








Ranking Feature-Block Importance in Artificial Multiblock Neural Networks

Anna Jenul^(✉) , Stefan Schrunner , Bao Ngoc Huynh , Runar Helin ,
Cecilia Marie Futsaether , Kristian Hovde Liland , and Oliver Tomic 

Norwegian University of Life Sciences, Universitetstunet 3, 1432 Ås, Norway
{anna.jenul,stefan.schranner,ngoc.huynh.bao,runar.helin,
cecilia.futsaether,kristian.liland,oliver.tomic}@nmbu.no

Abstract. In artificial neural networks, understanding the contributions of input features on the prediction fosters model explainability and delivers relevant information about the dataset. While typical setups for feature importance ranking assess input features individually, in this study, we go one step further and rank the importance of groups of features, denoted as feature-blocks. A feature-block can contain features of a specific type or features derived from a particular source, which are presented to the neural network in separate input branches (multiblock ANNs). This work presents three methods pursuing distinct strategies to rank feature-blocks in multiblock ANNs by their importance: (1) a composite strategy building on individual feature importance rankings, (2) a knock-in, and (3) a knock-out strategy. While the composite strategy builds on state-of-the-art feature importance rankings, knock-in and knock-out strategies evaluate the block as a whole via a mutual information criterion. Our experiments consist of a simulation study validating all three approaches, followed by a case study on two distinct real-world datasets to compare the strategies. We conclude that each strategy has its merits for specific application scenarios.

Keywords: Feature-block importance · Importance ranking · Multiblock neural network · Explainability · Mutual information

1 Introduction

In machine learning, datasets with an intrinsic block-wise input structure are common; blocks may represent distinct data sources or features of different types and are frequently present in datasets from industry [7], biology [3], or healthcare [5]. For example, in healthcare, heterogeneous data blocks like patient histology, genetics, clinical data, and image data are combined in outcome prediction models. However, good prediction models do not necessarily depend equally on each block. Instead, some blocks may be redundant or non-informative. Identifying the key data sources in multi-source treatment outcome models promises to deliver new insights into the behavior of black-box models like ANNs. In particular, potential benefits include improving model explainability, reducing costly

data acquisitions that do not contribute to the model prediction, and allowing domain experts to explore latent relations in the data. Thus, there is a need to measure the importances of feature-blocks, denoted as feature-block importance ranking (BIR).

In order to exploit the internal structure of the block-wise data in neural networks, a multiblock ANN (M-ANN) architecture is used. As depicted in Fig. 1, the M-ANN consists of a separate input branch for each block, a concatenation layer to merge information from all branches, and a blender network to map the information to the model output. The architecture allows for any type of network layer, depth, activation, or other network parameters, including the special case where the concatenation layer equals the input layer (block branches of depth 0).

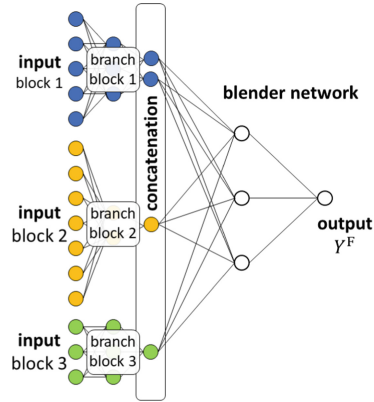


Fig. 1. M-ANN architecture.

Ranking individual features by their importances (feature importance ranking, FIR) has been studied for different types of ANNs [8, 14, 15]. An extensive evaluation [9] showed that versions of the variational gradient method (VarGrad) [1, 2] outperformed competitors such as guided backprop and integrated gradients. For BIR, however, a combination of features in one block may accumulate a larger amount of information than each feature separately due to informative non-linear relations between features. Hence, using FIR might oversimplify the problem of measuring block importance since interactions between features of the same block are disregarded. Nevertheless, the strategy of reducing BIR to a simple summary metric (sum, mean, max) over FIR scores is considered in our evaluation.

A related problem to FIR is feature selection, where the input dimensionality is reduced to the most influential features as part of the preprocessing. Feature selection is widely studied in ANNs. Furthermore, specialized feature selectors can account for block structures like UBayFS [11] or groupLasso [16]. Conceptually, these feature selectors aim to improve model performance and classify an entire block as important/unimportant in a binary way before model training. In contrast, our BIR problem is considered a post-processing procedure, focusing on analyzing the model after training without influence on the model performance.¹

This study presents and discusses three distinct approaches to quantify the importance of feature-blocks in M-ANNs. While exploiting the flexibility of ANNs and their capacities to learn complex underlying patterns in the data, the

¹ BIR may be used for block feature selection if deployed as filter method—however, this aspect is beyond the scope of the present work.

discussed methods aim to deliver insights into the trained network’s dependence structure on the distinct input blocks and thereby foster model explainability. We propose three paradigms for BIR: a block is considered as important if

1. it consists of features with high FIR scores (composition strategy), or
2. it explains a large part of the network output (knock-in strategy), or
3. its removal significantly changes the network output (knock-out strategy).

We evaluate and discuss the proposed paradigms in a simulation study and present two case studies on real-world datasets, where the behaviors of the proposed ranking strategies become apparent.

In the following, bold letters denote vectors and matrices; non-bold letters denote scalars, functions or sets. Square brackets denote index sets $[n] = \{1, \dots, n\}$.

2 Block Importance Ranking Methods

We assume data input \mathbf{x} from some feature space $D \subset \mathbb{R}^N$, $N \in \mathbb{N}$, following a probability distribution $\mathbf{X} \sim P_X$, and a univariate target variable $y \in T \subset \mathbb{R}$ following a probability distribution $Y \sim P_Y$. Given training data $(\mathbf{x}, y) \in D_{\text{train}} \times T_{\text{train}} \subset D \times T$, model parameters $\mathbf{w} \in W \subset \mathbb{R}^M$, $M \in \mathbb{N}$, are trained with respect to some loss term $e : D \times T \rightarrow \mathbb{R}^+$,

$$\mathbf{w}^* = \min_{\mathbf{w} \in W} e(f_{\mathbf{w}}(\mathbf{x}), y),$$

where the ANN is a function $f_{\mathbf{w}} : D \rightarrow T$ given weights \mathbf{w} .

In an M-ANN architecture, see Fig. 1, the block structure of the model input is represented by a direct sum of subspaces $D = \bigoplus_{b=1}^B D_b$, each corresponding to one block $b \in [B]$ with dimension $N_b = \dim(D_b)$, $N = \sum_{b=1}^B N_b$. Each block enters a distinct branch of the network that processes the block input. Afterwards, the outputs of all branches are merged in a concatenation layer, which consists of n_b nodes associated with each block b , respectively. A so-called blender network $f_{\mathbf{w}}^{\text{blender}}$ connects the concatenation layer to the network output. Network training is performed using backpropagation, where all block branches and the blender network are trained simultaneously in an end-to-end manner.

2.1 Composite Strategy

Our first paradigm composes block importance measures from FIR in a direct way. As a prototype of state-of-the-art FIR methods, we use VarGrad [2]. VarGrad builds on the idea that variations of an important feature provoke measurable variations in the output. Under the assumption that features are on a common scale, we estimate the gradient of the function $f_{\mathbf{w}}$ with respect to each feature by adding small random perturbations in the input layer. A large

variance in the gradient indicates that the network output depends strongly on a feature, i.e., the feature is important. We denote the importance of feature $n \in [N]$ as quantified by VarGrad, by $\alpha_n \in \mathbb{R}^+$.

To translate the feature-wise importance measure to feature-blocks in M-ANNs, we deploy a summary metric φ over all single-feature importances in a block $b \in [B]$. Thus, block importances $\gamma_\varphi^{(b)}$ are defined as

$$\gamma_\varphi^{(b)} = \varphi(\alpha_1^{(b)}, \dots, \alpha_{N_b}^{(b)}), \quad (1)$$

where $\alpha_n^{(b)}$ denotes the n^{th} feature associated with the b^{th} block. Intuitive choices for φ are either the sum, mean, or maximum operator, denoted as φ_{sum} , φ_{mean} , or φ_{max} , respectively. Rankings based on mean and sum are equal, if all blocks contain the same number of features. Operators φ_{sum} and φ_{mean} accumulate the individual feature importances: a block with multiple features of high average importances is preferred over blocks with few top features and numerous unimportant features. In contrast, φ_{max} compares the top-performing features out of each block, while neglecting all other's contributions. Statistical properties of block importance quantifiers implementing the composite strategy are transmitted from (i) the feature importance ranking method and (ii) the summary metric. Since this approach cannot capture between-feature relations, potentially impacting the importance of a block, we suggest two other paradigms.

2.2 Knock-In Strategy

The knock-in strategy is inspired by work on the information bottleneck [4], demonstrating that node activations can be exploited for model interpretation in ANNs. In the concatenation layer of the M-ANN (Fig. 1), where information from the blocks enters the blender network, activations are of particular importance since they represent an encoding of the block information. When passing model input \mathbf{x} through the network, we denote the activation of the n th node associated with block $b \in [B]$ in the concatenation layer by $c_{b,n}(\mathbf{x})$, $n \in [n_b]$. The average activation of the n th node in block $b \in [B]$ across all training data $\mathbf{x} \in D_{\text{train}}$ is denoted by $\bar{c}_{b,n}$.

For BIR, we compute a pseudo-output by passing data of only one block b through the network. For this purpose, we introduce a pseudo-input $\mathbf{v}^{(b)}(\mathbf{x})$ as

$$\mathbf{v}_{b',n}^{(b)}(\mathbf{x}) = \begin{cases} c_{b',n}(\mathbf{x}) & \text{if } b' = b \\ \bar{c}_{b',n} & \text{otherwise,} \end{cases} \quad (2)$$

where $b' \in [B]$, and $n \in [n_b]$. By propagating pseudo-input $\mathbf{v}^{(b)}(\mathbf{x})$ through the blender network, we obtain the pseudo-output $f_w^{\text{blender}}(\mathbf{v}^{(b)}(\mathbf{x}))$. The main assumption behind the knock-in strategy is that high agreement between output $f_w(\mathbf{x})$ and pseudo-output $f_w^{\text{blender}}(\mathbf{v}^{(b)}(\mathbf{x}))$ indicates a high importance of block b , since information from b is sufficient to recover most of the model output. In contrast, a large discrepancy between the two quantities indicates low explanatory power of the block b , and thus, a lower block importance. The concept to generate knock-in pseudo-outputs is illustrated in Fig. 2.

We implement the knock-in concept via the mutual information (MI) [6], an information-theoretic measure to quantify the level of joint information between two discrete random variables Z and Z' , defined as

$$\text{MI}(Z, Z') = \sum_z \sum_{z'} p_{Z, Z'}(z, z') \log_2 \left(\frac{p_{Z, Z'}(z, z')}{p_Z(z)p_{Z'}(z')} \right).$$

If Z and Z' are independent, $\text{MI}(Z, Z')$ is 0. Otherwise, $\text{MI}(Z, Z')$ is positive, where a high value indicates a large overlap in information. To quantify the joint and marginal distributions of continuous variables Z and Z' , two-dimensional and one-dimensional histograms can be used as non-parametric estimators for $p_{Z, Z'}$, p_Z , and $p_{Z'}$, respectively. We denote the number of equidistant histogram bins along each axis by $\ell \in \mathbb{N}$. It follows from the properties of entropy [6] that an upper bound to $\text{MI}(Z, Z')$ is given by $\log_2(\ell)$.

As shown in Fig. 2, the random variable of (full) model output, $Y^F = f_w(\mathbf{X})$, and the random variable of the pseudo-output with respect to block b , $Y^{(b)} = f_w^{\text{blender}}(\mathbf{v}^{(b)}(\mathbf{X}))$, where \mathbf{X} follows the input distribution $P_{\mathbf{X}}$, are used to measure knock-in (KI) block importance as

$$\gamma_{\text{KI}}^{(b)} = \frac{\text{MI}(Y^F, Y^{(b)})}{\log_2(\ell)}. \quad (3)$$

2.3 Knock-Out Strategy

The knock-out paradigm is an ablation procedure where one block at a time is removed from the model in order to measure the impact of the remaining blocks. We pursue a similar approach as in the knock-in paradigm and specify knock-out pseudo-inputs $\mathbf{v}^{(-b)}(\mathbf{x})$ as

$$v_{b', n}^{(-b)}(\mathbf{x}) = \begin{cases} \bar{c}_{b', n} & \text{if } b' = b \\ c_{b', n}(\mathbf{x}) & \text{otherwise,} \end{cases} \quad (4)$$

for an arbitrary block $b \in [B]$. Thus, the definition in Eq. 4 represents an opposite behavior of Eq. 2 in the knock-in case. In analogy to the knock-in notation, we denote the random variable of pseudo-outputs with respect to $\mathbf{v}^{(-b)}$ as $Y^{(-b)} = f_w^{\text{blender}}(\mathbf{v}^{(-b)}(\mathbf{X}))$. The knock-out concept is illustrated in Fig. 3. In contrast to knock-in, we assume that leaving out block b having a relevant impact on the final output delivers a more dissimilar pseudo-output to the full output since relevant information is lost. Removing an unimportant block preserves the relevant information and delivers a pseudo-output similar to the full output.

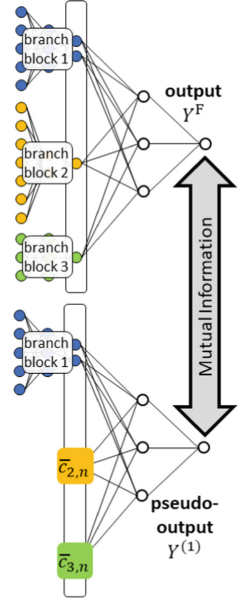


Fig. 2. Knock-in strategy: Pseudo-outputs for feature-block $b = 1$ are generated by activating block b , while imputing averaged activations for all other blocks.

Finally, we define the importance of block $b \in [B]$ with respect to the knock-out strategy (KO) as

$$\gamma_{\text{KO}}^{(b)} = \frac{\log_2(\ell) - \text{MI}(Y^{\text{F}}, Y^{(-b)})}{\log_2(\ell)}. \quad (5)$$

For both, KI and KO, importance scores $\gamma_{\text{KI}}^{(b)}$ and $\gamma_{\text{KO}}^{(b)}$ are bounded between 0 (unimportant block) and 1 (important block).

3 Experiments

As a proof of concept, we conduct two experiments to assess BIR in M-ANNs. The first experiment involves six simulated, non-linear regression problems, where our simulation setup delivers information on the ground truth block importances. This experiment verifies that our suggested measures can identify the ground truth block rankings, defined by their corresponding paradigms. Real-world datasets are evaluated in two case studies in experiment 2, where no exact ground truth block ranking is available. Instead, we compare BIR strategies to each other.

3.1 Simulation Experiment

We simulate a synthetic datasets along with six distinct target functions, denoted as setups S1a–S1c and S2a–S2c. The dataset consists of $N = 256$ features, divided randomly into $B = 8$ blocks (B1–B8) à $N_b = 32$ features. The sample size is set to $|D_{\text{train}}| = 10\,000$ and $|D_{\text{test}}| = 10\,000$. All features are simulated from a multivariate normal distribution with mean vector $\boldsymbol{\mu} = \mathbf{0}$ and a randomized covariance matrix $\boldsymbol{\Sigma}$; hence a non-trivial correlation structure is imposed.²

Setups S1a–S1c and S2a–S2c differ in the parameters used to compute the non-linear target variable y , which is simulated via a noisy linear combination of the squared features with coefficient matrix $\boldsymbol{\beta}^{(b)} \in \mathbb{R}^{N_b \times N_b}$, given as

$$y = \underbrace{\sum_{b=1}^8 \mathbf{x}^T \boldsymbol{\beta}^{(b)} \mathbf{x}}_{g(\mathbf{x})} + \varepsilon_{\text{noise}}, \quad \text{where} \quad \varepsilon_{\text{noise}} \underset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_{\text{noise}}^2), \quad \text{and} \quad \boldsymbol{\beta}^{(b)} = \begin{pmatrix} \beta_{\text{imp}} & 0 & \dots & 0 & 0 & 0 \\ \beta_{\text{int}} & \beta_{\text{imp}} & \dots & 0 & 0 & 0 \\ \dots & \dots & \ddots & \dots & \dots & \dots \\ \beta_{\text{int}} & \beta_{\text{int}} & \dots & \beta_{\text{imp}} & 0 & 0 \\ \hline 0 & 0 & \dots & 0 & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & 0 & 0 & 0 \end{pmatrix}. \quad (6)$$

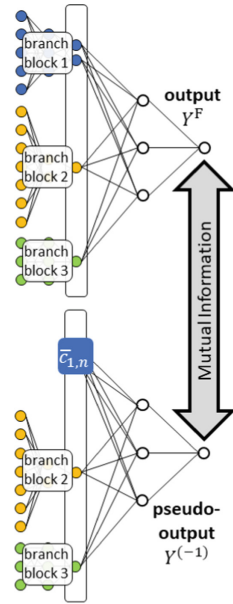


Fig. 3. Knock-out strategy: Pseudo-outputs are generated by activating all but one blocks.

² Code and details on simulation and network architecture are available at <https://github.com/annajenul/Block.Importance.Quantification>.

Table 1. Specifications for matrix $\beta^{(b)}$: block importance is steered via count N_{imp} , coefficient β_{imp} , and interaction β_{int} of the important features.

Setup	Block												
	B1			B2			...	B7			B8		
	N_{imp}	β_{imp}	β_{int}	N_{imp}	β_{imp}	β_{int}	...	N_{imp}	β_{imp}	β_{int}	N_{imp}	β_{imp}	β_{int}
S1a	2	7	0	2	6	0	...	2	1	0	0	0	0
S1b	7	2	0	6	2	0	...	1	2	0	0	0	0
S1c	1	7	0	2	6	0	...	7	1	0	0	0	0
S2a	2	7	1	2	6	1	...	2	1	1	0	0	1
S2b	7	2	1	6	2	1	...	1	2	1	0	0	1
S2c	1	7	1	2	6	1	...	7	1	1	0	0	1

The matrix $\beta^{(b)} \in \mathbb{R}^{N_b \times N_b}$ contains an $N_{\text{imp}} \times N_{\text{imp}}$ quadratic sub-matrix consisting of coefficients β_{imp} of important features, i.e. features with relevant contribution to the target, and interactions β_{int} . The noise parameter σ_{noise} is set to 10% of the standard deviation of the linear combination $g(\mathbf{x})$ across the generated samples \mathbf{x} . As shown in Table 1, block importances are varied between the setups and as follows

- S1a: varying coefficients of important features, but constant counts;
- S1b: varying counts of important features, but constant coefficients;
- S1c: varying counts and coefficients of important features;
- S2a–S2c: same as S1a–S1c, but with interaction terms between features.

Due to the randomized correlation matrix of the feature generation, unimportant features may be correlated with important features, as well as with the target y .

For each setup, we trained the described M-ANN model in 30 independent runs with distinct weight initializations after data standardization. Since BIR methods are deployed post-hoc and assume a model with appropriate performance, runs with poor performances ($R2 < 0.8$) were excluded from the analysis after outlier removal. Hence, the number of model runs in the analysis was 20 (S1a, S1b, S2a, S2b), 18 (S1c), and 19 (S2c), respectively. The remaining models achieved an average performance of ≥ 0.9 (R2 score) and ≤ 0.2 (RMSEIQR: root mean squared error scaled by inter-quartile range) on the test set.

For evaluation, importance scores across all model runs were tested for significant differences using a pairwise Wilcoxon-test with Bonferroni correction. If the p-value in a comparison between two blocks was above a significance level of 0.01, both were counted as tie. Figure 4 illustrates the distributions of BIR scores after min-max-normalization by setup and method, along with rankings (colors) based on significant group differences. All methods discovered the intrinsic ranking in dataset S1a. In dataset S1b, knock-in, knock-out, and VarGrad-mean identified the ranking by underlying important feature counts N_{imp} , while VarGrad-max failed to deliver a significant distinction between blocks with higher counts of

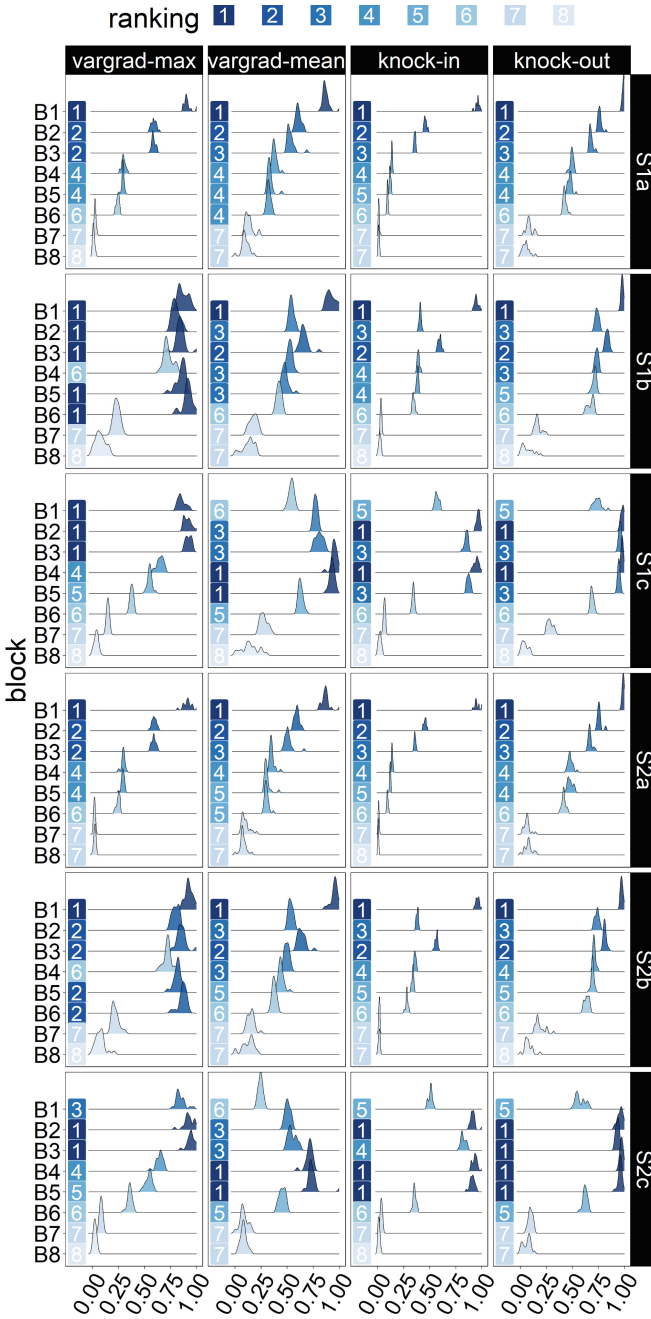


Fig. 4. Distributions of the normalized BIR scores across model runs. Rankings are indicated by colors and refer to significant group differences based on a pairwise Wilcoxon-test (significance level 0.01).

Table 2. Averaged Spearman’s rank correlation coefficients comparing each ranking to the ground truth BIR for each paradigm across model runs. Standard deviations were ≤ 0.03 for S1a, S1b, S2a and S2b, and ≤ 0.06 for S1c and S2c.

Paradigm	Dataset					
	S1a	S1b	S1c	S2a	S2b	S2c
Composite (VarGrad-max)	0.97	0.58	0.93	0.98	0.58	0.91
Composite (VarGrad-mean)	0.98	0.95	0.96	0.97	0.95	−0.40
Knock-in	0.99	0.98	0.85	0.97	0.99	0.89
Knock-out	0.99	0.98	0.81	0.99	0.99	0.89

important features. For dataset S1c, VarGrad-max mostly ranked by underlying β_{imp} and ignored N_{imp} , while knock-in, knock-out and VarGrad-mean delivered trade-offs between counts N_{imp} and coefficients β_{imp} of important features. In setups S2a, S2b, and S2c with between-feature interactions, the same rankings as in S1a, S1b, and S1c could be obtained by all methods with negligible deviations. Hence, we conclude that all metrics remain stable in more complex scenarios.

We further validated the paradigms by comparing the results to their corresponding ground truth block importances, determined by the real coefficients in the simulation setup. For the composite max and mean paradigms, the corresponding maxima and means over $\beta^{(b)}$, were used as references. Ground truth importances for knock-in (KI), and knock-out (KO) were based on the explained variances of the single block b in the underlying linear combination, given as

$$KI_b = \mathbb{E} \left(y - \left(\mathbf{x}^{(b)} \right)^T \boldsymbol{\beta}^{(b)} \mathbf{x}^{(b)} \right), \text{ and } KO_b = \mathbb{E} \left(y - \sum_{\substack{b'=1 \\ b' \neq b}}^8 \left(\mathbf{x}^{(b')} \right)^T \boldsymbol{\beta}^{(b')} \mathbf{x}^{(b')} \right),$$

where $\mathbf{x}^{(b)}$ denotes projection of input \mathbf{x} on the subspace of block b , D_b . The comparison between the rankings based on (average) predicted importance scores and ground truth rankings was made using Spearman’s correlation coefficient, see Table 2. With two exceptions, all correlation values were at a high level, indicating that our methods accurately predicted the ground truth. Spearman’s correlation coefficient is not representative in S1b, and S2b with respect to the maximum metric since the ground truth ranking is equal for blocks B1–B7. In S2c VarGrad-mean is distracted by decreasing β_{imp} and an increasing number of interaction terms, although underlying block importances are in increasing order with respect to the mean metric.

3.2 Real-World Experiment

Since verification on simulated data showed that the presented approaches match the ground truth according to their paradigms, we deployed the methods on two

real-world datasets, where underlying block importance is unknown. Prior to analysis, both datasets were standardized on the trained data. Again, we trained 30 independent model runs.

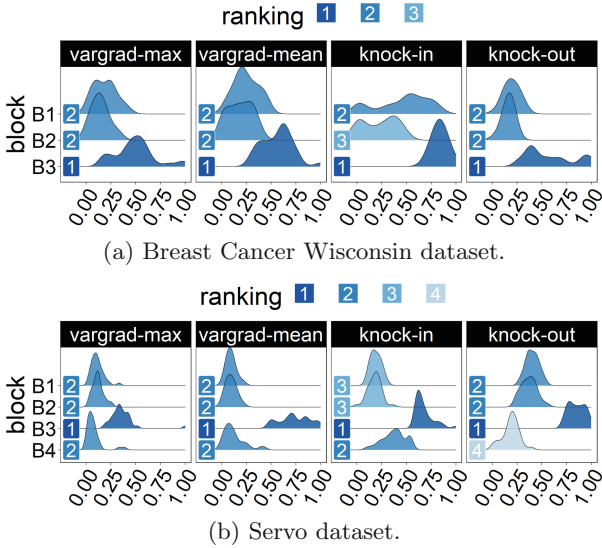


Fig. 5. Distributions of normalized BIR scores in experiment 2 across model runs.

The Breast Cancer Wisconsin dataset (BCW) [13] describes a binary classification problem (malignant or benign tumor) and consists of 569 samples (398 train, 171 test) and three blocks with ten features each, representing groups of distinct feature characteristics (mean values, standard deviations, and extreme values of measured parameters). The average performance was 0.95 (accuracy) and 0.96 (F1 score) without outliers. The average scores and rankings delivered by BIR methods are shown in Fig. 5a. All four paradigms discovered that block 3 is dominant, which agrees with previous research on the dataset [10]. However, knock-in was the only method that distinguished between the importances of B1 and B2. According to [10], block B1 contains overlapping information with B3, while B2 is rather non-informative. Thus, the experiment underlines a difference between knock-in and knock-out rankings in the presence of redundancies.

Servo [12] is a dataset containing 167 samples (120 train, 47 test), a univariate, numeric target variable, and four features, two of which are categorical variables with four levels each, and two are numerical variables. Each feature was assigned its own block. One-hot encoding was performed for the two blocks containing categorical features, leading to two blocks (B1 and B2) of four binary features, each. Blocks corresponding to numerical features (B3 and B4) contain one feature each. In the 30 M-ANN model runs, an average performance of 0.21 (RMSEIQR) and 0.87 (R2) was obtained without outliers. Figure 5b shows that

for all four paradigms, block B3 was most important. While VarGrad methods delivered a binary ranking, knock-in and knock-out suggested a ranking with 3 and 4 distinct importance levels, respectively—thus, the level of detail was higher in the MI-based rankings compared to VarGrad methods.

4 Discussion

Our experiments demonstrated several differences between the proposed strategies. While the composite strategy evaluates features individually and depends on two user-selected parameters (the feature-wise ranking scheme and the summary metric), the knock-in and knock-out strategies consider each block a closed unit. They require no selection of a summary statistic. MI-based rankings deliver a score in $[0, 1]$, while VarGrad has no upper bound. However, the discretization associated with the mutual information calculation may influence the importance scores and, thus, the rankings by knock-in and knock-out. All strategies are applicable for multivariate target variables, as well. However, an MI-based comparison between outputs and pseudo-outputs is prone to suffer from the curse of dimensionality since higher-dimensional probability distributions are compared to each other. On the contrary, the vanishing gradient problem can influence VarGrad in deep ANN architectures. All approaches delivered accurate experimental results, but only knock-in and knock-out provided a consistent ranking of blocks with minor importance in dominant blocks, such as for the servo dataset.

Even though knock-in and knock-out rely on the same concept of assessing pseudo-outputs related to each block, their properties and interpretations differ. The knock-in strategy determines whether a block can deliver a reasonable target description independently from the remaining blocks. This interpretation of block importance evaluates the performance achieved if we reduce the model to solely one block at a time. In contrast, knock-out quantifies whether the contribution of a block can be compensated by any other block. If two blocks contain redundant information about the target, knock-in delivers high values for both blocks since each block individually has high explanatory power. In contrast, knock-out penalizes redundant blocks since each of them can be removed without loss of information. This property became evident in the BCW experiment, where B3 was dominant but shared overlapping information with B1: knock-in was the only approach that discovered the higher information content in B1 compared to the uninformative B2.

5 Conclusion

We have demonstrated three strategies to rank the importance of feature-blocks as post-processing in ANNs with block-wise input structure. The composite strategy, which is a direct generalization of feature-wise importance rankings, provided promising results in most cases, but selecting the correct summary statistic was crucial. Knock-in and knock-out strategies, implemented using an information-theoretic measure on the model outputs, delivered a trade-off

between the extremes of maximum and mean feature importance in the composite case. All methods uncovered the true block importance with high accuracy and delivered new insights into the ANN's behavior. Still, computing multiple proposed metrics is advantageous for making informative block ranking decisions.

Acknowledgment. The authors gratefully acknowledge the financial support from internal funding scheme at Norwegian University of Life Sciences (project number 1211130114), which financed the international stay at the University of British Columbia, and thereby fostered the completion of this work.

References

1. Adebayo, J., Gilmer, J., Goodfellow, I., Kim, B.: Local explanation methods for deep neural networks lack sensitivity to parameter values. arXiv (2018)
2. Adebayo, J., Gilmer, J., Muelly, M., Goodfellow, I., Hardt, M., Kim, B.: Sanity checks for saliency maps. In: *Advances in Neural Information Processing Systems*, vol. 31 (2018)
3. Alnemer, L.M., et al.: Multiple sources classification of gene position on chromosomes using statistical significance of individual classification results. In: *International Conference on Machine Learning and Applications and Workshops*, vol. 1, pp. 7–12 (2011). <https://doi.org/10.1109/ICMLA.2011.101>
4. Amjad, R.A., Geiger, B.C.: Learning representations for neural network-based classification using the information bottleneck principle. *IEEE Trans. Pattern Anal. Mach. Intell.* **42**(9), 2225–2239 (2019). <https://doi.org/10.1109/TPAMI.2019.2909031>
5. Cao, B., He, L., Kong, X., Philip, S.Y., Hao, Z., Ragin, A.B.: Tensor-based multi-view feature selection with applications to brain diseases. In: *IEEE International Conference on Data Mining*, pp. 40–49 (2014)
6. Cover, T., Thomas, J.: *Elements of Information Theory*. Wiley, Hoboken (2012)
7. Dagnely, P., Tourwé, T., Tsiorkova, E.: Annotating the performance of industrial assets via relevancy estimation of event logs. In: *IEEE International Conference on Machine Learning and Applications (ICMLA)*, pp. 1261–1268 (2018). <https://doi.org/10.1109/ICMLA.2018.00205>
8. Ghorbani, A., Abid, A., Zou, J.: Interpretation of neural networks is fragile. In: *AAAI Conference on Artificial Intelligence*, vol. 33, pp. 3681–3688 (2019)
9. Hooker, S., Erhan, D., Kindermans, P.J., Kim, B.: A benchmark for interpretability methods in deep neural networks. In: *Advances in Neural Information Processing Systems*, vol. 32 (2019)
10. Jenul, A., Schrunner, S., Liland, K.H., Indahl, U.G., Futsaether, C.M., Tomic, O.: Rent-repeated elastic net technique for feature selection. *IEEE Access* **9**, 152333–152346 (2021)
11. Jenul, A., Schrunner, S., Pilz, J., Tomic, O.: A User-Guided Bayesian Framework for Ensemble Feature Selection in Life Science Applications (UBayFS). arXiv (2021)
12. Quinlan, J.R.: Combining instance-based and model-based learning. In: *International Conference on Machine Learning*, pp. 236–243 (1993)

13. Street, W.N., Wolberg, W.H., Mangasarian, O.L.: Nuclear feature extraction for breast tumor diagnosis. In: Acharya, R.S., Goldgof, D.B. (eds.) *Biomedical Image Processing and Biomedical Visualization*, vol. 1905, pp. 861–870. SPIE (1993). <https://doi.org/10.1117/12.148698>
14. Wojtas, M., Chen, K.: Feature importance ranking for deep learning. *Adv. Neural. Inf. Process. Syst.* **33**, 5105–5114 (2020)
15. Yu, R., et al.: NISP: pruning networks borisusing neuron importance score propagation. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9194–9203 (2018)
16. Yuan, M., Lin, Y.: Model selection and estimation in regression with grouped variables. *J. Roy. Stat. Soc. Ser. B (Stat. Methodol.)* **68**(1), 49–67 (2006)