target to be reconstructed by means of the NN model using climate data) and the data of direct and indirect climatic influences (our inputs). In particular, the latter were the annual mean temperature, annual mean cumulated precipitation, number of hours with a temperature higher than 30 °C during the growing season of crops, and, finally, cereal yields, which represent an indirect climate effect, as they are themselves influenced by the climate. We built a dataset composed of the annual patterns of four inputs and one output for each country.

The application of our tool led to several interesting results. Firstly, the reconstruction of migration data (using our non-linear neural model and climatic data as inputs) led to an explained variance close to 80% in terms of the ensemble mean performance, whereas the analogous linear multiple regression stopped at just over 60% (see Table 1 below). In Figure 2, we show a graphical view of the NN outputs in terms of the logarithm of the migration rate, defined as the annual migration flow from the origin country to Italy divided by the population of the country of origin in the same year. Given the lack of more social variables in our dataset that may have an influence on migration, the neural result surely looks good. Furthermore, the big difference in the explained variances of the NNs and multilinear regression already shows that the relationship between our inputs and the target is highly non-linear.

Table 1. Performance (in terms of R^2) of migration rate estimation by NNs (ensemble means) and multilinear regressions for the complete runs and for pruned models (from [7]).

Inputs	Target	NN (R^2)	Multilinear (R^2)
Prec.—temp.—# hours $T > 30$ °C—yield	ln(MigRate)	0.775	0.626
Prec.—temp.—# hours $T > 30$ °C	ln(MigRate)	0.671	0.611
Prec.—temp.—yield	ln(MigRate)	0.683	0.632
Prec.—# hours $T > 30$ °C—yield	ln(MigRate)	0.361	0.085
Temp.—# hours $T > 30$ °C—yield	ln(MigRate)	0.715	0.447

But is it possible to quantify the role of the individual input variables in the reconstruction of the target? In a linear model, their weight can be identified through the coefficients associated with the individual variables in the multiple regression. In a neural model, following this path—via the evaluation of the weights associated with the connections—is much more complicated and often does not lead to reliable results. Therefore, in our tool's applications, we follow a different path: that of pruning (or attribution, see later).




Figure 2. Modelling migration rates (black line) from ten countries of Sahel to Italy in the period 1995–2009 by NN models (red lines) and ensemble means (blue line) (from [7]).

Pruning means that we exclude the input variables in turn from the input layer and run the NN models again with more little networks, in order to test which inputs have a dominant role in achieving a correct estimate of the observed target data. The greater the decrease in performance when we extract a variable from the input, the more that variable ‘drives’ the behaviour of the target. In the meantime, by also pruning the linear model, we are able to investigate and identify specific linear and non-linear influences, too. If, when we remove a variable, the performances of the NNs and multilinear regressions become similar (and those of the multilinear model possibly change very little), then that variable plays an important non-linear role. Table 1 shows an example of the performance results of unpruned and pruned networks in the migration study (here, for simplicity, we adopted only R^2 as a measure of performance, but in the original papers, we also used others, e.g., the receiver operating characteristic—ROC—curve).

In this case, temperature is clearly the most influential variable on the migration rate, due to a strong decrease in the performance of both the NN and linear models when it was excluded from the input layer. The decrease in performance was less marked when other variables were extracted from the input layer. However, it is worthwhile to note that harvest yields and the number of hours with $T > 30\text{ }^\circ\text{C}$ have a clear non-linear role. In the case of their exclusion from the inputs, in fact, the performance of the multilinear model changed very little, and the performance of the NN model became very close to that of the linear one.

Of course, in each field of application, it is up to the researchers applying this tool to extract knowledge from these statistical results. In this case, for instance, it is clear that the average annual temperature strongly influences migration, but how? Another analysis of the same study shows that this mean temperature does not affect yields very much; therefore, its large observed increases could directly affect humans and animals. In this respect, it is very important to have noted the strong non-linear role of the hours of temperature exceeding $30\text{ }^\circ\text{C}$ during the harvest period, because this variable is a proxy for the duration of a climatic phenomenon—the heat wave—so that there may be a threshold effect on crops, and also on the heat tolerance of humans and animals, induced by these heat waves.

Studies are now in progress to better identify the values of the threshold in variables showing non-linear relationships with the target.

3.2. Importance and Roles of the Causal Variables (Attribution)

We now describe an alternative way of understanding the role of the individual causal variables on the target variable, which also allows simulations of the latter’s behaviour if the former have different values and trends. These simulations have been carried out for climate attribution studies, and we will use one of them here to exemplify the method; see [6] for further details.

A very important problem in recent climate research has been to understand what has caused the rapid rise in the global mean temperature over the last 60 years. To do this, dynamical models (global climate models—GCMs) have been used, which have shown that without the increase in CO₂ and other greenhouse gases (GHGs), recent global warming would not have occurred. But, given the complexity of these models, it was important to analyse this problem with very different (data-driven) models, such as neural networks and Granger causality models, to test the robustness of their results (see [21,22] and references therein).

As far as NN models are concerned, the concept with which we applied our tool to attribution studies is very simple. At the end of training, our networks found the best weights to reconstruct the trend of the target (over time) according to the input values that occurred in reality. Once these values have been fixed, we have thus found a ‘law’ that reconstructs the links between the inputs and the target, a law that we can consider valid in general. We can therefore apply it to values of the inputs other than those observed in historical reality, thus simulating different situations that could have occurred and obtaining the related values of the output variable. For instance, if we input to the network with fixed weights a constant or stationary value for one variable and real observed values for all the others, we will see the effect of real variations in the first variable on the target variable and will thus be able to understand its role in the evolution of the target (see the following example).

In the climate attribution example given in [6], the target was the global mean temperature, and the causal variables (expressed in terms of radiative forcings) were external natural influences (Sun and volcanoes) and anthropogenic ones (GHGs, black carbon (BC) and sulphate aerosols). In the latter forcings, GHGs and BC have a warming effect, while sulphates have a cooling one. Each of these causal variables has had a certain trend over time since 1850, and we estimated this from observations. But what would have happened if one of these trends had been different, e.g., constant or stationary? A first answer is in Figure 3, taken from [6].




Figure 3. Warming and cooling roles of the distinct anthropogenic forcings by an NN attribution experiment. Black line = observed T; red line = reconstructed T by all inputs with actual values; blue, light blue and green lines = attribution runs with constant forcing(s) (from [6]).

In terms of the ensemble means of our NN runs, Figure 3 shows the very good performance of our NN models in reconstructing the real behaviour of temperature when the real trends of all forcings were considered (see how the black and red lines are close). The light blue line shows the further increase in temperature that would have occurred over the last century if sulphates had remained constant at their 1850 value. The blue and green lines come from runs where constant values for the inputs of GHGs and GHGs + BC radiative forcings, respectively, were propagated along the network with fixed weights. This last result shows that without the increase in anthropogenic warming factors, the recent global warming would not have occurred, so humans are responsible for this warming. This result corroborates the analogous one coming from GCMs.

But what about the influence of observed past changes in solar forcing on the global temperature? The solar forcing always showed its typical eleven-year cycle but with an additional increase in power in the first decades of the last century, which continued into the last decades. We can simulate what would have happened if there had not been this increase by feeding our fixed-weight neural networks with a synthetic series of solar forcing (RFSOLSTAT) that is stationary on the values of the first decades of our time series. The results can be seen in Figure 4. Here, the increasing trend in temperature during the years 1910–1945 was not caught with stationary power from the Sun, and during 1945–1975 the hiatus temperature was about 0.2 K lower than observed. The hypothesis of a ‘stationary Sun’ has no detectable influence on the value of the increasing temperature trend of the last period; this is a new proof that it has been more deeply influenced by anthropogenic forcings.




Figure 4. Attribution experiment for solar influence. Black line = observed T, red lines = results of ensemble runs, blue line = ensemble mean. We mimicked stationary solar forcing by using RFSOLSTAT as an input for the validated NN models (from [6]).

In short, attribution studies allow us to examine the role of a causal variable by decoupling it from the others in a non-linear model that is able to correctly reconstruct the observed trend of the target variable.

3.3. Predictions in Future Scenarios

Once we have found a general transfer function from the inputs to the target and fixed the network weights, we have seen in attribution studies how we can simulate the

behaviour of the target variable for different values of the inputs, e.g., by placing a constant or stationary input together with the others with observed values.

At this point, however, nobody forbids exploiting the ‘law’ found through the neural model to simulate the behaviour of the target variable in very different situations, in ‘scenarios’ in which all input variables change their values. In particular, it is interesting to do this when one can hypothesise future changes in the trends of the input variables and wants to analyse what will happen to the values of the target variable.

This approach was recently followed in a climate application [8]. We will therefore refer to this example to briefly illustrate the potential of this method.

In addition to the finding that recent climate changes are largely driven by human influences, the climate also has an its own internal variability, due to natural cycles and oscillations. We are aware of many such examples of natural variability, and even the so-called Atlantic multidecadal oscillation (AMO) has long been considered a manifestation of natural variability. Recently, however, studies have shown how AMO may have been strongly influenced by certain anthropogenic forcings, such as sulphate aerosols.

These studies were always carried out using GCMs. In a recent study [6], however, we carried out an attribution analysis of AMO using our neural tool. This paper significantly contributed to the current debate on the natural or anthropogenic origin of AMO. The results showed how the AMO’s behaviour from 1866 to the present has been strongly influenced by anthropogenic forcings, with a particular emphasis on the radiative forcing of human-origin sulphate aerosols.

Now, as is well known, there are future scenarios for anthropogenic forcings, such as the SSP scenarios of the Coupled Model Intercomparison Phase 6 (CMIP6). We used the scenarios from one of these models [23] to make future predictions of the AMO’s behaviour by feeding these data into the networks that had best reconstructed the AMO’s historical oscillations. Future data of natural forcings were statistically extrapolated, with a couple of different options for solar radiative forcing; see [8].

The key finding of this study revealed that in all future scenarios, the AMO is expected to lose its cyclical pattern (see Figure 5). It could only be preserved in a scenario with highly unrealistic values for sulphate aerosols.




Figure 5. NN ensemble mean projections of the future behaviour of the AMO index under four SSP scenarios (SSP1-2.6, blue line; SSP2-4.5, yellow line; SSP3-7.0, brown line; SSP5-8.5, red line) in the case of high-power solar radiation hypothesised for the future (from [8]). Black line represents the historical behaviour of the AMO index.

This study showed how, by incorporating future scenarios for inputs, one can predict the behaviour of the target variable and come to very important conclusions. If these scenarios then depend on our actions, we can act to avoid the most damaging scenarios

(in the context of climate, think of the so-called ‘business as usual’ scenarios for future global warming).

4. Possible Applications in the Medical Field

The main application of our tool in medical research is probably in the search for robust non-linear relationships between genetic, environmental or lifestyle factors and the development of a disease or cancer. The pruning or attribution methods described above could make it possible to highlight the roles of various factors and possible non-linearities, e.g., for phenomena affected by thresholds in inputs. Moreover, most of these factors are modifiable; thus, once a non-linear transfer function between the causes (inputs) and a disease index (target) has been found, we can also study what would happen in different scenarios that are possibly conceivable for the future in the case of changes in lifestyles or, for example, in different pollution or climate scenarios.

Another case where our tool could be very useful is in the study of drug–disease interaction, especially in the very critical cases of cancer treatments, the success of which is often measured in the years of survival after surgery. Cases observed in patients with certain physical characteristics and treatments actually given (as inputs) using our tool could lead to reliable reports for the reconstruction of their survival years (the possible target). At this point, numerical experimentation using NNs could be conducted to see which treatments (and doses) are preferable for other patients with different characteristics in the input data, in order to increase their survival time. This can be done by the application of our forecasting strategy.

In general, neural modelling can analyse complex interactions between genetic, environmental and lifestyle factors to predict an individual’s cancer risk and estimate their prognosis. These models can capture non-linear relationships that traditional statistical methods may miss. At present, particular attention is paid to deep learning and its contribution to cancer research; see, for instance [24]. Here, however, we would like to briefly cite some research sectors in which our neural approach to small datasets can benefit cancer studies.

Neural modelling can assist in personalising cancer treatment by analysing patient-specific data, including genetic profiles and treatment responses. This approach can help to identify the most effective treatment strategies for individuals, for instance, in terms of personalised chemotherapy selection (see [25] for a problem of this type).

Neural networks can also analyse non-linear interactions within the tumour microenvironment, including interactions between cancer cells, immune cells and the extracellular matrix. This can shed light on tumour behaviour and immune responses, an important relationship which needs to be deepened; see [26] for a recent review on this topic.

Neural modelling can predict the non-linear responses of cancer cells to various drugs. Considering the complexity of drug–target interactions, these models can optimise drug selection for individual patients, including in cases of the use of “non-oncology drugs”, such as in [27].

Another field in which neural modelling can be very useful is that of non-linear resistance mechanisms. For instance, our tool can help uncover non-linear resistance mechanisms affecting cancer therapies, including targeted therapies and immunotherapies, by considering the interplay between tumour cells and the immune system (see [28] for a review on this important topic).

Finally, our neural tool performs multiple non-linear regressions, but with a little modification, it could be adapted for classification problems. In this case, it could also help with diagnostic analyses (inputs) and their use for the early detection of cancer (the target). In particular, if databases that are not too small are available, the linear and non-linear influences of certain inputs on the target could be identified in the analysis values. And then, along with understanding the role of individual inputs, one could also apply the ‘law’ found to cases of new patients, identifying the eventual onset of cancer. In this manner, a neural network tool, such as our one, can also analyse multi-modal data, including imaging,

genomic and clinical data, to improve the accuracy of early cancer detection and screening programs. Attempts have been made using deep learning (see, for instance [29]), but we are confident that an accurate preselection of inputs and the use of our robust tool may contribute to this field.

5. Conclusions

In this paper, we presented a neural network tool suitable for non-linear influence analyses in complex systems with small available datasets (as is often the case with environmental and health systems), together with its application methods. The latter were concretely exemplified in some climatic cases, but we believe that the methods discussed here of pruning, attribution and prediction can also be of great help for medical research and, in particular, cancer research, as we discussed in the previous section.

In particular, the discovery of the importance and role (linear, non-linear or threshold) of causal variables and the possibility of applying the relationships found to future scenarios could lead to very interesting applications in health sciences and the study and treatment of cancer.

We are, of course, aware that there are other methods, such as sensitivity analyses and generalised additive models (GAMs), that do similar things to our tool. However, what we consider ‘innovative’ here is the set of features of our tool, including ensemble runs, attribution analysis and forecasting activities for small available datasets. Obviously, we also know that the neural networks we used are quite simple (MLPs), and it is our intention in the future to include more powerful networks in this tool. In this perspective article, we only wanted to show the application potential of the current tool in the medical field.

Furthermore, our method should obviously be tested against the standard models used in medical research, and this will be done in our future works.

Finally, we hope that this perspective article can induce more researchers to adopt promising NN tools and their application methods in their influence analyses.

Author Contributions: Conceptualisation, A.P.; methodology, A.P. and S.A.; software, S.A.; validation, A.P. and S.A.; formal analysis, S.A.; investigation, A.P. and S.A.; data curation, S.A.; writing—original draft preparation, A.P.; writing—review and editing, A.P. and S.A.; supervision, A.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: We have made the Matlab[®] codes of our tool for the 3 applications discussed in the text (pruning, attribution and prediction) freely available. The relevant program files, together with the 3 datasets used for the related climate applications and a function called in the main programs, can be found at https://cnrsc-my.sharepoint.com/:f/g/person/antonello_pasini_cnr_it/Eu4DxbFPD2dPu1pZINdct_IBxgkTLMRLHNTgeUAxVkf5dg?e=7a9kvN (accessed on 1 February 2024) (see the folder ‘Free-codes&data’).

Acknowledgments: We warmly acknowledge L. Bertolaccini for helpful discussions.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Alzubaidi, L.; Zhang, J.; Humaidi, A.J.; Al-Dujaili, A.; Duan, Y.; Al-Shamma, O.; Santamaría, J.; Fadhel, M.A.; Al-Amidie, M.; Farhan, L. Review of deep learning: Concepts, CNN architectures, challenges, applications, future directions. *J. Big Data* **2021**, *8*, 53. [CrossRef]
2. Anderson, C. The End of Theory: The Data Deluge Makes the Scientific Method Obsolete. *Wired*. 23 June 2008, p. 16. Available online: <https://www.wired.com/2008/06/pb-theory/> (accessed on 2 February 2024).
3. Egger, J.; Gsaxner, C.; Pepe, A.; Pomykala, K.L.; Jonske, F.; Kurz, M.; Li, J.; Kleesiek, J. Medical deep learning—A systematic meta-review. *Comput. Methods Prog. Biomed.* **2022**, *221*, 106874. [CrossRef]
4. Piccialli, F.; Di Somma, V.; Giampaolo, F.; Cuomo, S.; Fortino, G. A survey on deep learning in medicine: Why, how and when? *Inform. Fusion* **2021**, *66*, 111–137. [CrossRef]

5. Pasini, A. Artificial neural networks for small dataset analysis. *J. Thorac. Dis.* **2015**, *7*, 953–960.
6. Pasini, A.; Racca, P.; Amendola, S.; Cartocci, G.; Cassardo, C. Attribution of recent temperature behaviour reassessed by a neural-network method. *Sci. Rep.* **2017**, *7*, 17681. [[CrossRef](#)]
7. Pasini, A.; Amendola, S. Linear and nonlinear influences of climatic changes on migration flows: A case study for the ‘Mediterranean bridge’. *Environ. Res. Commun.* **2019**, *1*, 011005. [[CrossRef](#)]
8. Pasini, A.; Amendola, S.; Federbusch, E. Is natural variability really natural? The case of Atlantic Multidecadal Oscillation investigated by a neural network model. *Theor. Appl. Clim.* **2022**, *150*, 881–892. [[CrossRef](#)]
9. Lualdi, M.; Cavalleri, A.; Battaglia, L.; Colombo, A.; Garrone, G.; Morelli, D.; Pignoli, E.; Sottotetti, E.; Leo, E. Early detection of colorectal adenocarcinoma: A clinical decision support tool based on plasma porphyrin accumulation and risk factors. *BMC Cancer* **2018**, *18*, 841. [[CrossRef](#)]
10. Bagher-Ebadian, H.; Janic, B.; Liu, C.; Pantelic, M.; Hearshen, D.; Elshaikh, M.; Movsas, B.; Chetty, I.J.; Wen, N. Detection of Dominant Intra-prostatic Lesions in Patients With Prostate Cancer Using an Artificial Neural Network and MR Multi-modal Radiomics Analysis. *Front. Oncol.* **2019**, *9*, 1313. [[CrossRef](#)]
11. Imai, S.; Takekuma, Y.; Kashiwagi, H.; Miyai, T.; Kobayashi, M.; Iseki, K.; Sugawara, M. Validation of the usefulness of artificial neural networks for risk prediction of adverse drug reactions used for individual patients in clinical practice. *PLoS ONE* **2020**, *15*, e0236789. [[CrossRef](#)]
12. Izonin, I.; Tkachenko, R.; Fedushko, S.; Koziy, D.; Zub, K.; Vovk, O. RBF-Based Input Doubling Method for Small Medical Data Processing. In *Advances in Artificial Systems for Logistics Engineering*; Hu, Z., Zhang, Q., Petoukhov, S., He, M., Eds.; Springer: Cham, Switzerland, 2021; pp. 23–31.
13. Sohn, J.; Jung, I.-Y.; Ku, Y.; Kim, Y. Machine-Learning-Based Rehabilitation Prognosis Prediction in Patients with Ischemic Stroke Using Brainstem Auditory Evoked Potential. *Diagnostics* **2021**, *11*, 673. [[CrossRef](#)]
14. Le, T.-D.; Noumeir, R.; Rambaud, J.; Sans, G.; Jouvét, P. Detecting of a Patient’s Condition From Clinical Narratives Using Natural Language Representation. *IEEE Open J. Eng. Med. Biol.* **2022**, *3*, 142–149. [[CrossRef](#)]
15. Sebros, R.; De la Garza-Ramos, C.; Peterson, J.J. Detecting whether L1 or other lumbar levels would be excluded from DXA bone mineral density analysis during opportunistic CT screening for osteoporosis using machine learning. *Int. J. Comput. Assist. Radiol. Surg.* **2023**, *18*, 2261–2272. [[CrossRef](#)]
16. Hertz, J.; Krogh, A.; Palmer, R.G. *Introduction to the Theory of Neural Computation*; Addison-Wesley: New York, NY, USA, 1991.
17. Bishop, C.M. *Neural Networks for Pattern Recognition*; Oxford University Press: New York, NY, USA, 1995.
18. Cybenko, G. Approximation by superposition of a sigmoidal function. *Math. Control Signal Syst.* **1989**, *2*, 303–314. [[CrossRef](#)]
19. Hornik, K.; Stinchcombe, M.; White, H. Multilayer feedforward networks are universal approximators. *Neural Netw.* **1989**, *2*, 359–366. [[CrossRef](#)]
20. LeCun, Y.; Bottou, L.; Orr, G.B.; Müller, K.-R. Efficient BackProp. In *Neural Networks: Tricks of the Trade*; Orr, G.B., Müller, K.-R., Eds.; Springer: New York, NY, USA, 1998; pp. 9–50.
21. Knutti, R.; Baumberger, C.; Hirsch Hadorn, G. Uncertainty Quantification Using Multiple Models—Prospects and Challenges. In *Computer Simulation Validation. Simulation Foundations, Methods and Applications*; Beisbart, C., Saam, N., Eds.; Springer: Cham, Switzerland, 2019; pp. 835–855.
22. Walmsley, L.D. The strategy of model building in climate science. *Synthese* **2021**, *199*, 745–765. [[CrossRef](#)]
23. Leimbach, M.; Bauer, N.; Baumstark, L.; Edenhofer, O. Mitigation costs in a globalized world: Climate policy analysis with REMIND-R. *Environ. Model. Assess.* **2010**, *15*, 155–173. [[CrossRef](#)]
24. Howard, F.M.; Pearson, A.T. Prognosis and Treatment of Non-Small Cell Lung Cancer in the Age of Deep Learning. *JAMA Netw. Open* **2020**, *3*, e206368. [[CrossRef](#)]
25. Yu, K.; Sang, Q.-X.A.; Lung, P.-Y.; Tan, W.; Lively, T.; Sheffield, C.; Bou-Dargham, M.J.; Liu, J.S.; Zhang, J. Personalized chemotherapy selection for breast cancer using gene expression profiles. *Sci. Rep.* **2017**, *7*, 43294. [[CrossRef](#)]
26. Binnewies, M.; Roberts, E.W.; Kersten, K.; Chan, V.; Fearon, D.F.; Merad, M.; Coussens, L.M.; Gaborilovich, D.I.; Ostrand-Rosenberg, S.; Hedrick, C.C.; et al. Understanding the tumor immune microenvironment (TIME) for effective therapy. *Nat. Med.* **2018**, *24*, 541–550. [[CrossRef](#)]
27. Corsello, S.M.; Nagari, R.T.; Spangler, R.D.; Rossen, J.; Kocak, M.; Bryan, J.G.; Humeidi, R.; Peck, D.; Wu, X.; Tang, A.A.; et al. Discovering the anti-cancer potential of non-oncology drugs by systematic viability profiling. *Nat. Cancer* **2020**, *1*, 235–248. [[CrossRef](#)]
28. Sharma, P.; Allison, J.P. The future of immune checkpoint therapy. *Science* **2015**, *348*, 56–61. [[CrossRef](#)]
29. Esteva, A.; Kuprel, B.; Novoa, R.A.; Ko, J.; Swetter, S.M.; Blau, H.M.; Thrun, S. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* **2017**, *542*, 115–118. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.