

mlr3summary: Concise and interpretable summaries for machine learning models

Susanne Dandl^{1,2}, Marc Becker^{1,2}, Bernd Bischl^{1,2}, Giuseppe Casalicchio^{1,2} and Ludwig Bothmann^{1,2,*}

¹Ludwig-Maximilians-Universität München (LMU Munich), Ludwigstrasse 33, 80539, Munich, Germany

²Munich Center for Machine Learning (MCML), Munich, Germany

Abstract

This work introduces a novel R package for concise, informative summaries of machine learning models. We take inspiration from the summary function for (generalized) linear models in R, but extend it in several directions: First, our summary function is model-agnostic and provides a unified summary output also for non-parametric machine learning models; Second, the summary output is more extensive and customizable – it comprises information on the dataset, model performance, model complexity, model’s estimated feature importances, feature effects, and fairness metrics; Third, models are evaluated based on resampling strategies for unbiased estimates of model performances, feature importances, etc. Overall, the clear, structured output should help to enhance and expedite the model selection process, making it a helpful tool for practitioners and researchers alike.

Keywords

Model summary, interpretable machine learning, resampling-based evaluation

1. Introduction

Machine learning (ML) increasingly supports decision-making processes in various domains. A data scientist has a wide range of models available, ranging from intrinsically interpretable models such as linear models to highly complex models such as random forests or gradient boosted trees. Intrinsically interpretable models can come at the expense of generalization performance, i.e., the model’s capability to predict accurately on future data. Being able to interpret predictive models is either often a strict requirement for scientific inference or at least a very desirable property to audit models in other (more technical) contexts. Many methods have been proposed for interpreting black-box ML models in the field of interpretable ML (IML).

For comparing (generalized) linear models (GLMs), the **stats** package in R offers a summary function, which only requires the model (fitted with `lm` or `glm`) as input. As an example, `glm` is

Late-breaking work, Demos and Doctoral Consortium, colocated with The 2nd World Conference on eXplainable Artificial Intelligence: 17-19 July, 2024, Valletta, Malta

*Corresponding author.

✉ susanne.dandl@stat.uni-muenchen.de (S. Dandl); marc.becker@stat.uni-muenchen.de (M. Becker); bernd.bischl@stat.uni-muenchen.de (B. Bischl); giuseppe.casalicchio@stat.uni-muenchen.de (G. Casalicchio); ludwig.bothmann@stat.uni-muenchen.de (L. Bothmann)

🆔 0000-0003-4324-4163 (S. Dandl); 0000-0002-8115-0400 (M. Becker); 0000-0001-6002-6980 (B. Bischl); 0000-0001-5324-5966 (G. Casalicchio); 0000-0002-1471-6582 (L. Bothmann)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

applied to a preprocessed version of the German credit dataset [1] (available in the package via `data("credit", package = "mlr3summary")`):

```
> logreg = glm(risk ~., data = credit, family = binomial(link = "logit"))
> summary(logreg)
```

Call:

```
glm(formula = risk ~., data = credit, family = binomial(link = "logit"))
```

Coefficients:

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	1.057e+00	3.646e-01	2.900	0.00373 **
age	9.103e-03	8.239e-03	1.105	0.26925

...

Residual deviance: 656.19 on 515 degrees of freedom

AIC: 670.19

...

This (shortened) summary informs about the significance of variables ($\text{Pr}(>|z|)$), their respective effect size and direction (*Estimate*), as well as the goodness-of-fit of the model (*Residual deviance* and *AIC*). Unfortunately, many other non-parametric ML models currently cannot be analyzed similarly: either targeted implementations exist for specific model classes, or an array of different model-agnostic interpretability techniques (e.g., to derive feature importance) scattered across multiple packages [2, 3, 4] must be employed. However, especially in applied data science, a user often performs model selection or model comparison across an often diverse pool of candidate models, so a standardized diagnostic output becomes highly desirable.

Another issue is that in the `glm`-based summary, the goodness-of-fit is only evaluated on the training data, but not on hold-out/test data. While this might be appropriate for GLM-type models – provided proper model diagnosis has been performed – this is not advisable for non-parametric and non-linear models, which can overfit the training data.¹ Here, hold-out test data or in general resampling techniques like cross-validation should be used for proper estimation of the generalization performance [5]. Such resampling-based performance estimation should also be used for loss-based IML methods. For interpretability methods that only rely on predictions, this might also be advisable but might not lead to huge differences in results [6, 7].

Contributions With the `mlr3summary` package, we provide a novel model-agnostic `summary` function for ML models and learning algorithms in R. This is facilitated by building upon `mlr3` [8, 9] – a package ecosystem for applied ML, including resampling-based performance assessment. The `summary` function returns a structured overview that gives information on the underlying dataset and model, generalization performances, complexity of the model, fairness metrics, and feature importances and effects. For the latter two, the function relies on model-agnostic methods from the field of IML. The output is customizable via a flexible `control` argument to allow adaptation to different application scenarios. The `mlr3summary` package is released under LGPL-3 on CRAN and GitHub (<https://github.com/mlr-org/mlr3summary>). Documentations in the form of help pages are available as well as unit tests. The code to reproduce this paper is part of the package and GitHub repository (`inst/` or `demo/`).

¹For completeness' sake: Overfitting *can* happen for GLMs, e.g., in high-dimensional spaces with limited sample size.

2. Related work

Most R packages that offer model summaries are restricted to parametric models and extend the `stats` summary method (e.g., `modelsummary` [10] or `broom` [11]). Performance is only assessed based on training data – generalization errors are not provided. Packages that can handle diverse ML models focus primarily on performance assessment (e.g., `mlr3` [8] or `caret` [12]). Packages that primarily consider feature importances and effects do not provide overviews in a concise, decluttered format but provide extensive reports (e.g., `modelDown` [13] and `modelStudio` [14] based on `DALEX` [3], or `explainer` [4]). While it is possible to base the assessment on hold-out/test data, assessment based on resampling is not automatically supported by these packages. Overall, to the best of our knowledge, there is no R package yet that allows for a concise yet informative overview based on resampling-based performance assessment, model complexity, feature importance and effect directions, and fairness metrics.

3. Design, functionality, and example

The core function of the `mlr3summary` package is the S3-based `summary` function for `mlr3` `Learner` objects. It has three arguments: `object` reflects a trained model – a model of class `Learner` fitted with `mlr3`; `resample_result` reflects the results of resampling – a `ResampleResult` object fitted with `mlr3`; `control` reflects some control arguments – a list created with `summary_control` (details in Section 3.2).

The `mlr3` package is the basis of `mlr3summary` because it provides a unified interface to diverse ML models and resampling strategies. A general overview of the `mlr3` ecosystem is given in Bischl et al. [9]. With `mlr3`, the modelling process involves the following steps: (1) initialize a regression or classification task, (2) choose a regression or classification learner,² (3) train a model with the specified learner on the initialized task, (4) apply a resampling strategy. The last step is necessary to receive valid estimates for performances, importances, etc., as mentioned in Section 1. The following lines of code illustrate steps (1)-(4) on the (preprocessed) credit dataset from Section 1 using a `ranger` random forest and 3-fold cross-validation.

```
> task = TaskClassif$new(id = "credit", backend = credit, target = "risk")
> rf = lrn("classif.ranger", predict_type = "prob")
> rf$train(task)
> cv3 = rsmp("cv", folds = 3L)
> rr = resample(task = task, learner = rf, resampling = cv3,
+   store_models = TRUE)
```

Internally, the `resample` function fits, in each iteration, the model on the respective training data, uses the model to predict the held-out test data, and stores the predictions in the result object. To compute performances, complexities, importances, and other metrics, the `summary` function iteratively accesses the models and datasets within the resulting object, which requires setting the parameter `store_models = TRUE` within `resample`. For the final summary output, the results of each iteration are aggregated (e.g., averages and standard deviations (sds)).

²Custom learners can also be added to `mlr3`, see <https://mlr3extralearners.ml-org.com/articles/extending.html>.

3.1. Summary function and output

This section shows the `summary` call and output for the random forest of the previous credit example and provides some details on each displayed paragraph.

```
> summary(object = rf, resample_result = rr)
```

```
— General —  
Task type: classif  
Target name: risk (good and bad)  
Feature names: age, credit.amount, duration, saving.accounts, and sex  
Model type: classif.ranger with num.threads=1  
Resampling: cv with folds=3  
  
— Residuals —  
      Min      1Q  Median      Mean      3Q      Max  
0.06633 0.28596 0.41013 0.43297 0.56701 0.94773  
  
— Performance [sd] —  
  
↑classif.auc (macro):  0.6812 [0.0489]  
↑classif.fbeta (macro): 0.6917 [0.0578]  
↓classif.bbrier (macro): 0.2262 [0.0239]  
↑classif.mcc (macro):  0.2622 [0.079]  
  
— Complexity [sd] —  
  
sparsity:                5 [0]  
interaction_strength: 0.5935 [0.1207]  
  
— Importance [sd] —  
  
duration      pdp      pfi.ce  
0.1653 [0.0248] 0.0843 [0.0489]  
credit.amount 0.1274 [0.0139] 0.0421 [0.0176]  
saving.accounts 0.0878 [0.0514] 0.0096 [0.0033]  
age           0.0485 [0.0135] -0.0038 [0.0317]  
sex          0.0301 [0.0232] -0.0153 [0.0088]  
  
— Effects —  
  
duration      pdp  ale  
credit.amount  pdp  ale  
saving.accounts pdp  ale  
age            pdp  ale  
sex            pdp  ale
```

General provides an overview of the task, the learner (including its hyperparameters), and the resampling strategy.³

Residuals display the distribution of residuals of *hold-out* data over the resampling iterations. For regression models, the residuals display the difference between true and predicted outcome. For classifiers that return class probabilities, the residuals are defined as the difference between predicted probabilities and a one-hot-encoding of the true class. For classifiers that return classes, a confusion matrix is shown.

Performance displays averages and sds (in `[]`) of performance measures over the iterations.⁴

³Currently, this is the only paragraph that is based on `object`, all other paragraphs are based on `resample_result`.

⁴Please note that there is no unbiased estimator of the variance, see [15] and Section 5 for a discussion.

The shown performance values are the area-under-the-curve (auc), the F-score (fbeta), the binary Brier score (bbrier), and Mathew's correlation coefficient (mcc). The arrows display whether lower or higher values refer to a better performance. "(macro)" indicates a macro aggregation, i.e., measures are computed for each iteration separately before averaging. "(micro)" would indicate that measures are computed across all iterations (see [9] for details).

Complexity displays averages and sds of two model complexity measures proposed by Molnar et al. [16]: `sparsity` shows the number of used features that have a non-zero effect on the prediction (evaluated by accumulated local effects (ale) [17]); `interaction_strength` shows the scaled approximation error between a main effect model (based on ale) and the prediction function.⁵

Importance shows the averages and sds of feature importances over the iterations. The first column (pdp) displays importances based on the sds of partial dependence curves [18, 19], the second column (pfi.ce) shows the results for permutation feature importance [20, 21].

Effects shows average effect plots over the iterations – partial dependence plots (pdp) [18] and ale plots [17]. For binary classifiers, the effect plots are only shown for the positively-labeled class (here, `task$positive = "good"`). For multi-class classifiers, the effect plots are given for each outcome class separately (one vs. all). For categorical features, the bars are ordered according to the factor levels of the feature.

The learner can also be a complete pipeline from **mlr3pipelines** [22], where the most common case would be an ML model with associated pre-processing steps. Then, the `summary` output also shows some basic information about the pipeline.⁶ Since preprocessing steps are treated as being part of the learner, the summary output is displayed on the original data (e.g., despite one-hot encoding of categorical features, importance results are not shown for each encoding level separately). The learner can also be an `AutoTuner` from **mlr3tuning**, where automatic processes for tuning the hyperparameters are conducted. Examples on pipelining and tuning are given in the demo of the package.

3.2. Customizations

The output of the `summary` function can be customized via a `control` argument which requires a list created with the function `summary_control` as an input. If no control is specified, the following default setting is used:

```
> summary_control(measures = NULL,  
+ complexity_measures = c("sparsity", "interaction_strength"),  
+ importance_measures = NULL, n_important = 15L,  
+ effect_measures = c("pdp", "ale"),  
+ fairness_measures = NULL, protected_attribute = NULL,  
+ hide = NULL, digits = max(3L, getOption("digits") - 3L))
```

Performances are adaptable via `measures`, complexities via `complexity_measures`, importances via `importance_measures` and effects via `effect_measures` within `summary_control`. Examples are given in the demo of the package. The default for `measures`

⁵The interaction strength has a value in $[0, 1]$, 0 means no interactions, 1 means no main effects but interactions.

⁶Linear pipelines can be displayed in the console, non-linear parts are suppressed in the output.

and `importance_measures` is `NULL`, which results in a collection of commonly reported measures being chosen, based on the task type – for concrete measures see the help page (`?summary_control`). `n_important` reflects that, by default, only the 15 most important features are displayed in the output. This is especially handy for high-dimensional data. With `hide`, paragraphs of the summary output can be omitted (e.g., "performance") and with `digits`, the number of printed digits is specified.

Fairness assessment for classification and regression models is also available in the `mlr3summary` package based on the `mlr3fairness` package [23]. Therefore, a protected attribute must be specified. This can be done either within the task by updating the feature roles or by specifying a `protected_attribute` in `summary_control`. The following shows the code and output when specifying `sex` as a protected attribute. The shown default fairness measures are demographic parity (dp), conditional use accuracy equality (cuae) and equalized odds (eod), other measures are possible via `fairness_measures` in `summary_control`.

```
> summary(object = rf, resample_result = rr,
+ control = summary_control(protected_attribute = "sex"))
```

```
— Fairness [sd] —
Protected attribute: sex

↓fairness.dp (macro):  0.1206 [0.1021]
↓fairness.cuae (macro): 0.0987 [0.0315]
↓fairness.eod (macro): 0.1224 [0.099]
```

4. Runtime assessment

To assess how the runtime scales with differing numbers of features $p \in \{5, 10, 25, 50, 100\}$ and numbers of observations $n \in \{50, 100, 500, 1000, 2000\}$, we conducted a simulation study. Given $X_1, X_2, X_3 \sim U(0, 1)$, $X_4 \sim \text{Bern}(0.75)$, the data generating process is $y = f(\mathbf{x}) + \epsilon$ with $f(\mathbf{x}) = 4x_1 + 4x_2 + 4x_4x_3^2$ and $\epsilon \sim N(0, 0.1 \cdot f(\mathbf{x}))$. As noise variables, X_5 as a categorical feature with five classes, and $X_6, \dots, X_p \sim N(0, 1)$ were added to the data. We trained random forests and linear main effect models on the datasets and conducted 3-fold cross-validation. The first two figures in Figure 1 show that runtimes of the linear model were lower compared to the random forest. To improve runtimes, we added parallelization over the resampling iterations (via the `future` package [24]) as another feature to `mlr3summary` – results for the random forest (with 3 cores) are on the right. Overall, scaling of runtimes is worse in p than in n .




Figure 1: Runtimes of the `summary` function for linear models (left), and random forests without (middle) and with (right) parallelization, for differing numbers of features p and observations n .

5. Outlook and discussion

In conclusion, this paper introduces a novel R package for concise model summaries. The summary output is highly adaptable due to a `control` argument and might be extended in the future. We also plan to offer a `report` function for detailed visualizations and model comparisons. To assess importance and effects of single features, **mlr3summary** builds upon the **iml** and **fastshap** packages. These packages only offer a limited set of interpretation methods. Recommended alternatives to permutation feature importances like conditional feature importance [6], are currently not available in a proper R package (published on CRAN). Our summary also currently lacks proper statistical tests for importances or confidence intervals for performances. This is because unbiased estimates of the variance are required which is a challenge for resampling strategies and the available methods that propose unbiased estimates are computationally infeasible (e.g., due to many model refits) [7, 25]. Addressing this issue requires some concerted efforts from the research community. If methods are readily available in R, we are happy to integrate them in **mlr3summary**.

Acknowledgments

This work has been partially supported by the Federal Statistical Office of Germany.

References

- [1] H. Hofmann, Statlog (German credit data), UCI Machine Learning Repository (1994). doi:10.24432/C5NC77.
- [2] C. Molnar, B. Bischl, G. Casalicchio, *iml*: An R package for interpretable machine learning, *JOSS* 3 (2018) 786. doi:10.21105/joss.00786.
- [3] P. Biecek, DALEX: Explainers for complex predictive models in R, *Journal of Machine Learning Research* 19 (2018) 1–5.
- [4] R. Zargari Marandi, *explainer*: Machine learning model explainer, 2023. URL: <https://CRAN.R-project.org/package=explainer>, R package version 1.0.0.
- [5] R. Simon, *Resampling strategies for model assessment and selection*, Springer US, Boston, MA, 2007, pp. 173–186. doi:10.1007/978-0-387-47509-7_8.
- [6] C. Molnar, G. König, J. Herbinger, T. Freiesleben, S. Dandl, C. A. Scholbeck, G. Casalicchio, M. Grosse-Wentrup, B. Bischl, General pitfalls of model-agnostic interpretation methods for machine learning models, in: A. Holzinger, et al. (Eds.), *xxAI - Beyond Explainable AI: International Workshop*, Springer International Publishing, Cham, 2022, pp. 39–68. doi:10.1007/978-3-031-04083-2_4.
- [7] C. Molnar, T. Freiesleben, G. König, J. Herbinger, T. Reisinger, G. Casalicchio, M. N. Wright, B. Bischl, Relating the partial dependence plot and permutation feature importance to the data generating process, in: L. Longo (Ed.), *Explainable Artificial Intelligence*, Springer Nature Switzerland, Cham, 2023, pp. 456–479. doi:10.1007/978-3-031-44064-9_24.
- [8] M. Lang, M. Binder, J. Richter, P. Schratz, F. Pfisterer, S. Coors, Q. Au, G. Casalicchio,

- L. Kotthoff, B. Bischl, *mlr3: A modern object-oriented machine learning framework in R*, *Journal of Open Source Software* (2019). doi:10.21105/joss.01903.
- [9] B. Bischl, R. Sonabend, L. Kotthoff, M. Lang, *Applied machine learning using mlr3 in R*, Chapman and Hall/CRC, 2024. doi:10.1201/9781003402848.
- [10] V. Arel-Bundock, *modelsummary: Data and model summaries in R*, *Journal of Statistical Software* 103 (2022) 1–23. doi:10.18637/jss.v103.i01.
- [11] D. Robinson, A. Hayes, S. Couch, *broom: Convert statistical objects into tidy tibbles*, 2023. URL: <https://CRAN.R-project.org/package=broom>, R package version 1.0.5.
- [12] Kuhn, Max, *Building predictive models in R using the caret package*, *Journal of Statistical Software* 28 (2008) 1–26. doi:10.18637/jss.v028.i05.
- [13] K. Romaszko, M. Tatarynowicz, M. Urbański, P. Biecek, *modelDown: Automated website generator with interpretable documentation for predictive machine learning models*, *Journal of Open Source Software* 4 (2019). doi:10.21105/joss.01444.
- [14] H. Baniecki, P. Biecek, *modelStudio: Interactive studio with explanations for ML predictive models*, *Journal of Open Source Software* 4 (2019) 1798. doi:10.21105/joss.01798.
- [15] C. Nadeau, Y. Bengio, *Inference for the generalization error*, in: S. Solla, et al. (Eds.), *Advances in Neural Information Processing Systems*, volume 12, MIT Press, 1999, pp. 1–7.
- [16] C. Molnar, G. Casalicchio, B. Bischl, *Quantifying model complexity via functional decomposition for better post-hoc interpretability*, Springer International Publishing, 2020, p. 193–204. doi:10.1007/978-3-030-43823-4_17.
- [17] D. W. Apley, J. Zhu, *Visualizing the effects of predictor variables in black box supervised learning models*, *Journal of the Royal Statistical Society Series B: Statistical Methodology* 82 (2020) 1059–1086. doi:10.1111/rssb.12377.
- [18] J. H. Friedman, *Greedy function approximation: A gradient boosting machine*, *The Annals of Statistics* 29 (2001). doi:10.1214/aos/1013203451.
- [19] B. M. Greenwell, B. C. Boehmke, A. J. McCarthy, *A simple and effective model-based variable importance measure*, *arXiv preprint arXiv:1805.04755* (2018). doi:10.48550/arXiv.1805.04755.
- [20] L. Breiman, *Random forests*, *Machine Learning* 45 (2001) 5–32. doi:10.1023/a:1010933404324.
- [21] A. Fisher, C. Rudin, F. Dominici, *All models are wrong, but many are useful: Learning a variable’s importance by studying an entire class of prediction models simultaneously*, *Journal of Machine Learning Research* 20 (2019).
- [22] M. Binder, F. Pfisterer, M. Lang, L. Schneider, L. Kotthoff, B. Bischl, *mlr3pipelines - Flexible machine learning pipelines in R*, *Journal of Machine Learning Research* 22 (2021) 1–7.
- [23] F. Pfisterer, W. Siyi, M. Lang, *mlr3fairness: Fairness auditing and debiasing for ‘mlr3’*, 2023. URL: <https://CRAN.R-project.org/package=mlr3fairness>, R package version 0.3.2.
- [24] H. Bengtsson, *A unifying framework for parallel and distributed processing in R using futures*, *The R Journal* 13 (2021) 208–227. doi:10.32614/RJ-2021-048.
- [25] T. H. Stephen Bates, R. Tibshirani, *Cross-validation: What does it estimate and how well does it do it?*, *Journal of the American Statistical Association* (2023) 1–12. doi:10.1080/01621459.2023.2197686.