



Flexible, non-parametric modeling using regularized neural networks

Downloaded from: <https://research.chalmers.se>, 2024-04-26 08:39 UTC

Citation for the original published paper (version of record):

Allerbo, O., Jörnsten, R. (2022). Flexible, non-parametric modeling using regularized neural networks. *Computational Statistics*, 37(4): 2029-2047. <http://dx.doi.org/10.1007/s00180-021-01190-4>

N.B. When citing this work, cite the original published paper.



Flexible, non-parametric modeling using regularized neural networks

Oskar Allerbo¹  · Rebecka Jörnsten¹

Received: 3 February 2021 / Accepted: 15 December 2021

© The Author(s) 2021

Abstract

Non-parametric, additive models are able to capture complex data dependencies in a flexible, yet interpretable way. However, choosing the format of the additive components often requires non-trivial data exploration. Here, as an alternative, we propose PrAda-net, a one-hidden-layer neural network, trained with proximal gradient descent and adaptive lasso. PrAda-net automatically adjusts the size and architecture of the neural network to reflect the complexity and structure of the data. The compact network obtained by PrAda-net can be translated to additive model components, making it suitable for non-parametric statistical modelling with automatic model selection. We demonstrate PrAda-net on simulated data, where we compare the test error performance, variable importance and variable subset identification properties of PrAda-net to other lasso-based regularization approaches for neural networks. We also apply PrAda-net to the massive U.K. black smoke data set, to demonstrate how PrAda-net can be used to model complex and heterogeneous data with spatial and temporal components. In contrast to classical, statistical non-parametric approaches, PrAda-net requires no preliminary modeling to select the functional forms of the additive components, yet still results in an interpretable model representation.

Keywords Additive models · Model selection · Non-parametric regression · Neural networks · Regularization · Adaptive lasso

✉ Oskar Allerbo
allerbo@chalmers.se

Rebecka Jörnsten
jornsten@chalmers.se

¹ Mathematical Sciences, University of Gothenburg and Chalmers University of Technology, Gothenburg, Sweden

1 Introduction

Non-parametric, additive models combine the flexibility of non-parametric regression (e.g. splines and smoothers) with the interpretability of an additive model structure (Hastie and Tibshirani 1990). For non-parametric functions $\{f_j\}_{j=1}^r$, the model can be formulated as

$$\mathbb{E}[y] = \sum_{j=1}^r f_j(x_j).$$

For more flexible integration of explanatory variables, Friedman and Stuetzle (1981) proposed projection pursuit regression (PPR), where the non-parametric additive components are formed from a component-specific projection of explanatories:

$$\mathbb{E}[y] = \sum_{j=1}^r f_j(\beta_j^T \mathbf{x}). \quad (1)$$

By allowing for the number of components, r , in Eq. 1 to grow, more and more complex data structures can be captured by the model.

PPR is very flexible, but loses some of the interpretability of the additive model as the projection parameters β_j may combine several covariates into additive components in a way that is often difficult to untangle. While a sparse projection could alleviate this fact, it is not trivial to combine with the selection and flexibility of the non-parametric functions, f_j . In practice, especially for high-dimensional and complex data, it is therefore common to use data exploration, marginal or partial modeling and domain knowledge to pre-define the subsets of covariates (e.g. potential interactions, component content) to combine into additive components, as well as to pre-select the family of functional components (e.g. which type of spline).

Neural networks also have the capacity to learn complex dependencies from data and can capture non-linear relationships between input and output variables, as well as interactions between variables. Contrary to the additive model, but similar to PPR, the results do not easily transfer to an interpretable model, since there will be links between all inputs and outputs through the inner layers. For a one-hidden-layer neural network, the neural network model resembles PPR, with the weight matrix elements (links) corresponding to the projection parameters β_j and an activation function σ that generates $f_j \propto \sigma(\beta_j^T \mathbf{x})$. Now, we recognize that some neural network parameters (links) contribute more to the prediction than others, and there is reason to believe that if the less important links are removed from the network architecture, an interpretable structure would appear, as sketched in Fig. 1. That is, sparse weights (projections) will enhance interpretability. Furthermore, by utilizing *multiple hidden nodes*, corresponding to squares in Fig. 1, to obtain the additive components, a diverse set of data-adaptive f_j , with varying complexity, can be obtained.

A popular approach to eliminate parameters from a model is l_1 -penalization, hereafter referred to as lasso (Tibshirani 1996), which, together with its extensions such as the group lasso (Yuan and Lin 2006) and the adaptive lasso (Zou 2006), has been used in neural networks to reduce complexity and increase interpretability. Group lasso is

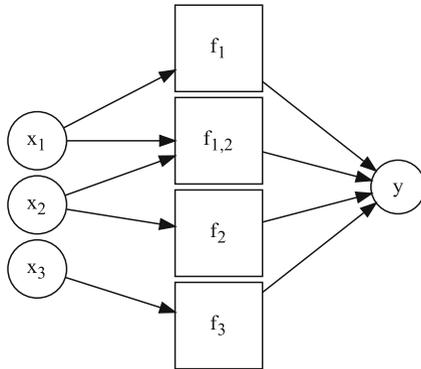


Fig. 1 Desired architecture for an interpretable neural network representing $\mathbb{E}[y] = f_1(x_1) + f_2(x_2) + f_3(x_3) + f_{1,2}(x_1, x_2)$. The circles represent single nodes and the squares represent combinations of nodes

used to penalize all the weights in a predefined group, typically all the inputs or outputs from a node or, in the case of convolutional neural networks, a channel (Scardapane et al. 2017; Wang et al. 2018a; Lym et al. 2019; Zhang et al. 2019; Tank et al. 2017; Ainsworth et al. 2018). An alternative to group lasso is to add extra parameters that are multiplied to all the weights in a group and then apply standard lasso to these parameters (Leclerc et al. 2018; Sun et al. 2017; Wang et al. 2018b; Vaughan et al. 2018). Adaptive lasso, where each parameter is assigned a unique penalty, has been used by Xu and Han (2016) and Qiao et al. (2018) to improve estimation stability, but without focus on interpretability. Works cited above that focus on interpretability are those by Tank et al. (2017), where Granger causality between time series is restrained; Ainsworth et al. (2018), where multiple decoders are connected only to a subset of the latent variables generated by variational autoencoders; and, most relevant for our work, by Wang et al. (2018b), Distillation Guided Routing (DGR) and Vaughan et al. (2018), Explainable Neural Networks (xNN).

In DGR, unimportant channels in a trained convolutional neural network are removed by multiplying the outputs from each channel with a lasso penalized parameter, balancing removal of channels with maintained performance. Since DGR removes entire channels, it can never remove unimportant links from important nodes. In Fig. 1 this corresponds to penalizing only the weights between the f 's and y .

In xNN, parallel neural networks with lasso penalized inputs are used to model each function in an additive model. xNN removes individual links, but only between the inputs and the, still black-box, parallel networks. In Fig. 1, each f -box would correspond to a neural network of pre-defined capacity, where the inputs of the functions are chosen by the model, but where function complexities are pre-defined.

Here, we propose PrAda-net, a novel method that combines adaptive lasso for individual penalties of network parameters with proximal gradient descent. PrAda-net addresses two problems with lasso regularization of neural networks. First, lasso penalizes all model parameters equally, yet it is reasonable to assume that in an over-parameterized neural network some weights contribute more to the final result than others, and should ideally be penalized less. Second, gradient descent algorithms cannot handle the non-differentiability of the lasso penalty at zero, something that has to

be circumvented, e.g. by defining $\frac{\partial|x|}{\partial x}\Big|_{x=0} := 0$ or by using some smooth approximation, such as $|x| \approx \sqrt{x^2 + \varepsilon}$. However, such approximations result in no parameters set exactly to zero, introducing the need for thresholding, which requires tuning as well.

PrAda-net can be thought of as a generalization of xNN, where the number of hidden nodes to realize a function is adaptively chosen, depending on the function complexity. Therefore, it both improves the interpretability and reduces the number of parameters. Compared to DGR, PrAda-net penalizes not only the outputs, but also the inputs, of the functions, resulting in a clear-cut identification of functional components. Compared to non-parametric, additive models, PrAda-net is completely data adaptive and neither the input, nor the form of the functional components, f_j , need to be pre-specified.

The rest of this paper is structured as follows: In Sect. 2, we present the PrAda-net algorithm. In Sect. 3 we illustrate the method on simulated data, comparing the predictive and interpretable performance of PrAda-net to that of standard lasso and DGR. We also apply PrAda-net to U.K. black smoke data. We compare the model components automatically selected by PrAda-net to the large-scale generalized additive model presented by Wood et al. (2017). We show how the method is able to automatically select the network complexity to generate a concise, interpretable architecture, thus making it suitable for non-parametric additive modelling. Finally, in Sect. 4, we discuss our findings and possible future directions for research.

2 Method

In this paper the following notation will be used: The input nodes of a neural network will be denoted by the vector \mathbf{x} , with elements x_j , and the single output node will be denoted by y . Functions realized by hidden nodes of the network will be denoted by $f_j(\mathbf{x}|\boldsymbol{\theta})$ where $\boldsymbol{\theta}$ denotes the vector of network parameters, which will interchangeably be referred to as parameters, weights and links. Usually, $\boldsymbol{\theta}$ will be implicit, writing only $f_j(\mathbf{x})$. A hat over a parameter or a function, e.g. \hat{y} , means that this is the reconstructed parameter or function inferred by the network; when it is clear from the context, the hat will be omitted.

As sketched in Fig. 1, the architecture of a sparse neural network can be interpreted as an additive model, where each function is defined as the set of hidden nodes connected to the same set of input nodes. In order to use lasso to obtain such an architecture, the two limitations stated above have to be overcome. First, lasso penalizes all parameters equally, while ideally, if the true model were known, one would want a very small penalty on the true model parameters, and a large penalty on parameters that should not be included. In the neural network setting, it is better to think of parameters as “relevant” or “important” rather than “true”, but the rationale for data-adaptively eliminating parameters still stands. It was shown by Zou (2006) that under certain conditions, standard lasso is not consistent and as a remedy the adaptive lasso was proposed. In adaptive lasso, each parameter obtains an individual penalty based on its ordinary least square (OLS) estimate, penalizing parameters with small OLS values more.

Second, lasso regularized models trained with gradient descent do not obtain explicit zeroes, which is a consequence of the non-uniqueness of the derivative of the lasso penalty at zero. Proximal gradient descent (Rockafellar 1976), on the other hand, leverages on this non-uniqueness and obtains exact zeros. Like standard gradient descent, proximal gradient descent is an iterative method, with the difference that the non-differentiable parts of the objective function are handled in an additional, proximal, step. If the objective function, $l(\theta)$, can be decomposed into $l(\theta) = g(\theta) + h(\theta)$, where g is differentiable (typically a reconstruction error) and h is not (typically a lasso regularization term), then a standard gradient descent step followed by a proximal gradient descent step is defined as

$$\theta^{t+1} = \text{prox}_{\alpha h}(\theta^t - \alpha \nabla g(\theta^t))$$

where α is a step size and prox is the proximal operator, that depends on h . In the case of lasso with penalty parameter λ , when $h(\theta) = \lambda \|\theta\|_1 = \lambda \sum_j |\theta_j|$, the proximal operator decomposes component-wise and is, with $\eta_j := (\theta_j - \alpha \nabla g(\theta_j))$ denoting the output of the standard gradient descent step,

$$\text{prox}_{\alpha h}(\eta_j) = \text{sign}(\eta_j) \cdot \max(|\eta_j| - \alpha \lambda, 0).$$

I.e., each η_j is adaptively shrunk towards zero, and once it changes sign it is set to exactly zero.

2.1 The PrAda-net algorithm

PrAda-net builds on the assumption that once a neural network is trained, even if all nodes in two subsequent layers are connected, some of the weights are more important than others, and that these important weights are larger in absolute value than the less important ones. Thus, adding an adaptive lasso penalty, i.e. an individual penalty based on the current value of the weight, and continue training will penalize the small, unimportant, weights more than the important ones, ideally setting all unimportant weights to zero, while penalizing the important weights less, thus reducing bias. Adding the adaptive lasso penalty changes the objective function as below:

$$\underset{\theta}{\text{argmin}} l(\mathbf{x}|\theta) \implies \underset{\theta}{\text{argmin}} \left(l(\mathbf{x}|\theta) + \lambda \sum_j \frac{|\theta_j|}{|\hat{\theta}_j|^\gamma} \right), \tag{2}$$

where $\hat{\theta}_j$ denotes the value of the weight just before adding the penalty and $\gamma > 0$ is a tunable parameter that we here set to 2, which is common practice in adaptive lasso (Zou 2006).

In order to get exact zeros one would like to use proximal gradient descent for this phase of the training. However, proximal gradient descent is not compatible with momentum based optimization algorithms, such as Adam (Kingma and Ba 2014). Furthermore, the choice of the step size α , is not as obvious as in Adam, where the

default parameters have been shown to work very well in most cases. We circumvent these issues by training PrAda-net in three stages: In the first stage, the network is trained with a standard optimizer, but without the lasso penalty, until convergence or some other stopping criterion. In the second stage, the adaptive lasso penalty is turned on, using the same optimizer as in the first stage, i.e. not utilizing proximal gradient descent, thus obtaining some weights close to, but distinct from, zero. Before the third stage the adaptive lasso penalty is updated based on the current weight values and then proximal gradient descent is used. Since some weights are now very close to zero, even with small values of α and λ , these weights will be heavily penalized and proximal gradient descent converges in relatively few iterations. For stage three, we thus used 10^{-5} for both α and λ , since the algorithm provided a stable behaviour around this value. PrAda-net is summarized in Algorithm 1. The more crucial selection of tuning parameters in stage 2 is discussed separately in Sect. 2.2.

Algorithm 1 PrAda-net

- 1: Train the neural network until convergence or other stopping criterion, with optimizer of choice.
 - 2: Add an adaptive lasso penalty to each weight according to Eq. 2. Continue training with the chosen optimizer.
 - 3: Update the adaptive lasso penalty. Continue training with proximal gradient descent.
-

2.2 Choosing the lasso penalty parameter

Choosing the optimal regularization parameter, λ^* , is a trade off between interpretability and model performance. When performance is the sole priority, the optimal regularization is often chosen so that it minimizes the mean test error across multiple splits of the data into training and test sets (cross-validation). Here, we are willing to sacrifice some performance to gain interpretability and use a similar approach as in standard lasso regression packages (Friedman et al. 2010). We thus choose our λ^* as the largest one of those whose resulting mean test error is within two standard deviations of the minimum test error across all λ . Figure 2c in Sect. 3.1 illustrates this.

2.3 Choosing stopping criterion

The most obvious stopping criterion for all three stages in PrAda-net is to train until the performance on the test set stops improving. However, it is known that early stopping can work as a regularizer, which is sometimes desirable. Nonetheless, our experience is that using PrAda-net with early stopping is inferior compared to training until convergence. Too early stopping in the first stage tends to result in $\hat{\theta}_j$ in Eq. 2 to exhibit an insufficient spread of magnitude for the adaptive lasso penalty have an impact. This thus reduces the ability of PrAda-net to generate compact model components. Too early stopping in one of the two last stages tends to result in unnecessarily complex models, using more hidden nodes than necessary. In all our experiment full convergence was used as the stopping criterion.

2.4 Identifying linear functions

As activation function for the hidden nodes, tanh was used. The slope of tanh varies smoothly between 0 and 1 on its support, which makes it suitable to approximate an unknown, smooth function. Another popular activation function is the rectified linear unit, $\max(0, x)$. In contrast to tanh, this function is piecewise linear, and using it for approximating a smooth function would in general require more hidden nodes, or a rougher approximation, than when using tanh.

However, as can be seen from its Maclaurin expansion $\tanh(x) = x + \mathcal{O}(x^3)$, is almost linear if x is small in absolute value. For x_1 close to zero, $\tanh(x_1 + x_2) \approx x_1 + \tanh(x_2)$, and therefore, if the output is linear in some input variable (x_1 in this example), no separate hidden node is needed to model this linear function. Instead, the function can be incorporated into some other node, thus increasing compactness but decreasing interpretability. Ideally, a separate penalty or regularization technique should be included to avoid this, but since it is far from obvious how this penalty should be formulated while retaining computational feasibility, instead the following postprocessing step was added: For each identified function, the partial derivatives with respect to all its inputs were calculated and if, for some input, the derivative was close to constant, i.e. with variance lower than some σ_{\max}^2 , that input was removed from the function and put into a new, linear, function, realized by a new node. σ_{\max}^2 was calibrated on synthetic data to 0.01. Since for standardized data, σ_{\max}^2 translates between data sets, this value was consistently used throughout the paper.

3 Experiments

All experiments in this section were done on neural networks with one hidden layer, with tanh as activation function, and one output node with a linear activation function. The mean squared error was used as the loss function and the Adam optimizer was used in the two first stages of the PrAda-net algorithm. The data was randomly split into 90 % training and 10 % testing data. When choosing λ^* , 20 random splits of the data were made for each λ value. In order to escape suboptimal minima, five different initializations were used for each training. All computations were done on a NVIDIA V100, 32GB, GPU.

Three experiments were performed. In the first experiment, synthetic data was generated from the first four Legendre polynomials; in the second experiment, we used a data set with black smoke levels in the U.K. recently analyzed by using a large-scale generalized additive model by Wood et al. (2017); and in the third experiment synthetic data was generated from the model inferred by PrAda-net on the U.K. black smoke data.

3.1 Legendre polynomials

Our first experiment was done on the sum of the first four Legendre polynomials, which are orthogonal on the interval $[-1, 1]$ and are given by

$$\begin{aligned} P_1(x) &= x \\ P_2(x) &= \frac{1}{2}(3x^2 - 1) \\ P_3(x) &= \frac{1}{2}(5x^3 - 3x) \\ P_4(x) &= \frac{1}{8}(34x^4 - 30x^2 + 3). \end{aligned} \quad (3)$$

For the simulations, five random variables, x_1, \dots, x_5 , were generated from $\mathcal{U}[-1, 1]$, each with 1000 realizations, resulting in a 1000×5 x -matrix, from which a 1000×1 y -matrix was created according to

$$y = P_1(x_1) + P_2(x_2) + P_3(x_3) + P_4(x_4) + \varepsilon \quad (4)$$

where $\varepsilon \sim \mathcal{N}(0, 0.1^2)$ is added noise. Note that y is independent of x_5 . A neural network with five input nodes and 50 hidden nodes was trained by applying Algorithm 1 with λ^* chosen as explained in Sect. 2.2.

The results are shown in Fig. 2. As seen in 2a, out of the original 50 hidden nodes, only 10 are used in the final architecture and the hidden units are split into four different functions, each with only one x_j as input. Figure 2b shows the reconstructed $\hat{P}_j(x)$'s (left-most panels) and also how the model decomposes each $\hat{P}_j(x)$ into a sum of up to four subfunctions, each represented by one node in the hidden layer. Figure 2c demonstrates the penalty parameter selection for this simulation.

We compared the test error and the variable importance, measured using saliency maps (Simonyan et al. 2013) for PrAda-net, standard lasso and DGR. To make standard lasso more competitive, we did not apply lasso regularization from the start. Instead, we utilized the PrAda-net algorithm with the same penalty for all links, i.e. $\gamma = 0$ in Eq. 2. DGR was adapted to feedforward regression networks by looking at nodes instead of channels and by using the squared error instead of the cross-entropy loss.

For saliency maps, the importance of variable j , evaluated at \mathbf{x} , is defined as $I_j(\mathbf{x}) := \left| \frac{\partial f(\mathbf{x})}{\partial x_j} \right|$, which in our case reduces to $I_j(\mathbf{x}) = \left| \frac{\partial P_j(x_j)}{\partial x_j} \right|$. Since for non-linear functions the value of the derivative depends on \mathbf{x} , in order to get a global importance measure, we averaged over all \mathbf{x} 's in the test set \mathcal{X} , $I_j := \frac{1}{|\mathcal{X}|} \sum_{\mathbf{x} \in \mathcal{X}} I_j(\mathbf{x})$. This can be seen as a Monte Carlo integral, meaning that as $|\mathcal{X}| \rightarrow \infty$, $I_j \rightarrow \frac{1}{2} \int_{-1}^1 \left| \frac{\partial P_j(x_j)}{\partial x_j} \right| dx_j$, where $\frac{1}{2}$ is the probability density function of the uniform distribution that x_j is sampled from. Thus, for large sample sizes, the true value of the variable importance is given by the analytical integral.

50 runs with different noise realizations were performed with PrAda-net penalized with λ^* , and the other two models penalized to equal complexity. In the case of standard

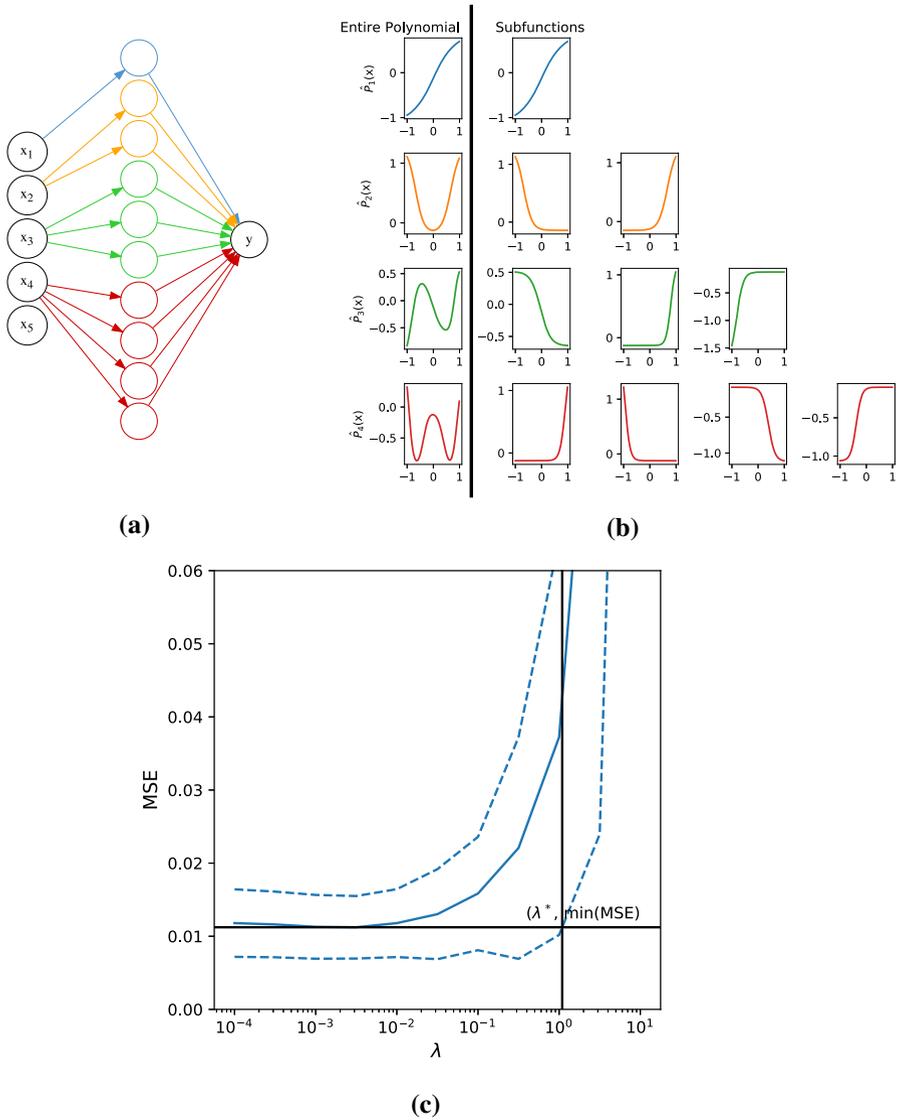


Fig. 2 Identifying the sum of the four first Legendre polynomials. **a** shows the inferred reduced neural network, with the hidden nodes and links color coded according to the identified functions. In **b**, the leftmost column shows, from top to bottom, $\hat{P}_1(x_1)$, $\hat{P}_2(x_2)$, $\hat{P}_3(x_3)$ and $\hat{P}_4(x_4)$, while the other columns show their decompositions into the subfunctions realized by single hidden nodes. **c** shows the mean \pm two standard deviations of the test error obtained by 20 bootstrap runs for each λ . λ^* is chosen as the value where the lower interval intersects the lowest mean value

Table 1 Legendre polynomials: Mean and one standard deviation of test error and estimated variable importance ($\frac{1}{|\mathcal{X}|} \sum_{\mathcal{X}} \left| \frac{\partial f(x)}{\partial x_j} \right|$) together with the true asymptotic variable importance ($\frac{1}{2} \int_{-1}^1 \left| \frac{\partial P_j(x_j)}{\partial x_j} \right| dx_j$)

	True Var.Imp.	PrAda-net	Standard Lasso	DGR
Test error	–	0.04 ± 0.02	0.13 ± 0.03	0.16 ± 0.03
x_1	1	0.75 ± 0.33	0.69 ± 0.43	0.63 ± 0.09
x_2	1.5	1.26 ± 0.09	1.21 ± 0.05	0.65 ± 0.12
x_3	1.89	1.42 ± 0.16	0.96 ± 0.11	0.32 ± 0.12
x_4	2.23	1.72 ± 0.3	0.43 ± 0.24	0.25 ± 0.08
x_5	0	0.0 ± 0.0	0.01 ± 0.01	0.07 ± 0.04

lasso, the regularization was chosen to obtain the same number of parameters in both models, while for DGR this would be an unfair comparison, since penalization is done at node level. Instead, for DGR, regularization was chosen to have the same number of nodes as PrAda-net, allowing it to use many more parameters than PrAda-net and standard lasso.

The results are summarized in Table 1. PrAda-net outperforms both standard lasso and DGR, both in terms of test error and variable importance. This effect is especially notable for the higher order functions. While all methods did fairly well at identifying the noise variable as such, PrAda-net was the only method to do so in all 50 runs.

We also investigated how well the algorithms could identify the four different functions in the true model. In this comparison, DGR was left out since all input nodes are connected to all nodes in the hidden layer, resulting in only $f(x_1, x_2, x_3, x_4, x_5)$ being identified. Standard lasso was penalized with two different strengths, one to get the same number of parameters as for PrAda-net and one to get the same number of functions.

Table 2 shows that PrAda-net is able to identify the true functions in the model as well as their complexity, measured in number of nodes. That is, PrAda-net assigns more nodes to the higher order polynomials. The true functions are identified much more often by PrAda-net than by standard lasso. Standard lasso with the same complexity erroneously includes higher-order interaction terms (e.g. $f(x_1, x_2, x_3)$) rather than increasing the presence or function complexity of the main effect terms (e.g. $f(x_2)$). Standard lasso with the same number of functions successfully identifies $f_1(x_1)$ and $f_2(x_2)$ but fails to identify $f_3(x_3)$ and $f_4(x_4)$.

3.2 Black smoke data

To illustrate the usefulness of PrAda-net on real, large-scale data with a complex structure, we applied it to the U.K. black smoke network daily data set, analyzed by (Wood et al. 2017) and available on the main author's homepage¹. The data set, collected over four decades, is massive, comprising 10 million observations of air pollution data (daily concentration of black smoke particulates in $\mu\text{g m}^{-3}$) measured at

¹ https://www.maths.ed.ac.uk/~swood34/data/black_smoke.RData.

Table 2 Legendre polynomials: Presence (proportion of simulations where function is identified) and average complexity of functions across 50 simulation runs. The true model components (main effects) are marked with an asterisk. Complexity of identified functions is measured in number of nodes

Function	PrAda-net		Standard Lasso (same number (of functions))		Standard Lasso (same complexity)	
	Presence	Complexity	Presence	Complexity	Presence	Complexity
$f(x_1)^*$	1.0	1.0	1.0	1.72	0.12	1.0
$f(x_1, x_2)$	-	-	-	-	0.12	1.17
$f(x_1, x_2, x_3)$	-	-	-	-	0.14	1.0
$f(x_1, x_2, x_3, x_4)$	-	-	-	-	0.82	1.46
$f(x_1, x_2, x_3, x_4, x_5)$	-	-	-	-	0.44	1.0
$f(x_1, x_2, x_3, x_5)$	-	-	-	-	0.02	1.0
$f(x_1, x_3)$	-	-	-	-	0.96	1.06
$f(x_1, x_3, x_4)$	-	-	-	-	0.64	1.25
$f(x_1, x_3, x_4, x_5)$	-	-	-	-	0.06	1.0
$f(x_1, x_4)$	-	-	-	-	0.42	1.0
$f(x_2)^*$	1.0	1.96	0.92	2.0	0.14	1.14
$f(x_2, x_3)$	-	-	0.08	2.0	0.02	1.0
$f(x_2, x_3, x_4)$	-	-	-	-	0.04	1.0
$f(x_2, x_3, x_5)$	-	-	-	-	0.02	1.0
$f(x_2, x_4)$	0.04	1.0	-	-	0.08	1.0
$f(x_2, x_5)$	-	-	-	-	0.02	1.0
$f(x_3)^*$	1.0	2.62	0.1	1.6	0.2	1.6
$f(x_3, x_4)$	0.04	1.0	0.02	1.0	0.1	1.2
$f(x_4)^*$	1.0	3.98	0.06	1.0	0.96	1.83

more than 2000 monitoring stations. In addition to the pollution data, bs , the following eleven covariates are available: year, \underline{y} ; day of year, doy ; day of week, dow ; location as kilometers east, e and north, n ; height, h ; cubic root transformed rainfall, r ; daily minimum and maximum temperature T^0 and T^1 ; and mean temperatures the previous two days, T^1 and T^2 .

Wood et al. (2017) first perform a separate spatial modeling and then a separate temporal modeling of the data to propose some candidate model components. In a final modeling step, the components are combined through a generalized additive modeling approach, including interactions between the proposed model components. Specifically, the log transformed black smoke level is modelled as a sum of fourteen functions, each containing up to three of the eleven covariates. In addition to these functions, the model includes an offset for each station site type k (rural, industrial, commercial, city/town center or mixed), α_k ; a station specific random effect, b ; and a time correlation model for the error term following an AR process, e . The final model is given by

Table 3 Identified functions in 20 bootstrap runs using PrAda-net and standard lasso. Only the functions present in more than half of the runs are shown

	PrAda-net		Standard lasso	
Identified functions in total	52		143	
Functions with more than 3 covariates	6		103	
Most frequent functions with frequencies	$f(\text{doy})$	1.0	$f(e, n)$	0.8
	$f(y)$	1.0	$f(\text{dow})$	0.8
	$f(\text{T}^1)$	1.0	$f(n)$	0.7
	$f(n)$	1.0	$f(n, h)$	0.6
	$f(e)$	1.0	$f(h)$	0.55
	$f(e, n, h)$	1.0	$f(e)$	0.55
	$f(e, n)$	1.0		
	$f(\text{dow})$	0.95		
	$f(h)$	0.9		
	$f(n, h)$	0.85		
	$f(e, h)$	0.8		
	$f(\text{T}^0)$	0.75		
	$f(\text{doy}, \text{T}^1)$	0.6		
	$f(\text{T}^2)$	0.55		

$$\begin{aligned}
 \log(\text{bs}) = & f_1(y) + f_2(\text{doy}) + f_3(\text{dow}) + f_4(y, \text{doy}) + f_5(y, \text{dow}) \\
 & + f_6(\text{doy}, \text{dow}) + f_7(n, e) + f_8(n, e, y) + f_9(n, e, \text{doy}) \\
 & + f_{10}(n, e, \text{dow}) + f_{11}(h) + f_{12}(\text{T}^0, \text{T}^1) + f_{13}(\text{T}^1, \text{T}^2) \\
 & + f_{14}(x) + \alpha_k + b + e.
 \end{aligned} \tag{5}$$

For functional forms, cubic splines were used for the temporal components, thin-plate splines for the spatial components and tensor product smoothers for the interactions. In summary, this analysis builds on a substantial and non-trivial preliminary screening and modeling of the data.

In contrast, our approach is to let PrAda-net with 100 hidden nodes automatically decide the format of the additive model functions. To make the results reasonably comparable we use a penalization parameter for PrAda-net to obtain, approximately, the same number of functions as in Eq. 5. Note, however, that this generally produces slightly fewer functions for PrAda-net than the model by Wood et al. (2017) since we do not post-process to combine functions that are decomposed into main and interaction terms by PrAda-net. E.g. the single function $f(e, n)$ in Wood et al. (2017) is decomposed into two main effects, $f(e)$ and $f(n)$, and one interaction effect, $f(e, n)$, by PrAda-net (see Eq. 6 and Table 3) and is thus realized by three functions. It is non-trivial to conduct such post-processing for all possible main and higher order interactions so we elect to choose a conservative model for PrAda-net.

We ran PrAda-net for 20 different random splits of the data which resulted in, in total, 52 different functions, 6 of which contained more than three covariates, see Table 3. The medoid model, where the distance between models was measured by the Jaccard distance between the set of included functions, is given by

$$\begin{aligned}
 \log(\text{bS})_{\text{PrAda}} = & \underbrace{f_1(\text{y})}_{10} + \underbrace{f_2(\text{doy})}_{4} + \underbrace{f_3(\text{dow})}_{2} + \underbrace{f_4(\text{T}^0)}_1 + \underbrace{f_5(\text{T}^1)}_2 + \underbrace{f_6(\text{T}\bar{2})}_1 \\
 & + \underbrace{f_7(\text{r})}_1 + \underbrace{f_8(\text{e})}_3 + \underbrace{f_9(\text{n})}_8 + \underbrace{f_{10}(\text{h})}_2 + \underbrace{f_{11}(\text{doy}, \text{T}^1)}_2 + \underbrace{f_{12}(\text{e}, \text{n})}_{12} \\
 & + \underbrace{f_{13}(\text{e}, \text{h})}_1 + \underbrace{f_{14}(\text{n}, \text{h})}_1 + \underbrace{f_{15}(\text{T}^0, \text{T}^1, \text{T}\bar{1})}_1 + \underbrace{f_{16}(\text{e}, \text{n}, \text{h})}_3 \\
 & + \alpha_k + b + e,
 \end{aligned} \tag{6}$$

where the number under a function denotes its complexity, measured in the number of hidden nodes used to realize it. Six of the sixteen functions in the PrAda-net medoid model overlap with the manually selected functions in Wood et al. (2017). However, the pre-identified interaction terms between temporal covariates and between temporal and spatial covariates are not selected by PrAda-net, which, on the other hand, selects a more complex spatial interaction, $f_{16}(\text{e}, \text{n}, \text{h})$.

For standard lasso, the 20 data splits resulted in 143 different functions, 103 of which contained more than 3 covariates, see Table 3, with medoid function

$$\begin{aligned}
 \log(\text{bS})_{\text{Lasso}} = & \underbrace{f_1(\text{dow})}_2 + \underbrace{f_2(\text{e})}_1 + \underbrace{f_3(\text{n})}_1 + \underbrace{f_4(\text{h})}_1 + \underbrace{f_5(\text{doy}, \text{T}\bar{2})}_1 + \underbrace{f_6(\text{e}, \text{n})}_1 \\
 & + \underbrace{f_7(\text{r}, \text{h})}_1 + \underbrace{f_8(\text{n}, \text{h})}_1 + \underbrace{f_9(\text{y}, \text{r}, \text{e})}_1 + \underbrace{f_{10}(\text{y}, \text{e}, \text{n})}_1 \\
 & + \underbrace{f_{11}(\text{y}, \text{doy}, \text{T}^1, \text{r})}_1 + \underbrace{f_{12}(\text{y}, \text{T}^1, \text{T}\bar{1}, \text{r})}_1 + \underbrace{f_{13}(\text{y}, \text{e}, \text{n}, \text{h})}_1 \\
 & + \underbrace{f_{14}(\text{T}^0, \text{T}\bar{1}, \text{T}\bar{2}, \text{h})}_1 + \underbrace{f_{15}(\text{T}^1, \text{r}, \text{n}, \text{h})}_1 + \underbrace{f_{16}(\text{T}^0, \text{T}^1, \text{e}, \text{n}, \text{h})}_1 \\
 & + \underbrace{f_{17}(\text{y}, \text{doy}, \text{T}^0, \text{T}^1, \text{r}, \text{h})}_1 + \underbrace{f_{18}(\text{y}, \text{dow}, \text{T}\bar{1}, \text{r}, \text{n}, \text{h})}_1 + \alpha_k + b + e.
 \end{aligned} \tag{7}$$

Compared to standard lasso, PrAda-net has higher precision, i.e. it detects fewer unique functions more frequently, indicating a more stable selection performance. This is confirmed by the medoid model being more similar to the other models for PrAda-net than for standard lasso. The average Jaccard similarity of the PrAda-net medoid model is 0.63, compared with 0.18 for standard lasso. PrAda-net also identifies functions containing fewer covariates (order of interaction) thus resulting in more interpretable

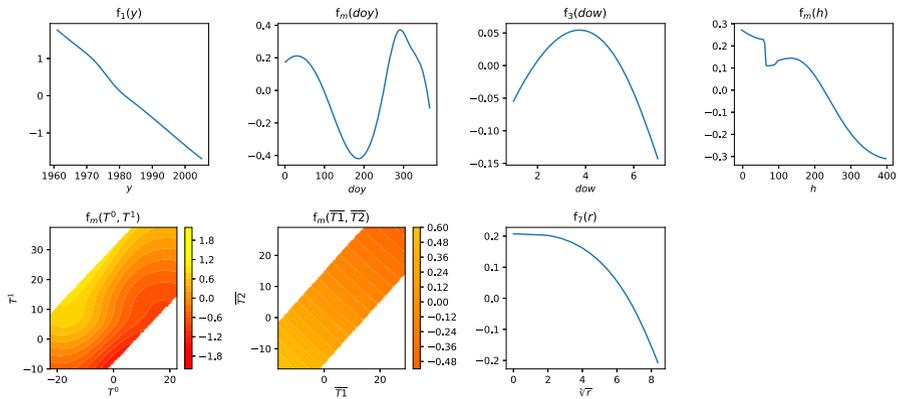


Fig. 3 Functions identified by PrAda-net to be compared to functions by Wood et al. (2017), Supplemental Fig. 2, <https://www.tandfonline.com/doi/suppl/10.1080/01621459.2016.1195744>. Since day of year, height and the temperature covariates are present in multiple functions in PrAda-net, in these cases the marginal function is displayed. Subscript numbers refer to functions in Eq. 6, while subscript m denotes marginal function

models compared to standard lasso. All the $2 \cdot 20$ models had an explained variance of $R^2 \approx 0.79$, which is the same as reported by Wood et al. (2017).

In Figure 3 we visually summarize a subset of the model components identified by PrAda-net in Eq. 6. These same model components are presented in Fig. 2 in the Supplemental material of Wood et al. (2017)². In the cases where PrAda-net splits the dependency of a covariate across multiple functions, the marginal function is plotted to allow for direct comparison with Wood et al. (2017). The generated functions in Fig. 3 bear a strong resemblance to those in Supplemental Fig. 2 of Wood et al. (2017), indicating that PrAda-net was able to automatically identify a highly similar model to the manually curated one. With the exception of height, h , PrAda-net tends to identify smoother functions than Wood et al. (2017).

3.3 Synthetic black smoke data

To test PrAda-net on a more complex model than the synthetic data in Sect. 3.1, we generated data from the PrAda-net medoid model for the black smoke data. 10000 data points were sampled, according to $\mathcal{N}(0, \Sigma_S)$, where Σ_S is the Spearman rank correlation matrix of the variables in the black smoke data set, resulting in a 10000×11 x -matrix, from which a 10000×1 y -matrix was created according to Eq. 6, without α_k , b , and e and with added noise distributed as $\mathcal{N}(0, 0.1^2)$. 50 simulations, with different noise realizations, were performed using a hidden layer with 100 units. For PrAda-net, λ^* was chosen as described in Sect. 2.2 and for standard lasso, the regularization strength was chosen to obtain the same number of functions as for PrAda-net. The results are summarized in Table 4, where all functions present in more than 20 % of

² Available at <https://www.tandfonline.com/doi/suppl/10.1080/01621459.2016.1195744>.

Table 4 Presence and average complexity of functions across 50 simulation runs for PrAda-net (left) and standard lasso (right), with true model components above the line and falsely included components below. Only functions present in more than 20 % of the runs are presented, functions present in more than 50 % of the runs are marked with an asterisk

PrAda-net			Standard lasso		
Function	Presence	Complexity	Function	Presence	Complexity
$f(y)^*$	1.0	5.06	$f(y)^*$	1.0	15.8
$f(\text{doy})^*$	1.0	20.92	$f(e, n)^*$	1.0	4.16
$f(\text{dow})^*$	1.0	2.9	$f(e, n, h)^*$	0.96	1.27
$f(e)^*$	1.0	2.36	$f(\text{doy})^*$	0.82	1.76
$f(n)^*$	1.0	3.0	$f(n)^*$	0.82	1.37
$f(\text{doy}, T^1)^*$	1.0	2.22	$f(h)^*$	0.56	1.0
$f(e, n)^*$	1.0	13.44	$f(\text{dow}, e, h)^*$	0.7	1.2
$f(n, h)^*$	1.0	1.18	$f(T^0, e, n, h)^*$	0.58	1.03
$f(e, n, h)^*$	1.0	2.22	$f(\text{doy}, T^1, \bar{T}2, e, h)^*$	0.52	1.27
$f(T^1)^*$	0.98	1.27	$f(\text{doy}, T^0, T^1, \bar{T}2, r, e, n, h)^*$	0.52	1.0
$f(T^0)^*$	0.94	1.23	$f(\text{dow}, e)$	0.44	1.05
$f(e, h)^*$	0.72	1.44	$f(\text{doy}, T^0, T^1, \bar{T}1, \bar{T}2, r, e, n, h)$	0.34	1.18
$f(h)$	0.48	1.04	$f(\text{doy}, T^0, T^1, e, h)$	0.32	1.0
$f(r)$	0.22	1.0	$f(\text{doy}, \text{dow}, e, h)$	0.3	1.07
$f(\text{doy}, \bar{T}1)$	0.4	1.05	$f(\text{doy}, T^1, \bar{T}2, e)$	0.3	1.0
$f(\text{doy}, e)$	0.3	1.07	$f(\text{doy}, \bar{T}1, e, h)$	0.3	1.0
$f(\text{dow}, e)$	0.24	1.17	$f(\text{doy}, T^0, T^1, \bar{T}1, \bar{T}2, h)$	0.3	1.0
$f(T^1, \bar{T}1)$	0.2	1.0	$f(\text{doy}, T^0, T^1, \bar{T}2, e, h)$	0.3	1.2
			$f(T^0, T^1, \bar{T}1, \bar{T}2, r, e, n, h)$	0.26	1.0
			$f(T^1, n)$	0.24	1.0
			$f(T^0, e, n)$	0.24	1.0
			$f(\text{doy}, T^0, T^1, \bar{T}1, \bar{T}2, r, e, h)$	0.24	1.0
			$f(\text{doy}, T^1, e, h)$	0.22	1.0

the simulations are shown and functions present in more than 50 % of the simulations are marked with an asterisk.

PrAda-net identifies 14 of the 16 true model components together with 4 false positives, the corresponding numbers for standard lasso are 6 and 17, respectively. There is a tendency that more complex functions, i.e. functions consisting of more than one hidden node, are more easily identified. Except for $f_5(T^1)$ and $f_{10}(h)$, PrAda-net identifies all the functions realized by more than one hidden node in 100 % of the simulations; for both algorithms all the falsely included functions also have a complexity close to one node. Apart from $f(\text{doy})$, PrAda-net seems to be closer to the true complexity than standard lasso.

In Figs. 4 and 5, the means and standard deviations for the most frequently identified functions are plotted together with the true functions. Since the intercept can be

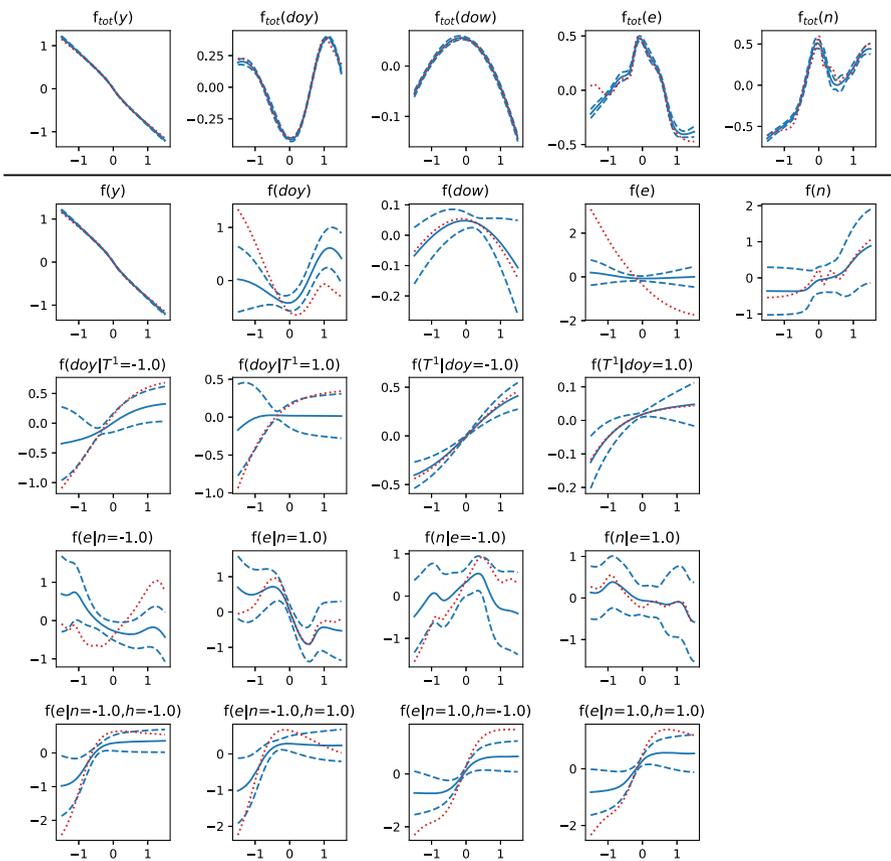


Fig. 4 PrAda-net: Comparison of true (red and dotted) and the mean \pm two standard deviations of the identified functions (blue). For the multi-dimensional functions, in each plot all but one of the input variables are kept fixed. The first row shows the total dependency of the covariate, including all functions, while the rest of the rows show the most frequently identified functions

realized in any function, all functions were translated to have zero mean. For functions depending on more than one covariate, all covariates except one were fixed at ± 1 to be able to plot in one dimension. For PrAda-net we see that for total the contribution of a covariate, the true and identified functions coincide very well (first row). This is also true for most identified functions with a few exceptions. However, the differences between true and identified functions seem to be balanced by those of other functions including the same covariate. See e.g. $f(\text{doy})$ and $f(\text{doy}|T^1)$, where, for small values of doy , the true $f(\text{doy})$ is much larger than the estimated function, while the true $f(\text{doy}|T^1)$ is much smaller than the estimated function.

Standard lasso exhibits poor overlap with the true model components, even in terms of total contribution. The functions identified by standard lasso align poorly with the true model components.

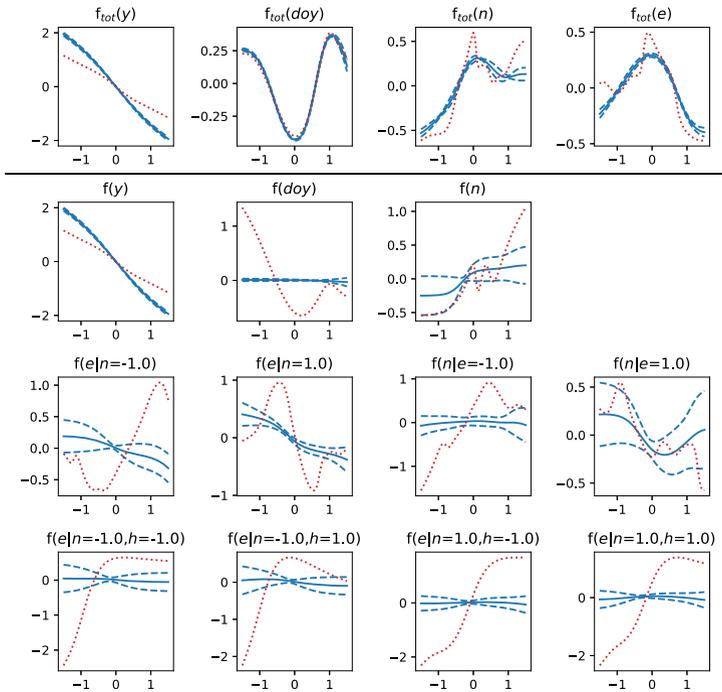


Fig. 5 Standard lasso: Comparison of true (red and dotted) and the mean \pm two standard deviations of the identified functions (blue). For the multi-dimensional functions, in each plot all but one of the input variables are kept fixed. The first row shows the total dependency of the covariate, including all functions, while the rest of the rows show the most frequently identified functions

4 Conclusions and discussion

For complex and large data sets, non-parametric additive models often require substantial exploratory work to identify candidate model components, such as sets of potential interactions and format or family of additive function, before any fitting strategy can be applied. Here, we proposed a fully data-adaptive alternative based on a simple neural network trained with proximal gradient descent and adaptive lasso, PrAda-net. PrAda-net was found to improve on lasso penalized neural networks both in terms of test error performance and in terms of generating interpretable models. For additive models, PrAda-net was able to identify the function components of the models as well as to express the function complexity by using multiple hidden nodes. We illustrated how PrAda-net could be used to model a massive and complex air pollution data set with weather, temporal and spatial covariates. PrAda-net was able to identify a highly similar model to one recently presented in the state-of-the-art literature, but, in contrast, required no pre-selection of modeling components or pre-processing of the data. This data driven strategy thus reduces the dependency on subjective choices and preliminary or partial modeling of complex data, while retaining the interpretability of the classical statistical methods.

While there is no explicit limitation to the depth of the network that PrAda-net can be applied to, we only used one hidden layer in this paper. We thereby loose the possibility of making the network even more compact. However, preliminary results from applying PrAda-net to deep networks indicated that interpretability was much reduced in favor of compact representation. It is not trivial to untangle the compact representation to obtain an interpretable additive model representation. In addition, according to the universal approximation theorem (Cybenko 1989) the utilization of a one-hidden-layer network does not limit the complexity of the functions that can be modelled. It would, nonetheless, be interesting to apply extensions of PrAda-net to deeper networks, with regularization penalties constructed to e.g. guarantee an ordering of main to higher order effects in the layers. Indeed, structural constraints may be what is needed to generate interpretable networks from more complex architectures. However, this is left for future work.

Code is available at https://github.com/allerbo/prada_net.

Acknowledgements This research was supported by grants from the Swedish Research Council (VR), the Swedish Foundation for Strategic Research (SSF) and the Chalmers AI Research Center (CHAIR).

Funding Open access funding provided by University of Gothenburg.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ainsworth SK, Foti NJ, Lee AK, Fox EB (2018) oi-vae: output interpretable vaes for nonlinear group factor analysis. In: international conference on machine learning, pp. 119–128
- Cybenko G (1989) Approximation by superpositions of a sigmoidal function. *Math Control Signals Syst* 2(4):303–314
- Friedman J, Hastie T, Tibshirani R (2010) Regularization paths for generalized linear models via coordinate descent. *J Stat Softw* 33(1):1
- Friedman JH, Stuetzle W (1981) Projection pursuit regression. *J Am Stat Assoc* 76(376):817–823
- Hastie TJ, Tibshirani RJ (1990) Generalized additive models, vol 43. CRC Press, London
- Kingma DP, Ba J (2014) Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*
- Leclerc G, Vartak M, Fernandez RC, Kraska T, Madden S (2018) Smallify: Learning network size while training. *arXiv preprint arXiv:1806.03723*
- Lym S, Choukse E, Zangeneh S, Wen W, Erez M, Shanghavi S (2019) Prunetrain: Gradual structured pruning from scratch for faster neural network training. *arXiv preprint arXiv:1901.09290*
- Qiao J, Wang L, Yang C (2018) Adaptive lasso echo state network based on modified bayesian information criterion for nonlinear system modeling. *Neural Comput Appl* 31(10):6163–6177
- Rockafellar RT (1976) Monotone operators and the proximal point algorithm. *SIAM J Control Optim* 14(5):877–898
- Scardapane S, Comminiello D, Hussain A, Uncini A (2017) Group sparse regularization for deep neural networks. *Neurocomputing* 241:81–89
- Simonyan K, Vedaldi A, Zisserman A (2013) Deep inside convolutional networks: visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*

- Sun K, Huang S-H, Wong DS-H, Jang S-S (2017) Design and application of a variable selection method for multilayer perceptron neural network with lasso. *IEEE Trans Neural Netw Learn Syst* 28(6):1386–1396
- Tank A, Cover I, Foti NJ, Shojaie A, Fox EB (2017) An interpretable and sparse neural network model for nonlinear granger causality discovery. *arXiv preprint* [arXiv:1711.08160](https://arxiv.org/abs/1711.08160)
- Tibshirani R (1996) Regression shrinkage and selection via the lasso. *J R Stat Soc Ser B (Methodol)* 58(1):267–288
- Vaughan J, Sudjianto A, Brahim E, Chen J, Nair VN (2018) Explainable neural networks based on additive index models. *arXiv preprint* [arXiv:1806.01933](https://arxiv.org/abs/1806.01933)
- Wang J, Xu C, Yang X, Zurada JM (2018) A novel pruning algorithm for smoothing feedforward neural networks based on group lasso method. *IEEE Trans Neural Netw Learn Syst* 29(5):2012–2024
- Wang Y, Su H, Zhang B, Hu X (2018b) Interpret neural networks by identifying critical data routing paths. In: proceedings of the IEEE conference on computer vision and pattern recognition, pp. 8906–8914
- Wood SN, Li Z, Shaddick G, Augustin NH (2017) Generalized additive models for gigadata: modeling the uk black smoke network daily data. *J Am Stat Assoc* 112(519):1199–1210
- Xu M, Han M (2016) Adaptive elastic echo state network for multivariate time series prediction. *IEEE Trans Cybern* 46(10):2173–2183
- Yuan M, Lin Y (2006) Model selection and estimation in regression with grouped variables. *J R Stat Soc Ser B (Stat Methodol)* 68(1):49–67
- Zhang H, Wang J, Sun Z, Zurada JM, Pal NR (2019) Feature selection for neural networks using group lasso regularization. *IEEE Trans Knowl Data Eng* 32(4):659–673
- Zou H (2006) The adaptive lasso and its oracle properties. *J Am Stat Assoc* 101(476):1418–1429

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.