








**ADVANCED REVIEW** OPEN ACCESS

# A Systematic Categorization of Performance Measures for Estimated Non-Linear Associations Between an Outcome and Continuous Predictors

Theresa Ullmann<sup>1</sup>  | Georg Heinze<sup>1</sup>  | Michal Abrahamowicz<sup>2</sup>  | Aris Perperoglou<sup>3</sup>  | Willi Sauerbrei<sup>4</sup>  | Matthias Schmid<sup>5</sup>  | Daniela Dunkler<sup>1</sup>  | TG2 of the STRATOS Initiative

<sup>1</sup>Institute of Clinical Biometrics, Center for Medical Data Science, Medical University of Vienna, Vienna, Austria | <sup>2</sup>Department of Epidemiology & Biostatistics, McGill University, Montreal, Quebec, Canada | <sup>3</sup>Statistics and Data Science Innovation Hub, GSK, Stevenage, UK | <sup>4</sup>Institute of Medical Biometry and Statistics, Faculty of Medicine and Medical Center, University of Freiburg, Freiburg, Germany | <sup>5</sup>Department of Medical Biometry, Informatics, and Epidemiology, University Hospital Bonn, Bonn, Germany

**Correspondence:** Daniela Dunkler ([daniela.dunkler@meduniwien.ac.at](mailto:daniela.dunkler@meduniwien.ac.at))

**Received:** 19 March 2025 | **Revised:** 29 July 2025 | **Accepted:** 10 August 2025

**Commissioning Editor:** David W Scott | **Editor-in-Chief:** David W Scott

**Funding:** This work was supported by Austrian Science Fund, I-4739-B; Deutsche Forschungsgemeinschaft, SA580/10-3.

**Keywords:** fractional polynomials | functional associations | simulation | smoothing | splines

## ABSTRACT

In regression analysis, associations between continuous predictors and the outcome are often assumed to be linear. However, modeling the associations as non-linear can improve model fit. Many flexible modeling techniques, like (fractional) polynomials and spline-based approaches, are available. Such methods can be systematically compared in simulation studies, which require suitable performance measures to evaluate the accuracy of the estimated curves against the true data-generating functions. Although various measures have been proposed in the literature, no systematic overview exists so far. To fill this gap, we introduce a categorization of performance measures for evaluating estimated non-linear associations between an outcome and continuous predictors. This categorization includes many commonly used measures. The measures can not only be used in simulation studies, but also in application studies to compare different estimates to each other. We further illustrate and compare the behavior of different performance measures through some examples and a Shiny app.

This article is categorized under:

Statistical and Graphical Methods of Data Analysis > Modeling Methods and Algorithms


## 1 | Introduction

Regression models often describe associations between an outcome and one or multiple predictors (also referred to as explanatory variables, covariates, or independent variables). Determining the functional forms of the continuous predictors in such models is an important aspect of model building

and may substantially affect the validity and interpretation of the model. Often, the associations of the continuous predictors with the outcome are simply assumed to be linear (potentially after transformation of the expected outcome via a link function). However, modeling the effect of one or multiple predictors as *non-linear* can be more appropriate in many applications. Consider the example in Figure 1a, which shows the

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2025 The Author(s). *WIREs Computational Statistics* published by Wiley Periodicals LLC.



**FIGURE 1** | Association of age in years with BMI, estimated from 9377 participants of NHANES. The black line connects the mean BMI values at each age. (a) The green line shows the linear fit. The shaded areas show the 95% confidence intervals. (b) Blue line: Fit with fractional polynomials, orange line: Fit with linear B-splines with three basis functions, pink line: Fit with cubic B-splines with 15 basis functions. For the B-spline fits, the inner knots are positioned at suitable quantiles of age.


association of age with body mass index (BMI) based on data from 9377 participants of the National Health and Nutrition Examination Survey (NHANES) (CDC and NCHS 2024). The black curve connects the mean BMI values at each age. To account for the sampling variation reflected in the excessive wiggleness of the black curve, we can try to find a smooth model fit to better describe the systematic changes in the mean BMI values with increasing age. The linear fit, shown in green color, does not describe the mean BMI values well (by definition, the linear fit is unable to account for the clear non-monotonicity of the relationship), such that options for a non-linear functional form should be explored.

There are many different methods to extend the framework of linear predictor models in order to model non-linear associations. For example, the effect of a predictor can be modeled with a polynomial function. The *fractional polynomial approach* (Royston and Sauerbrei 2008) extends the simple polynomial approach; here one or several power transformations are applied to a predictor. Such transformations are global, that is, a (fractional) polynomial function is defined over the whole range of the predictor. In contrast, *spline-based approaches* are local smoothers: the range of the predictor is divided into segments, a different polynomial is defined on each segment, and the polynomials (all having the same degree) are smoothly joined at the knots, that is, the endpoints of the segments, to ensure continuity of the fitted function. Spline-based approaches include, for example, B-splines (De Boor 1978), restricted cubic splines (Stone and Koo 1985),

P-splines (Eilers and Marx 1996), and thin-plate regression splines (including smoothing splines) (Wood 2003); see Perperoglou et al. (2019) for an overview. Figure 1b again shows the NHANES example, now with a fractional polynomial fit and two different B-spline fits (linear B-splines with three basis functions and cubic B-splines with 15 basis functions). The fitted curves have different properties: all curves reflect the aforementioned non-monotonicity, but they differ with respect to (local) extrema and the degree of wiggleness. It is not immediately clear which curve is the “best one” nor how one should define a criterion for determining such a “best” curve.

Currently, there is a lack of knowledge regarding the performance of different methods for modeling functional associations. Studies comparing different methods to each other are relatively rare (Binder et al. 2013; Strasak et al. 2011). More systematic comparison studies are needed to better understand the properties of different methods and to derive recommendations about which methods might be preferable in which application contexts, depending on the relevant characteristics of the data being analyzed and/or the complexity of the underlying true functional forms. Ideally, comparison studies should be *neutral*, that is, the authors should be (as a group) approximately equally familiar with all compared methods and they should not have a vested interest in one of the methods (Boulesteix et al. 2013, 2017, 2024). Typically, such studies include simulated data with known true data-generating functions (“ground truths”). As an illustrative example, suppose that we perform a simulation study which systematically compares different flexible estimation methods. We choose a true data-generating mechanism inspired by the observed association of age on BMI in the NHANES data. This scenario is depicted in Figure 2. The black line represents the assumed ground truth curve. The fractional polynomial fit and the two B-spline fits represent the curves obtained in one simulation repetition from three methods to be compared. However, it is unclear which criteria should be used to assess their success in describing the true association of the predictor and the outcome. This illustrates a general issue: it is not clear which *performance measures* should be used in systematic comparison studies, that is, which measures should be used to compare the functional forms (linear or non-linear) estimated by the different methods to the ground truth. Several performance measures have been proposed in the literature (Binder et al. 2011; Buchholz et al. 2014; Govindarajulu et al. 2007), but so far we are not aware of a systematic overview and a classification of the measures. In particular, it is not clear which characteristics of the estimated functional forms are addressed by different measures and how similar certain measures might be to each other.

Therefore, the present article aims to characterize and compare performance measures for estimated non-linear associations with an outcome. The choice of suitable performance measures depends on the goal of the statistical modeling. Researchers often distinguish between three different types of statistical modeling: descriptive, predictive, and explanatory (Shmueli 2010). In *descriptive modeling*, the goal is to identify if and how variables are associated with an outcome, which



**FIGURE 2** | Illustration of a hypothetical simulation study. The black line depicts the ground truth curve. The blue line shows a fit with fractional polynomials, the orange line a fit with linear B-splines with three basis functions, and the pink line a fit with cubic B-splines with 15 basis functions. The three fitted lines can be interpreted as the curves obtained in one simulation repetition from three different methods that are compared in the simulation study. The similarity of the estimated curves with the ground truth curve can be quantified by some suitable measure, but it is not immediately clear which measure(s) should be used. Different measures are able to capture different aspects of the curves, for example, the bias of the curves estimated with fractional polynomials or linear B-splines with three basis functions, or the wiggleness of the curve estimated with cubic B-splines with 15 basis functions.

implies for continuous variables to find functional forms that fit the data best (Carlin and Moreno-Betancur 2025). Estimating non-linear associations is relevant for the latter aspect. In *predictive modeling*, the objective is to predict an outcome as accurately as possible. In *explanatory modeling*, the objective is to estimate the causal effect of a specific explanatory variable on the outcome adjusted for confounders. In our article, we focus on appropriate performance measures for non-linear associations in the context of *descriptive modeling*. Estimating non-linear functional forms can also be relevant for predictive and explanatory modeling. We refer to Section 4 for a brief discussion of appropriate performance measures in the context of predictive and explanatory modeling.

Although our main motivation in the present article is the use of performance measures in the context of comparison studies/simulation studies in methodological research, some measures from our categorization can also be helpful in applied studies. For example, when analyzing a specific real-world dataset, researchers might apply two or more different methods to estimate non-linear associations (see, e.g., Govindarajulu et al. (2007)). They can then compare the results using suitable similarity measures to assess how similar they are to each other, rather than determining which result is closer to the ground truth (which is not known in applied studies). In the example in Figure 1b, suitable measures could be used to compare the different fits to each other.

The remainder of this article is structured as follows: in Section 2, we introduce our categorization of performance measures. In Section 3, we illustrate the behavior of different measures based on examples which are available in a Shiny app. We conclude the article with a discussion in Section 4.

## 2 | Categorization of Performance Measures for Estimated Non-Linear Associations

### 2.1 | General Notation

First, we introduce some notation. For observations  $i = 1, \dots, n$ , we assume a data-generating mechanism of the form

$$y_i = \beta_0 + \sum_{j=1}^p f_j(x_{ij}) + \varepsilon_i$$

with continuous outcome  $y_i$ , predictors  $x_{ij}$ ,  $j = 1, \dots, p$  (realizations of random variables  $X_1, \dots, X_p$ ), and error term  $\varepsilon_i$ . The notation easily generalizes to other outcome types by using a suitable link function and modifying the error term. While the model may contain categorical predictors in the form of dummy variables, from now on we only consider continuous predictors without loss of generality. Some of the functions  $f_j$  might be linear, that is,  $f_j(x_{ij}) = \beta_j x_{ij}$ , while others could be non-linear.

The fitted model is denoted as

$$\hat{y}_i = \hat{\beta}_0 + \sum_{j=1}^p \hat{f}_j(x_{ij})$$

Again, some of the estimated functions  $\hat{f}_j$  might be linear, that is,  $\hat{f}_j(x_{ij}) = \hat{\beta}_j x_{ij}$ , while others might be non-linear. Note that for some covariates  $j$ , the fitted function  $\hat{f}_j$  may be linear while the true function  $f_j$  is non-linear, and vice versa.

In the following, we assume that any type of model selection (if performed) has been completed, such that the fitted model(s) under consideration are fixed. Moreover, we only consider “univariable” performance measures in the sense that the estimated curve of each predictor of interest is evaluated individually. More precisely, the estimated curve  $\hat{f}_j$  for each predictor  $X_j$  of interest is compared to the “true” curve  $f_j$ . For notational simplicity, we will assume that the model is univariable with only one continuous predictor ( $p = 1$ ), which allows us to write  $f = f_1$ ,  $\hat{f} = \hat{f}_1$  and  $X = X_1$ . The predictor  $X$  takes values in a space  $\mathcal{X} \subseteq \mathbb{R}$ .  $P_X$  is the probability distribution of  $X$ ,  $F_X$  is the cumulative distribution function, and  $F_X^{-1}$  the corresponding quantile function. The formulas for the performance measures given further below also apply to multivariable models, by considering each predictor of interest separately and setting  $f = f_j$ ,  $\hat{f} = \hat{f}_j$ ,  $X = X_j$  and so on for each index  $j \in \{1, \dots, p\}$  of interest in turn. In that case, each estimate  $\hat{f} = \hat{f}_j$  is adjusted for one or more other variables, that is, the estimated curves of the other predictors in the model.

### 2.2 | Description of the Categorization

The general idea of our categorization of performance measures is that the measures can be described by different aspects. Consider, for example, the following measure:

$$\int_{\mathcal{X}} |\hat{f}(x) - f(x)| dx \quad (1)$$

This measure integrates the absolute difference between the estimate  $\hat{f}$  and the true  $f$  over  $\mathcal{X}$ . There are several characteristics of this measure. For example, the integral is used for summarizing the absolute difference between  $\hat{f}$  and  $f$ , but different options for aggregation would be possible (e.g., calculating the integral with respect to  $dF_X$  instead of  $dx$ , or taking the maximum absolute difference over  $\mathcal{X}$ ). The measure aggregates the information across the entire range  $\mathcal{X}$ . Alternatives are possible, for example, restricting the range to a subset of  $\mathcal{X}$  (such as the range between the 5% and 95% quantile of  $F_X$ ), with the extreme case that the difference between  $\hat{f}$  and  $f$  is only evaluated at a specific point  $x^* \in \mathcal{X}$ . Moreover, instead of the absolute loss, other types of losses might be considered, such as the squared loss. Finally, instead of comparing  $\hat{f}$  with  $f$ , one might also wish to compare the derivatives  $\hat{f}'$  with  $f'$  and/or  $\hat{f}''$  with  $f''$ .

This example demonstrates that there are several aspects of performance measures and that different performance measures can be defined by specific combinations of these aspects. This leads to our proposed categorization of performance measures as depicted in a schematic overview in Figure 3. First, we describe the three main aspects shown in the blue panels.

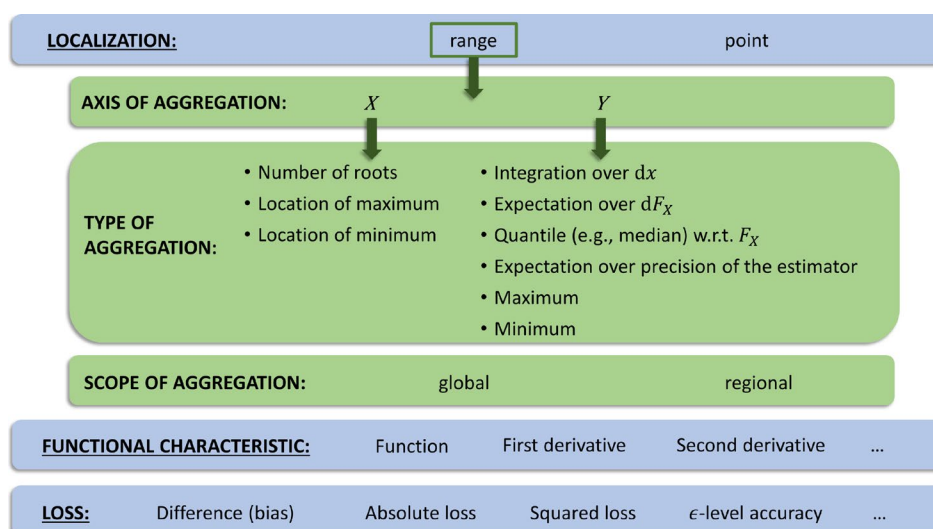
- The aspect *localization* distinguishes between measures that aggregate information over a *range* (there are different options for the scope of aggregation, as described below) and measures that only take performance at a specific *point*  $x^* \in \mathcal{X}$  into account.
- We can also categorize performance measures based on the *functional characteristic*, that is, whether they evaluate the *function*, the *first derivative of the function*, or the *second derivative of the function*, provided that these derivatives exist. When the functional characteristic is the first derivative, the slopes of the curves are compared; for the second derivative, the “wiggleness” of the curves is compared (Green and Silverman 1993). While the third, fourth, etc. derivatives could also be considered, these are

generally more difficult to interpret and are thus not commonly used in practice.

- Finally, different types of *losses* can be chosen to measure the deviation of the estimated curve from the true curve: *difference loss (bias)*, *absolute loss*, *squared loss*, or  *$\epsilon$ -level accuracy*. The  $\epsilon$ -level accuracy (Hirji et al. 1989) evaluates in a binary way whether two values are “close” to each other, and is defined as  $\mathbb{1}_{\{|z_1 - z_2| \leq \epsilon\}}$  for two real numbers  $z_1, z_2$ , with  $\epsilon > 0$  a suitable (small) value, see below for further details. We consider these four types of losses due to their popularity; other loss functions could, of course, also be used.

For measures that aggregate the information over a *range*, there are some further aspects to consider (panels in green color in Figure 3):

- Most measures quantify the deviation from the truth on the *Y axis*. For this purpose, there are different *types of aggregation*: integration over  $dx$ , expectation over  $dF_X$ , a quantile (e.g., the median) with respect to  $F_X$ , expectation over the precision of the estimator, maximum, or minimum.
- Other measures compare the curves on the *X axis*, that is, how the estimated curve and the true curve differ from each other with respect to the number of roots (i.e., the points where the functions or their derivatives take the value zero), the location of the global maximum in the range, or the location of the global minimum in the range.
- For both types of performance measures (axis of aggregation = “Y” and axis of aggregation = “X”), the scope of aggregation has to be specified. *Global* measures consider the information on the whole range  $\mathcal{X}$  of  $X$ , while *regional* measures consider a subrange of  $X$ . A typical subrange is the range between a lower quantile  $l$  and an upper quantile  $u$  of the distribution  $F_X$  where the variable often takes values (e.g., the 5% quantile  $F_X^{-1}(0.05)$  and the 95% quantile



**FIGURE 3** | Categorization of performance measures. Each performance measure can be obtained by choosing specific options for each of the aspects (localization, functional characteristic, etc.).

$F_X^{-1}(0.95)$ ). Here, the motivation is that the behavior of the curves might be less relevant in the boundaries of the range, that is, beyond  $l$  and  $u$ . Restricting the range might also be necessary to obtain well-defined and finite values for the measure, see below for more details. While the range between a lower and an upper quantile appears to be the most common type of subset considered in the literature (Binder et al. 2011; Buchholz et al. 2014), the scope of aggregation could in principle also be another kind of subset of  $\mathcal{X}$ . For example, in medical applications, the region above a critical threshold of a clinical predictor may be of particular interest.

### 2.3 | Some Examples of Measures

Each performance measure arises from choosing specific options for each of the aspects stated above. For example, the measure given in (1) is the measure with localization = “range”, axis of aggregation = “Y”, type of aggregation = “integration over  $dx$ ”, scope of aggregation =  $\mathcal{X}$ , functional characteristic = “function”, and loss = “absolute”. If the type of aggregation is the expectation over  $dF_X$  instead, we obtain the following measure:

$$\int_{\mathcal{X}} |\hat{f}(x) - f(x)| dF_X(x) \quad (2)$$

Compared to measure (1), measure (2) gives more weight to the areas with more data, and less weight to the areas where the data is sparse.

As another example, the following measure results from setting localization = “range”, axis of aggregation = “X”, type of aggregation = “location of maximum”, scope of aggregation = “global”, functional characteristic = “first derivative”, and loss = “absolute”:

$$|\arg \max_{x \in \mathcal{X}} \hat{f}'(x) - \arg \max_{x \in \mathcal{X}} f'(x)| \quad (3)$$

This measure compares the locations of the maxima of the first derivatives, that is, the positions of the steepest ascents of  $\hat{f}$  versus  $f$ . Suppose we are less interested in the absolute difference of the locations of the maxima per se, but only in whether the maxima are located within some distance  $\epsilon$  of each other. If so, then switching from the absolute loss to the  $\epsilon$ -level accuracy results in the measure

$$\mathbb{1}_{\left\{ \left| \arg \max_{x \in \mathcal{X}} \hat{f}'(x) - \arg \max_{x \in \mathcal{X}} f'(x) \right| \leq \epsilon \right\}} \quad (4)$$

which evaluates whether the positions of the steepest ascents of  $\hat{f}$  and  $f$  are sufficiently close. For all measures involving the  $\epsilon$ -level accuracy, suitable values for  $\epsilon$  have to be chosen. This depends on how “strict” a researcher wishes the measure to be, as well as on the domain of  $f$  (if axis of aggregation = “X”) or the range of  $f$  (if axis of aggregation = “Y”). For example, suppose that  $f$  is defined on the domain  $[0, 1]$ . Then for a measure as in (4), where axis of aggregation = “X”, for example  $\epsilon = 0.05$  could be a reasonable choice (evaluating whether the positions of steepest ascents differ by at most 0.05). Generally, one could

choose  $\epsilon$ , for example, as 0.05 times the length of the domain/range of  $f$ , if the domain/range is a bounded interval, or 0.05 times the length of the interval between suitable quantiles of  $F_X$  (as described above) if the domain of  $f$  is unbounded.

The examples of measures given so far all have localization = “range”. In contrast, point-specific measures compare the estimated curve and the true curve at a single point  $x^* \in \mathcal{X}$ . For example, the following measure is defined by localization = “point”, functional characteristic = “function”, and loss = “absolute”:

$$|\hat{f}(x^*) - f(x^*)| \quad (5)$$

This measure compares the values of  $\hat{f}$  and  $f$  at the point  $x^*$ .

When using point-specific measures, it typically makes sense to consider multiple points  $x^* \in \mathcal{X}$ . These points could correspond, for example, to different quantiles of the distribution  $F_X$ . Returning to the NHANES example from Section 1, it could be of interest to compare the estimated to the true BMI values at different quantiles of age.

As mentioned above, when comparing the estimated curve to the ground truth, the primary focus is often on the region where the variable  $X$  typically takes its values. In such cases, regional measures (e.g., with scope of aggregation =  $[F_X^{-1}(0.05), F_X^{-1}(0.95)]$ ) can be used to aggregate the deviation of the estimate from the truth across that region. However, in other cases (such as when extrapolation is a key goal), the behavior of the curves in the boundaries of the range of  $X$  becomes important. For these situations, point-specific measures can be particularly useful, with the points chosen from the tails of the distribution  $F_X$ .

### 2.4 | Examples From the Literature

Our categorization of measures includes several measures that were already proposed in the literature. For example, Buchholz et al. (2014) (see also Govindarajulu et al. 2007) discussed the following measure:

$$\int_{F_X^{-1}(0.01)}^{F_X^{-1}(0.99)} |\hat{f}(x) - f(x)| \hat{p}(x) dx \quad (6)$$

In the terminology of our categorization, this is the measure with localization = “range”, axis of aggregation = “Y”, type of aggregation = “expectation over the precision of the estimator”, scope of aggregation =  $[F_X^{-1}(0.01), F_X^{-1}(0.99)]$ , functional characteristic = “function”, and loss = “absolute”. Here,  $\hat{p}(x)$  denotes the estimate for the precision of the estimator  $\hat{f}$  at point  $x$  (typically defined via the inverse of the variance of the estimator, Buchholz et al. 2014). Note that Buchholz et al. (2014) proposed this measure in a slightly different context, namely for modeling time-varying effects in survival analysis with non-linear functional forms, but the measure can also be used in our context, that is, for modeling the effects of predictors. The idea behind taking the expectation over the precision of the estimator is to give more weight to the areas where the estimation is more precise (i.e., where the

standard error of  $\hat{f}(x)$  is small). In case of cubic regression splines with pre-specified knot positions, the precision of  $\hat{f}(x)$  could, for example, be obtained by evaluating the inverse of the expression  $B(x)\hat{\Sigma}B(x)^\top$ , where  $B(x)$  is the row vector of basis functions (evaluated at  $x$ ) and  $\hat{\Sigma}$  is the estimated covariance matrix of the respective spline coefficient estimates (obtained from standard maximum likelihood fitting of an additive model, see Harrell 2015). In case of P-splines, an estimate of the precision of  $\hat{f}(x)$  could be derived analogously, taking a Bayesian view of the smoothing process and using the covariance of the posterior distribution of the spline coefficients (see section 6.10 of Wood 2017). The precision of fractional polynomials could be evaluated using bootstrap methods (see, e.g., Royston and Sauerbrei 2009).

Measures such as (1), (2), (5) and (6) are defined with functional characteristic = "function". In some cases, the values  $\hat{f}(x)$  of the estimated curve might be quite close to the values  $f(x)$  of the true curve for most  $x$  values, but the curves still have rather different shapes. This will not be detected by measures with functional characteristic = "function", but could be detected by measures which consider the first and/or second derivatives of the curves. Binder et al. (2011) therefore proposed the following measure:

$$\int_{F_X^{-1}(0.05)}^{F_X^{-1}(0.95)} (\hat{f}'(x) - f'(x))^2 dF_X(x) \quad (7)$$

This is the measure with localization = "range", axis of aggregation = "Y", type of aggregation = "expectation over  $dF_X$ ", scope of aggregation =  $[F_X^{-1}(0.05), F_X^{-1}(0.95)]$ , functional characteristic = "first derivative", and loss = "squared". The measure evaluates how similar the slopes of the two curves are, with larger differences being punished by the squared loss. An analogous measure could be defined for functional characteristic = "second derivative" to compare the amount of "wiggleness" of the two curves.

## 2.5 | Formulas for All Measures

The formulas for all performance measures considered in this paper are given in Table 1 (measures with localization = "range" and axis of aggregation = "Y"), Table 2 (measures with localization = "range" and axis of aggregation = "X") and Table 3 (measures with localization = "point"). Formulas are given only for functional characteristic = "function"; the analogous formulas for the derivatives can be obtained by replacing  $f, \hat{f}$  with the respective derivatives. For the tables, some further notation is required. We use the symbol  $S$  for the scope of aggregation, that is,  $S \subseteq \mathcal{X}$  where either (i)  $S = \mathcal{X}$  (if the scope of aggregation is the whole range of  $X$ ) or (ii)  $S = [F_X^{-1}(l), F_X^{-1}(u)]$  (if the scope of aggregation is between the  $l$ 'th and  $u$ 'th quantile of  $F_X$ ). The definition of  $S$  could also be generalized to denote an arbitrary subset of  $\mathcal{X}$ .

For measures with type of aggregation = "quantile with respect to  $F_X$ ", let  $F_{g(X_S)}$  denote the distribution of a random variable  $g(X_S)$  that arises by sampling  $X$  values from  $P_X$  and discarding values outside of  $S$  (i.e., sampling from the distribution

truncated to  $S$ ), and then applying some transformation  $g$  to the remaining values (e.g.,  $g(x) = |\hat{f}(x) - f(x)|$ ). The corresponding quantile function is denoted as  $F_{g(X_S)}^{-1}$ . For example, this notation allows us to express the performance measure with localization = "range", axis of aggregation = "Y", type of aggregation = "median with respect to  $F_X$ ", scope of aggregation =  $S = [F_X^{-1}(l), F_X^{-1}(u)]$ , functional characteristic = "function", and loss = "absolute" with the following formula:

$$F_{|\hat{f}(X_S) - f(X_S)|}^{-1} \quad (0.5)$$

That is, values of  $X$  are sampled between  $F_X^{-1}(l)$  and  $F_X^{-1}(u)$ , the absolute differences  $|\hat{f}(x) - f(x)|$  for all sampled values  $x$  are calculated, and the median of these absolute differences is determined.

For measures with axis of aggregation = "X", we define the following quantities:

Location of maxima of  $f, \hat{f}$  on  $S$ :

$$x^{f, \max S} = \arg \max_{x \in S} f(x), x^{\hat{f}, \max S} = \arg \max_{x \in S} \hat{f}(x)$$

Location of minima of  $f, \hat{f}$  on  $S$ :

$$x^{f, \min S} = \arg \min_{x \in S} f(x), x^{\hat{f}, \min S} = \arg \min_{x \in S} \hat{f}(x)$$

Number of roots of  $f, \hat{f}$  on  $S$ :

$$\# \text{ roots}^S(f), \# \text{ roots}^S(\hat{f})$$

For Tables 1 and 2, we generally assume that the respective values (integrals, optima, etc.) exist and are well-defined and finite. In practice, this might not always be the case. For example, global measures might be ill-defined if the curves or their derivatives tend towards infinity at the boundaries of the full range. Moreover, some estimators tend to be very unstable in the tails of the distribution  $F_X$ , inducing a risk that global performance measures will be excessively impacted by the values estimated in regions where there is little data. In these cases, it can help to switch from global to regional measures by restricting the scope of aggregation from the whole range  $\mathcal{X}$  to the interval  $[F_X^{-1}(l), F_X^{-1}(u)]$  with suitable quantiles  $l, u$ .

## 2.6 | Using Performance Measures in Applied Studies

Performance measures for estimated non-linear associations are not only relevant in simulation studies with a known ground truth, but could also be used as similarity measures in application studies to compare two curves estimated by two different methods. In this case, one can replace  $f, \hat{f}$  in the formulas in Tables 1–3 with  $\hat{f}^{(1)}, \hat{f}^{(2)}$ , where  $\hat{f}^{(1)}$  denotes the curve estimated by the first method and  $\hat{f}^{(2)}$  the curve estimated by the second method. The precision  $\hat{p}(x)$  then refers to the precision of the estimator of the difference  $\hat{f}^{(1)}(x) - \hat{f}^{(2)}(x)$  (Buchholz et al. 2014).

**TABLE 1** | Performance measures for localization = “range” and axis of aggregation = “Y”.

Type of aggregation	Loss	Formula
Integration over dx	Difference	$\int_S (\hat{f}(x) - f(x)) dx$
	Absolute	$\int_S  \hat{f}(x) - f(x)  dx$
	Squared	$\int_S (\hat{f}(x) - f(x))^2 dx$
	e-level accuracy	$\int_S \mathbb{1}_{\{ \hat{f}(x) - f(x)  \leq \epsilon\}} dx$
Expectation over $dF_X(x)$	Difference	$\int_S (\hat{f}(x) - f(x)) dF_X(x)$
	Absolute	$\int_S  \hat{f}(x) - f(x)  dF_X(x)$
	Squared	$\int_S (\hat{f}(x) - f(x))^2 dF_X(x)$
	e-level accuracy	$\int_S \mathbb{1}_{\{ \hat{f}(x) - f(x)  \leq \epsilon\}} dF_X(x)$
Quantile with respect to $F_X$ (e.g., median: $q = 0.5$ )	Difference	$F_{\hat{f}(X_S) - f(X_S)}^{-1}(q)$
	Absolute	$F_{ \hat{f}(X_S) - f(X_S) }^{-1}(q)$
	Squared	$F_{(\hat{f}(X_S) - f(X_S))^2}^{-1}(q)$
	e-level accuracy	$F_{\mathbb{1}_{\{ \hat{f}(X_S) - f(X_S)  \leq \epsilon\}}}^{-1}(q)$
Expectation over precision of the estimator	Difference	$\int_S (\hat{f}(x) - f(x)) \hat{p}(x) dx$
	Absolute	$\int_S  \hat{f}(x) - f(x)  \hat{p}(x) dx$
	Squared	$\int_S (\hat{f}(x) - f(x))^2 \hat{p}(x) dx$
	e-level accuracy	$\int_S \mathbb{1}_{\{ \hat{f}(x) - f(x)  \leq \epsilon\}} \hat{p}(x) dx$
Maximum	Difference	$\max_{x \in S} (\hat{f}(x) - f(x))$
	Absolute	$\max_{x \in S}  \hat{f}(x) - f(x) $
	Squared	$\max_{x \in S} (\hat{f}(x) - f(x))^2$
	e-level accuracy	$\max_{x \in S} \mathbb{1}_{\{ \hat{f}(x) - f(x)  \leq \epsilon\}}$
Minimum	Difference	$\min_{x \in S} (\hat{f}(x) - f(x))$
	Absolute	$\min_{x \in S}  \hat{f}(x) - f(x) $
	Squared	$\min_{x \in S} (\hat{f}(x) - f(x))^2$
	e-level accuracy	$\min_{x \in S} \mathbb{1}_{\{ \hat{f}(x) - f(x)  \leq \epsilon\}}$

Note: The symbol  $S$  is a placeholder either for  $S = \mathcal{X}$  or  $S = [F_X^{-1}(l), F_X^{-1}(u)]$ . Formulas are given for functional characteristic = “function”, analogous formulas for  $f'$ ,  $f''$  can be obtained by replacing  $f, \hat{f}$  with the respective derivatives.

**TABLE 2** | Performance measures for localization = “range” and axis of aggregation = “X”.

Type of aggregation	Loss	Formula
Number of roots	Difference	$\#\text{roots}^S(\hat{f}) - \#\text{roots}^S(f)$
	Absolute	$ \#\text{roots}^S(\hat{f}) - \#\text{roots}^S(f) $
	Squared	$(\#\text{roots}^S(\hat{f}) - \#\text{roots}^S(f))^2$
	e-level accuracy	$\mathbb{1}_{\{ \#\text{roots}^S(\hat{f}) - \#\text{roots}^S(f)  \leq \epsilon\}}$
Location of maximum	Difference	$\mathcal{X}^{\hat{f}, \max_S} - \mathcal{X}^{f, \max_S}$
	Absolute	$ \mathcal{X}^{\hat{f}, \max_S} - \mathcal{X}^{f, \max_S} $
	Squared	$(\mathcal{X}^{\hat{f}, \max_S} - \mathcal{X}^{f, \max_S})^2$
	e-level accuracy	$\mathbb{1}_{\{ \mathcal{X}^{\hat{f}, \max_S} - \mathcal{X}^{f, \max_S}  \leq \epsilon\}}$
Location of minimum	Difference	$\mathcal{X}^{\hat{f}, \min_S} - \mathcal{X}^{f, \min_S}$
	Absolute	$ \mathcal{X}^{\hat{f}, \min_S} - \mathcal{X}^{f, \min_S} $
	Squared	$(\mathcal{X}^{\hat{f}, \min_S} - \mathcal{X}^{f, \min_S})^2$
	e-level accuracy	$\mathbb{1}_{\{ \mathcal{X}^{\hat{f}, \min_S} - \mathcal{X}^{f, \min_S}  \leq \epsilon\}}$

Note: The symbol  $S$  is a placeholder either for  $S = \mathcal{X}$  or  $S = [F_X^{-1}(l), F_X^{-1}(u)]$ . Formulas are given for functional characteristic = “function”, analogous formulas for  $f'$ ,  $f''$  can be obtained by replacing  $f, \hat{f}$  with the respective derivatives.

**TABLE 3** | Performance measures for localization = “point” (evaluation at a specific  $x^* \in \mathcal{X}$ ).

	Loss
Difference	$\hat{f}(x^*) - f(x^*)$
Absolute	$ \hat{f}(x^*) - f(x^*) $
Squared	$(\hat{f}(x^*) - f(x^*))^2$
e-level accuracy	$\mathbb{1}_{\{ \hat{f}(x^*) - f(x^*)  \leq \epsilon\}}$

Note: Formulas are given for functional characteristic = “function”, analogous formulas for  $f'$ ,  $f''$  can be obtained by replacing  $f, \hat{f}$  with the respective derivatives.

Note, however, that measures with difference loss are not symmetric with respect to the two estimated curves. Mostly, these measures can be symmetrized by taking the absolute value of the measure, with the exception of measures with localization = “range” and axis of aggregation = “Y” where the type of aggregation is either the maximum or minimum. In this case, both the maximum and minimum difference loss should be considered.

## 2.7 | Calculating the Measures

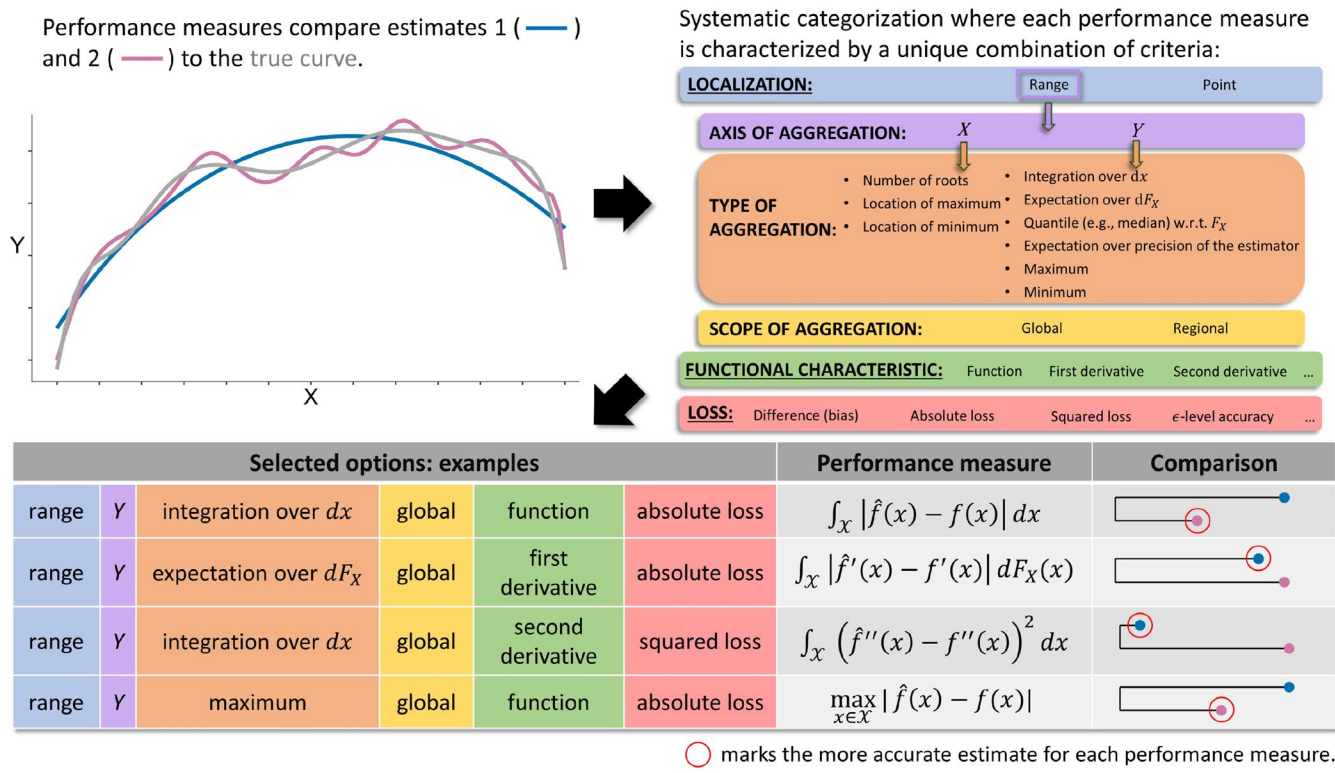
When calculating the measures with statistical software, some computational considerations have to be taken into account. Derivatives can sometimes be determined analytically, for example, for (fractional) polynomials. For splines, the choice between calculating derivatives analytically or numerically depends on the specific type of spline and the software implementation. Point-specific measures are generally easier to compute than measures with localization = “range”. For measures that involve integrals, the integrals can be approximated with Riemann sums. Note that some authors (Govindarajulu et al. 2007; Strasak et al. 2011) present formulas for this type of measure directly as Riemann sums, not as integrals. Measures with the type of aggregation = “expectation over  $dF_X$ ” can be approximated either by using the density  $\frac{d}{dx}F_X(x)$  of  $X$  in combination with an approximation by Riemann sums, or by sampling i.i.d. observations  $x$  from  $P_X$  and taking the mean over the loss values, for example, the mean over the differences  $|\hat{f}(x) - f(x)|$  in measure (2) (Binder et al. 2011).

For measures with localization = “range”, axis of aggregation = “Y” and type of aggregation = “maximum/minimum”, the maximum/minimum can be determined with a grid search. For measures with localization = “range”, axis of aggregation = “X” and type of aggregation = “location of maximum/minimum”, the locations of the maximum/minimum can also be determined with a grid search.


## 3 | Illustration of the Behavior of Different Performance Measures

To show how the categorization works in practice, Figure 4 provides an illustration. The top left panel shows an example similar to Figure 2 in the Introduction. The true curve  $f$  is shown in gray. Two estimates (estimate 1 in blue and estimate 2 in pink) are compared to the true curve. For this purpose, we can select various performance measures from the categorization (top right panel). As described before, each measure results from selecting specific options for each of the aspects. The table in the bottom panel presents four resulting examples of different performance measures. Notably, whether estimate 1 or estimate 2 is ranked higher (right-most column in the table) depends on the specific performance measure: two performance measures (rows 2 and 3) rank estimate 1 higher, while the other two favor estimate 2. This illustrates that different performance measures capture different properties of estimated curves. For example, the third performance measure penalizes the wiggleness of estimate 2 and therefore ranks estimate 1 as more accurate. In contrast, the fourth performance measure considers only the maximum deviation from the true curve, making estimate 2 the preferred estimate.

We now systematically demonstrate that different performance measures capture distinct properties of estimated curves, by varying one aspect of each measure at a time—for example, modifying the functional characteristic from “function” to “first derivative” or “second derivative” while holding all other components of the measure constant. For this systematic illustration, we switch to other examples and measures. We use several examples of ground truths and estimates that are available in



**FIGURE 4** | Illustration of using the categorization in practice. Top left panel: Example with a true curve and two estimates. Top right panel: For evaluating the accuracy of the estimates, various performance measures from the systematic categorization can be used. Bottom panel: Four examples for selecting options from the categorization, resulting in four different performance measures. The ranking of the two estimates depends on the concrete performance measure.



**FIGURE 5** | Density of the Beta(2,2) distribution.

a Shiny app at <https://thullmann.shinyapps.io/presplinesshinyapp/>. The selection of ground truths was inspired by applications in medical research, considering as wide a range of archetypical forms as possible. In the app, users can select a performance measure by choosing a specific option for each aspect of the categorization. The values of the chosen performance measure are then shown for four different examples. Each example consists of a ground truth curve  $f$  and five different hypothetical estimates  $\hat{f}$ . The first estimate is always a linear fit. The other, non-linear estimates were derived from hand-drawn curves, representing various shapes where determining the optimal curve is not immediately obvious.

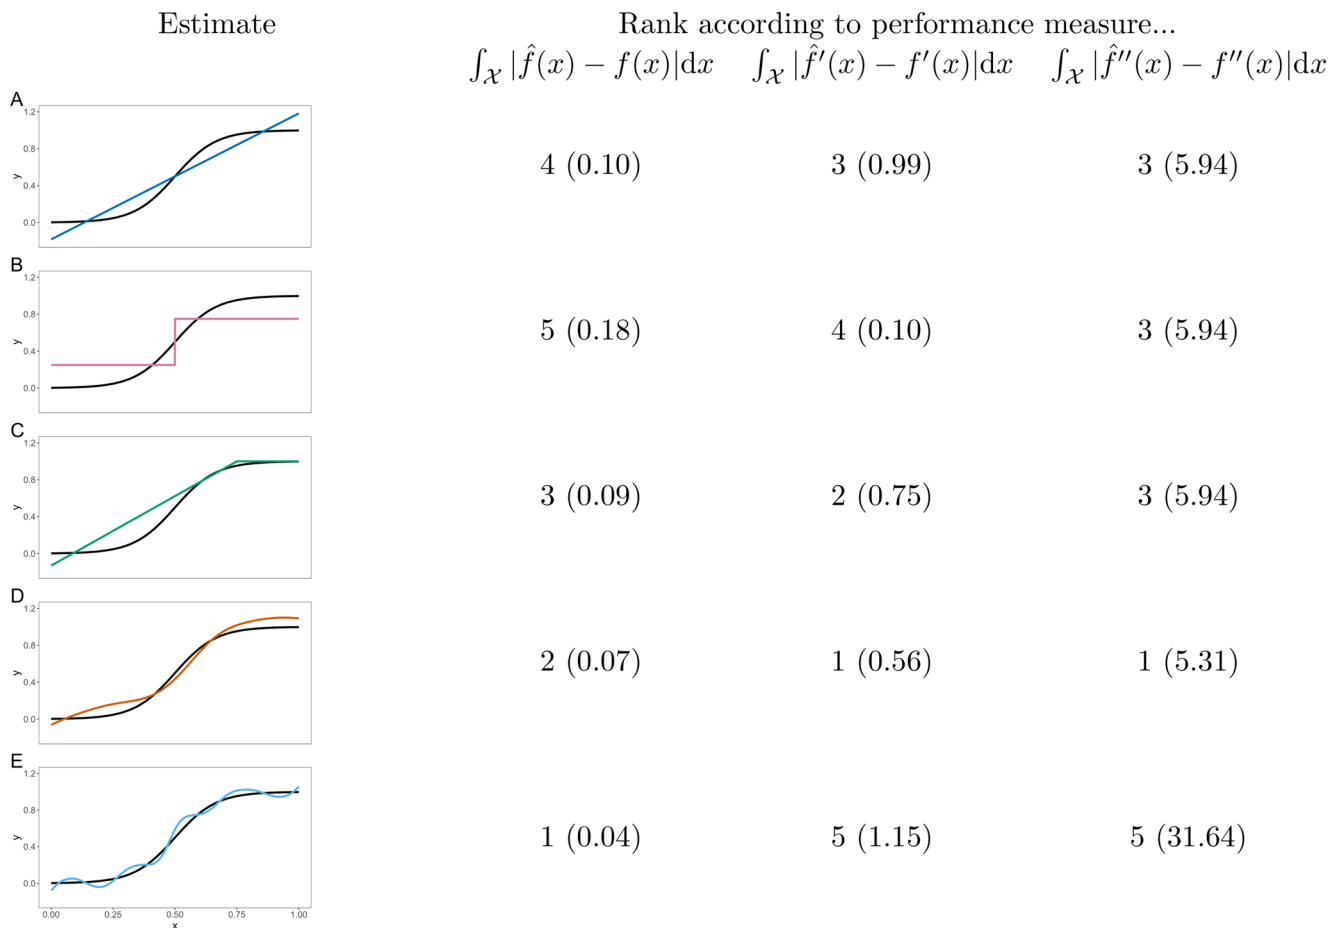
All curves and estimates are defined on the interval  $\mathcal{X} = [0, 1]$ . Three different options for the distribution  $F_X$  are available: Beta(2, 2), Beta(2, 5), and Beta(5, 2). For measures with localization = “range”, there are two options for the scope of aggregation: the full range  $\mathcal{X} = [0, 1]$  or the subrange  $[F_X^{-1}(0.05), F_X^{-1}(0.95)]$ . For measures with loss = “ $\epsilon$ -level accuracy”, different values for  $\epsilon$  can be chosen. For measures with type of aggregation = “quantile with respect to  $F_X$ ”, different quantiles can be chosen as well. Measures with type of aggregation = “expectation over precision of the estimator” are not available in the app due to the manual definition of the hypothetical estimates.

The code for the Shiny app is openly available on GitHub at <https://github.com/thullmann/performance-meas>. It includes an implementation of the performance measures from scratch in the statistical software R.

In the following, we use three examples from the app and a subset of 11 measures from the categorization in Section 2.  $F_X$  is set to Beta(2, 2). The density is shown in Figure 5.

### 3.1 | Measures Evaluating the Function, the First Derivative, or the Second Derivative

Figure 6 shows an example from the Shiny app with three measures, all with localization = “range”, axis of aggregation = “Y”, type of aggregation = “integral over  $dx$ ”, scope of aggregation =  $\mathcal{X}$ , and loss = “absolute”. The measures differ with respect to the functional characteristic (function, first derivative, and second



**FIGURE 6** | Example: Comparing measures that differ with respect to the functional characteristic (function, first derivative, or second derivative). The black curve is the true curve, the colored lines (A–E) are five hypothetical estimates. Next to each estimate, its rank among all five estimates according to the respective performance measure is shown. The values of the performance measures are shown in brackets, rounded to two decimal places.

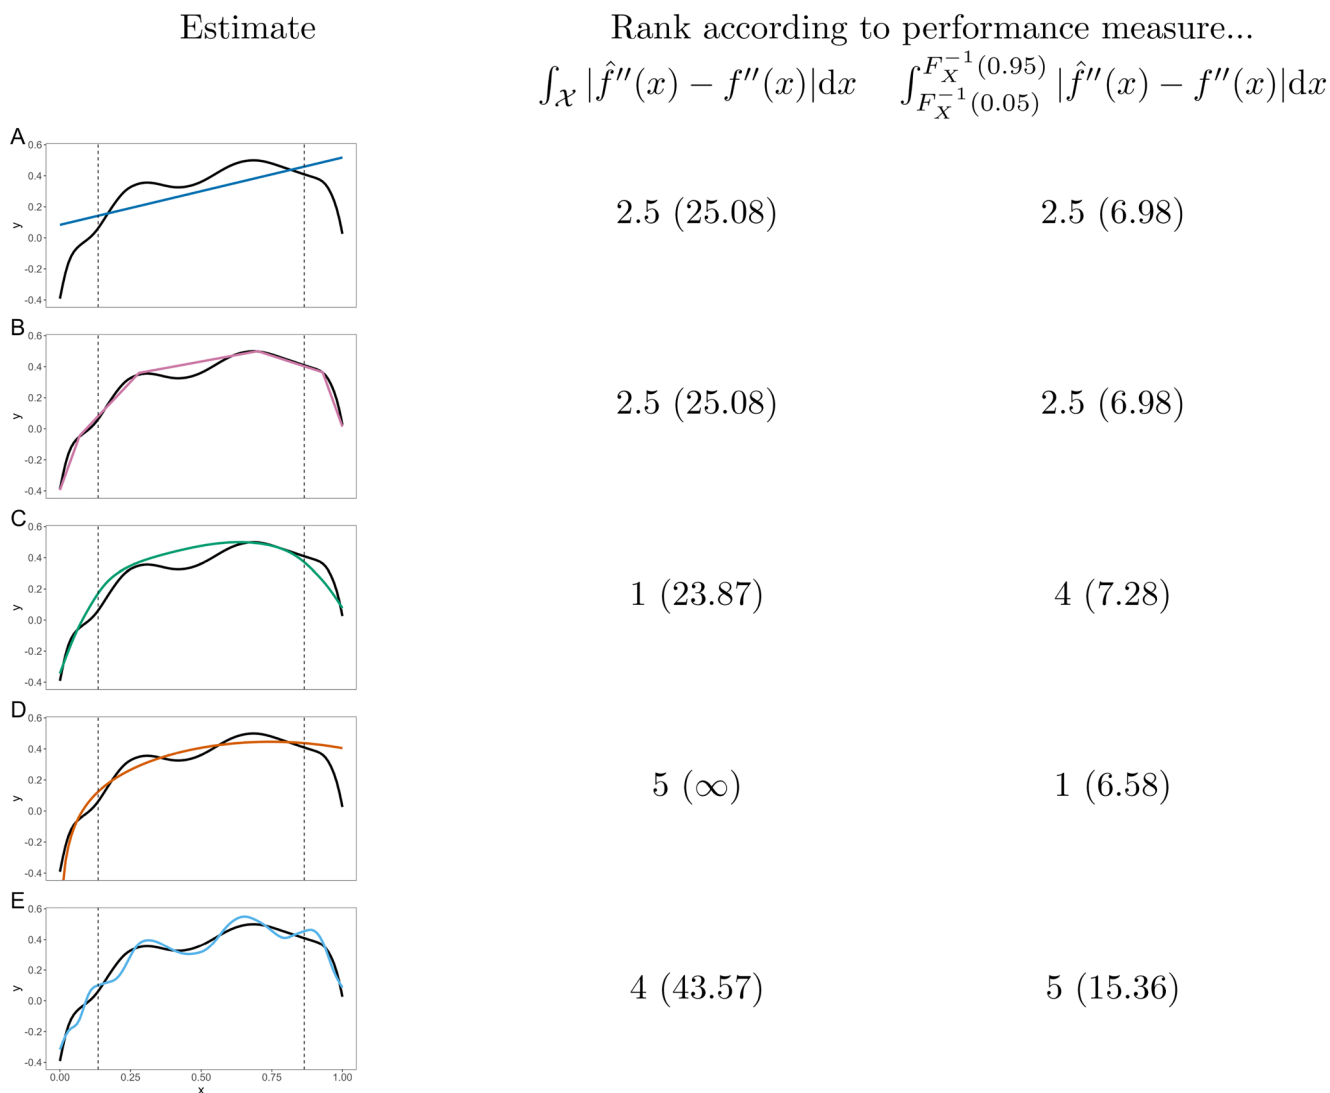
derivative). Next to each hypothetical estimate, its rank among the five estimates according to the respective performance measure is given. For the measure with functional characteristic = “function” (leftmost column in Figure 6), the most wiggly estimate E is ranked as the best one, because it is rather close to the true curve. However, when switching to the first or second derivative, estimate E is ranked as the worst estimate, because the slope and the wiggleness of the estimated curve are very different from the true curve, which is penalized by the respective measures. In the S1, screenshots of the Shiny app demonstrate how the first two measures can be selected in the app, yielding the ranks given in Figure 6.

### 3.2 | Global Versus Regional Measures

Figure 7 shows another example from the app. Two different measures are compared, both with localization = “range”, axis of aggregation = “Y”, type of aggregation = “integral over dx”, functional characteristic = “second derivative” and loss = “absolute”. The first measure uses the full range  $\mathcal{X} = [0, 1]$  as the scope of aggregation, the second measure the subrange  $[F_X^{-1}(0.05), F_X^{-1}(0.95)]$ . According to the first measure, estimate

C is the best one. Estimate D is ranked last because the integral diverges, that is, the performance value takes the value  $\infty$  (due to the second derivative of estimate D tending to  $\infty$  at a fast rate as  $x \rightarrow 0$ ). However, according to the second measure, estimate C has only rank 4 and estimate D is the best one (because now the area around zero where the second derivative tends to  $-\infty$  is excluded).

Figure 8 shows a further example where different scopes of aggregation make a difference. Two different measures are compared, both with localization = “range”, axis of aggregation = “Y”, type of aggregation = “maximum”, functional characteristic = “function” and loss = “absolute”. As before, the first measure uses the full range  $\mathcal{X} = [0, 1]$  as the scope of aggregation, the second measure the subrange  $[F_X^{-1}(0.05), F_X^{-1}(0.95)]$ , which excludes both tails of the distribution. According to the first measure, estimate B (the piecewise linear fit) is the best, while estimate D has rank 3. When restricting the scope of aggregation to the subrange  $[F_X^{-1}(0.05), F_X^{-1}(0.95)]$ , estimate B is ranked as second best and estimate D becomes the best one, because the areas towards the end points of the interval where estimate D has larger distances to the ground truth are excluded.



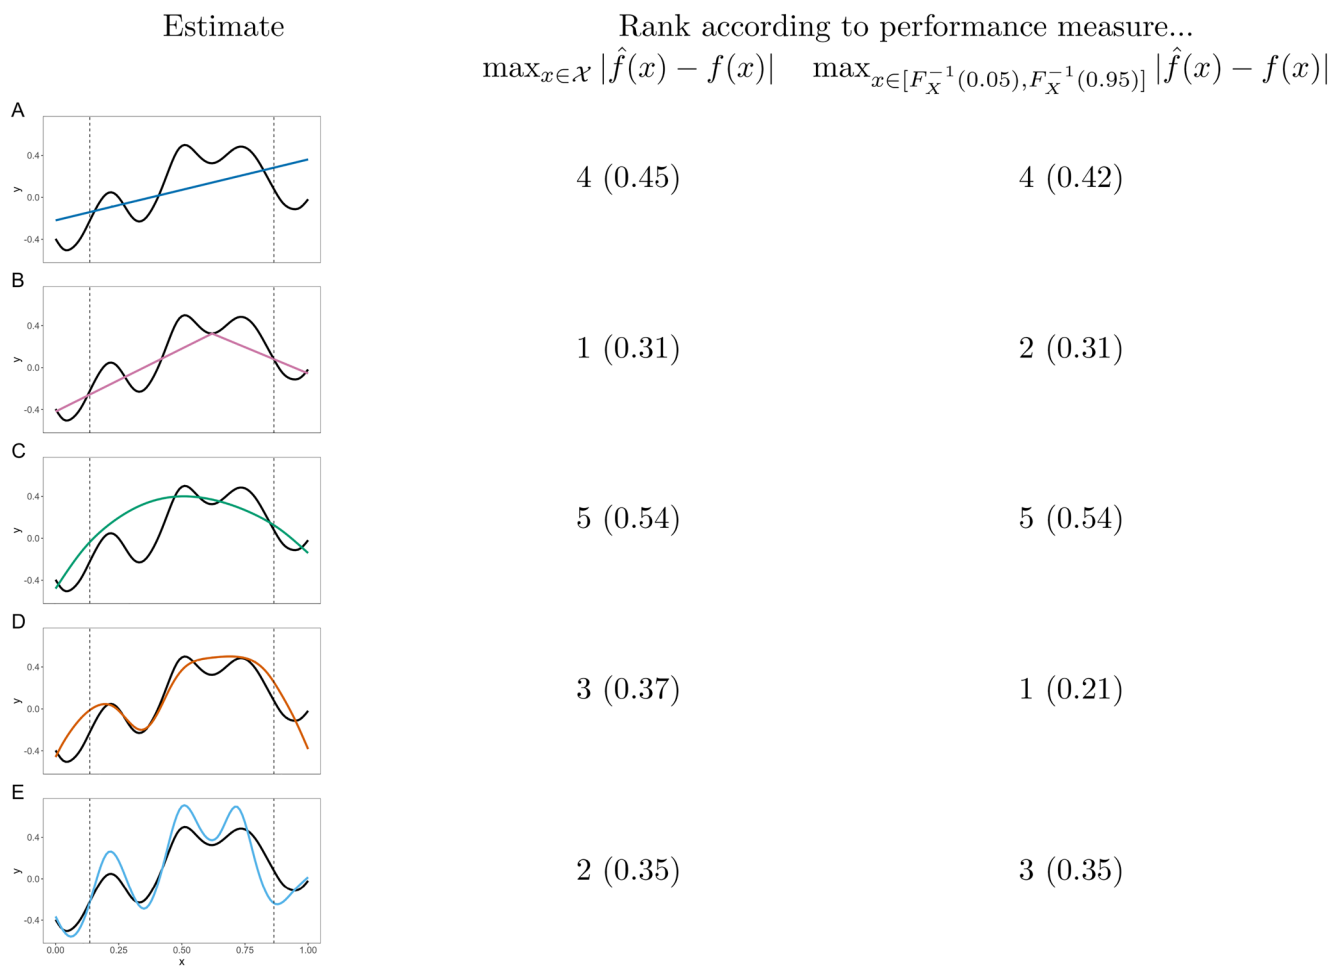
**FIGURE 7** | Example: Comparing measures where the scope of aggregation is either the whole range  $[0, 1]$  or the subrange  $[F_X^{-1}(0.05), F_X^{-1}(0.95)]$ . The quantiles  $F_X^{-1}(0.05)$  and  $F_X^{-1}(0.95)$  are indicated by dashed vertical lines. The black curve is the true curve, the colored lines (A–E) are five hypothetical estimates. Next to each estimate, its rank among all five estimates according to the respective performance measure is shown. The values of the performance measures are shown in brackets, rounded to two decimal places.

### 3.3 | Measures With Different Types of Aggregation: Integral Over $dx$ Versus the Expectation Over $dF_X$

Figure 9 compares two different measures, both with localization = “range”, axis of aggregation = “Y”, scope of aggregation =  $\mathcal{X}$ , functional characteristic = “function” and loss = “squared”. The type of aggregation is either the standard integral or the expectation over  $dF_X$ . For the standard integral, estimate B (the piecewise linear fit) is ranked better than estimate E. For the expectation over  $dF_X$ , it is the other way around. Estimates C and D also switch ranks. This is because the expectation over  $dF_X$  puts more weight on the area at the center of the interval (where the density of  $F_X$  peaks, see Figure 5). At the center, estimates B and C have more distance to the ground truth compared to estimates E and D, respectively.

### 3.4 | Measures With Different Types of Aggregation: Integral Over $dx$ Versus Maximum

Figure 10 again compares different types of loss aggregation, now the integral over  $dx$  versus the maximum. The measures have localization = “range”, axis of aggregation = “Y”, scope of aggregation =  $\mathcal{X}$ , functional characteristic = “function”, and loss = “absolute”. The first measure integrates the loss over  $\mathcal{X}$ . According to this measure, estimate B (the piecewise linear fit) has rank 3 and estimate D is the best estimate. The second measure considers the maximum loss. Here, estimate B is the best estimate and estimate D now has rank 3. This can be explained as follows: on average, estimate D is closer to the ground truth compared to estimate B, but estimate D has a relatively large difference from the ground truth at the point  $x = 1$ .



**FIGURE 8** | Example: Comparing measures where the scope of aggregation is either the whole range  $[0, 1]$  or the subrange  $[F_X^{-1}(0.05), F_X^{-1}(0.95)]$ . The quantiles  $F_X^{-1}(0.05)$  and  $F_X^{-1}(0.95)$  are indicated by dashed vertical lines. The black curve is the true curve, the colored lines (A–E) are five hypothetical estimates. Next to each estimate, its rank among all five estimates according to the respective performance measure is shown. The values of the performance measures are shown in brackets, rounded to two decimal places.

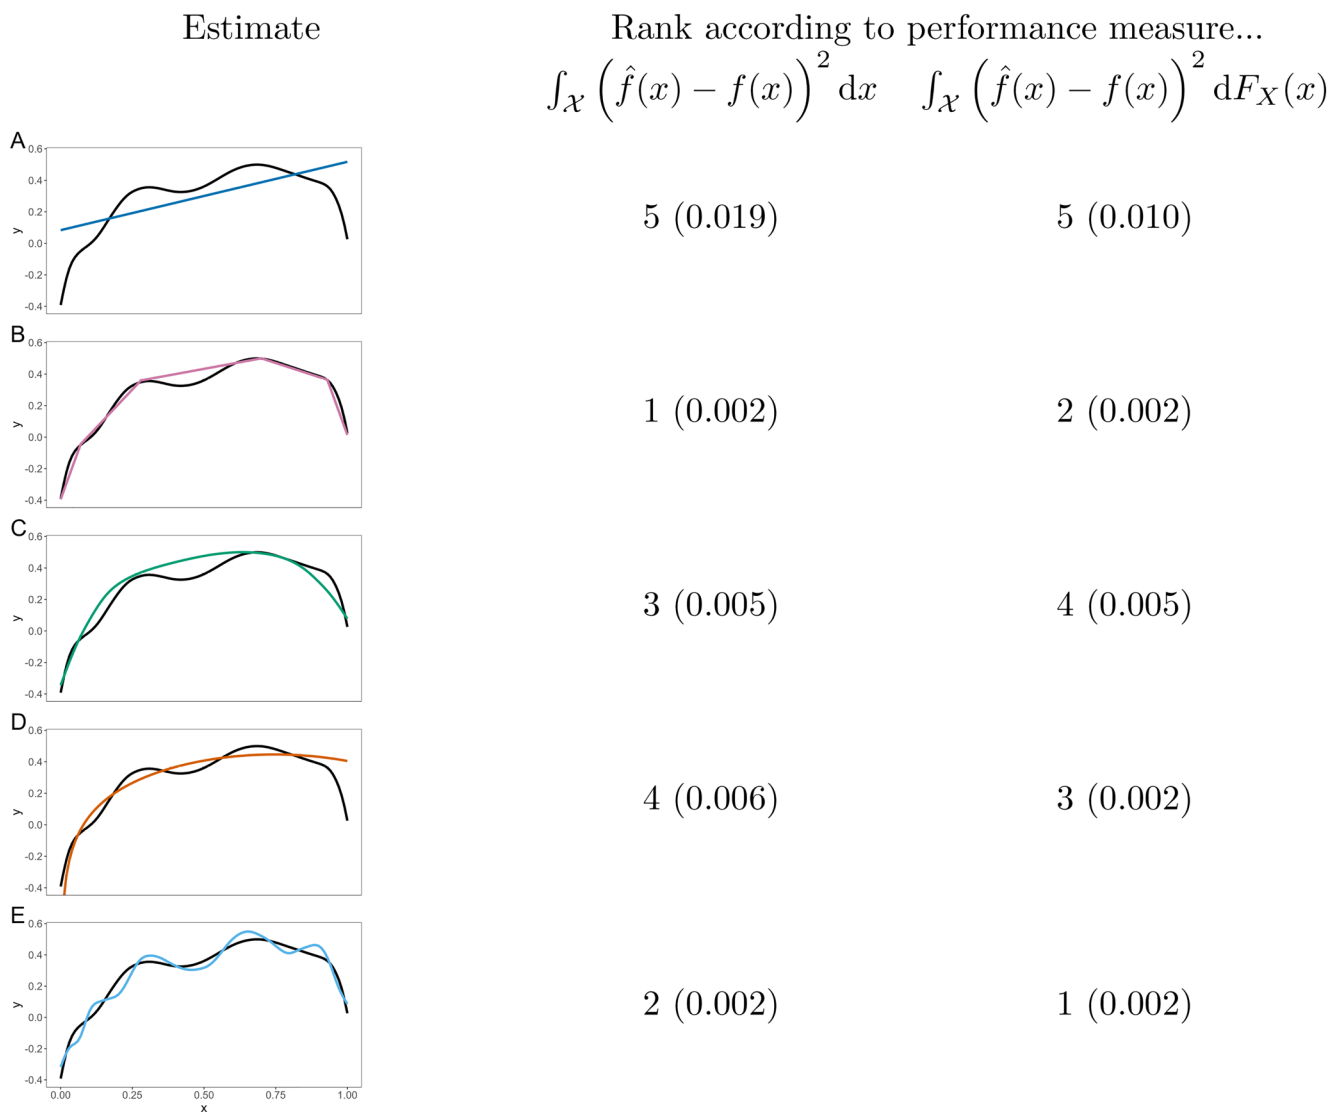
### 3.5 | Measures With Different Types of Losses

Different types of losses can lead to different rankings of estimates. Figure 11 compares the difference loss to the absolute loss. Both measures have localization = “range”, axis of aggregation = “Y”, type of aggregation = “integral over  $dx$ ”, scope of aggregation =  $\mathcal{X}$ , and functional characteristic = “function”. The first measure uses the difference loss, the second measure the absolute loss. With the difference loss, estimates A and B are best with an integrated bias of zero, because the positive and negative signs of the differences from the ground truth cancel each other out. However, according to the measure with the absolute loss, these estimates are the worst ones, because they are rather different from the ground truth in absolute terms.

Figure 12 shows an example for the absolute loss versus the squared loss. The two measures have localization = “range”, axis of aggregation = “Y”, type of aggregation = “expectation over  $dF_X$ ”, scope of aggregation =  $\mathcal{X}$ , and functional characteristic = “first derivative”. The first measure uses the absolute loss, the second measure the squared loss. With respect to the absolute loss, estimate D is ranked third, but only fifth with respect to the squared loss. The derivative of estimate D tends

to  $\infty$  for  $x \rightarrow 0$ , but at such a rate that the expectation over  $dF_X$  still converges when using the absolute loss. For the squared loss, the integral diverges to  $\infty$ , making estimate D the worst estimate.

In summary, the examples in Figures 6–12 showed that different performance measures often rank estimates in different ways. Different measures emphasize different aspects of the estimated curves. This is relevant when choosing performance measures for a simulation study aiming to compare different methods for modeling functional associations. Using a single performance measure might not capture the full complexity of the results. Existing comparison studies have often focused on only a few measures or even just one measure. Consider, for example, the study of Strasak et al. (2011) who compared fractional polynomials with spline-based estimators. The authors focused on the global performance measure which integrates the squared loss over  $dx$ , that is,  $\int_{\mathcal{X}} (\hat{f}(x) - f(x))^2 dx$ . As Figure 9 demonstrates, this can yield a different ranking of estimated curves compared to using the measure which integrates over  $dF_X$ , that is,  $\int_{\mathcal{X}} (\hat{f}(x) - f(x))^2 dF_X(x)$ . The latter measure, in turn, can yield a different ranking than the measure using



**FIGURE 9** | Example: Comparing measures that use the integral over  $dx$  versus the expectation over  $dF_X$  for aggregating the loss. The black curve is the true curve, the colored lines (A–E) are five hypothetical estimates. Next to each estimate, its rank among all five estimates according to the respective performance measure is shown. The values of the performance measures are shown in brackets, rounded to three decimal places.

the absolute loss instead of the squared loss, see Figure 12. Moreover, analogously to Figure 7, the examples in the Shiny app also show that the ranking of the curves can change when switching from the measure  $\int_{\mathcal{X}} (\hat{f}(x) - f(x))^2 dx$  to the measures  $\int_{\mathcal{X}} (\hat{f}'(x) - f'(x))^2 dx$  and  $\int_{\mathcal{X}} (\hat{f}''(x) - f''(x))^2 dx$ , with wigglier curves being more penalized by the measures considering the derivatives. Of course, these considerations do not invalidate the results of Strasak et al. (2011) (in any case, the curves in their study are different from the curves in our examples). Still, for future comparison studies, it is important to carefully choose a range of performance measures that capture different aspects of the estimated curves.


Using different performance measures does not always result in entirely different rankings, however. As can be seen in the Shiny app, many measures are similar to each other with respect to the resulting rankings. We will return to this point in the Discussion.

## 4 | Discussion

In the literature, there is a lack of consensus on how to evaluate the accuracy of non-linear functional form estimates in additive regression. Here, we have suggested a categorization of performance measures for estimated non-linear associations that could, for example, be used to evaluate and compare fits obtained from different types of spline regression and/or fractional polynomial regression.

### 4.1 | Performance Measures for Multivariable Models

We have considered only “univariable” performance measures, that is, the estimated curve of each predictor of interest is evaluated individually, though the curve can be adjusted for the estimated effects of other predictors in a multivariable model.



**FIGURE 10** | Example: Comparing measures with the integrated loss versus the maximum loss as the type of aggregation. The black curve is the true curve, the colored lines (A–E) are five hypothetical estimates. Next to each estimate, its rank among all five estimates according to the respective performance measure is shown. The values of the performance measures are shown in brackets, rounded to two decimal places.

This approach is typical for descriptive modeling, where one is interested in correctly capturing the functional form of a predictor of interest. Although it could in principle be possible to aggregate individual performance values across predictors to obtain a single performance metric for the whole model, this only makes sense if the ranges of the performance measures are comparable across predictors with different distributions. Also, after aggregation, the quality of the individual curves cannot be assessed anymore. In contrast, calculating a single metric for the whole model is more relevant for predictive modeling where the prediction error is such a metric.

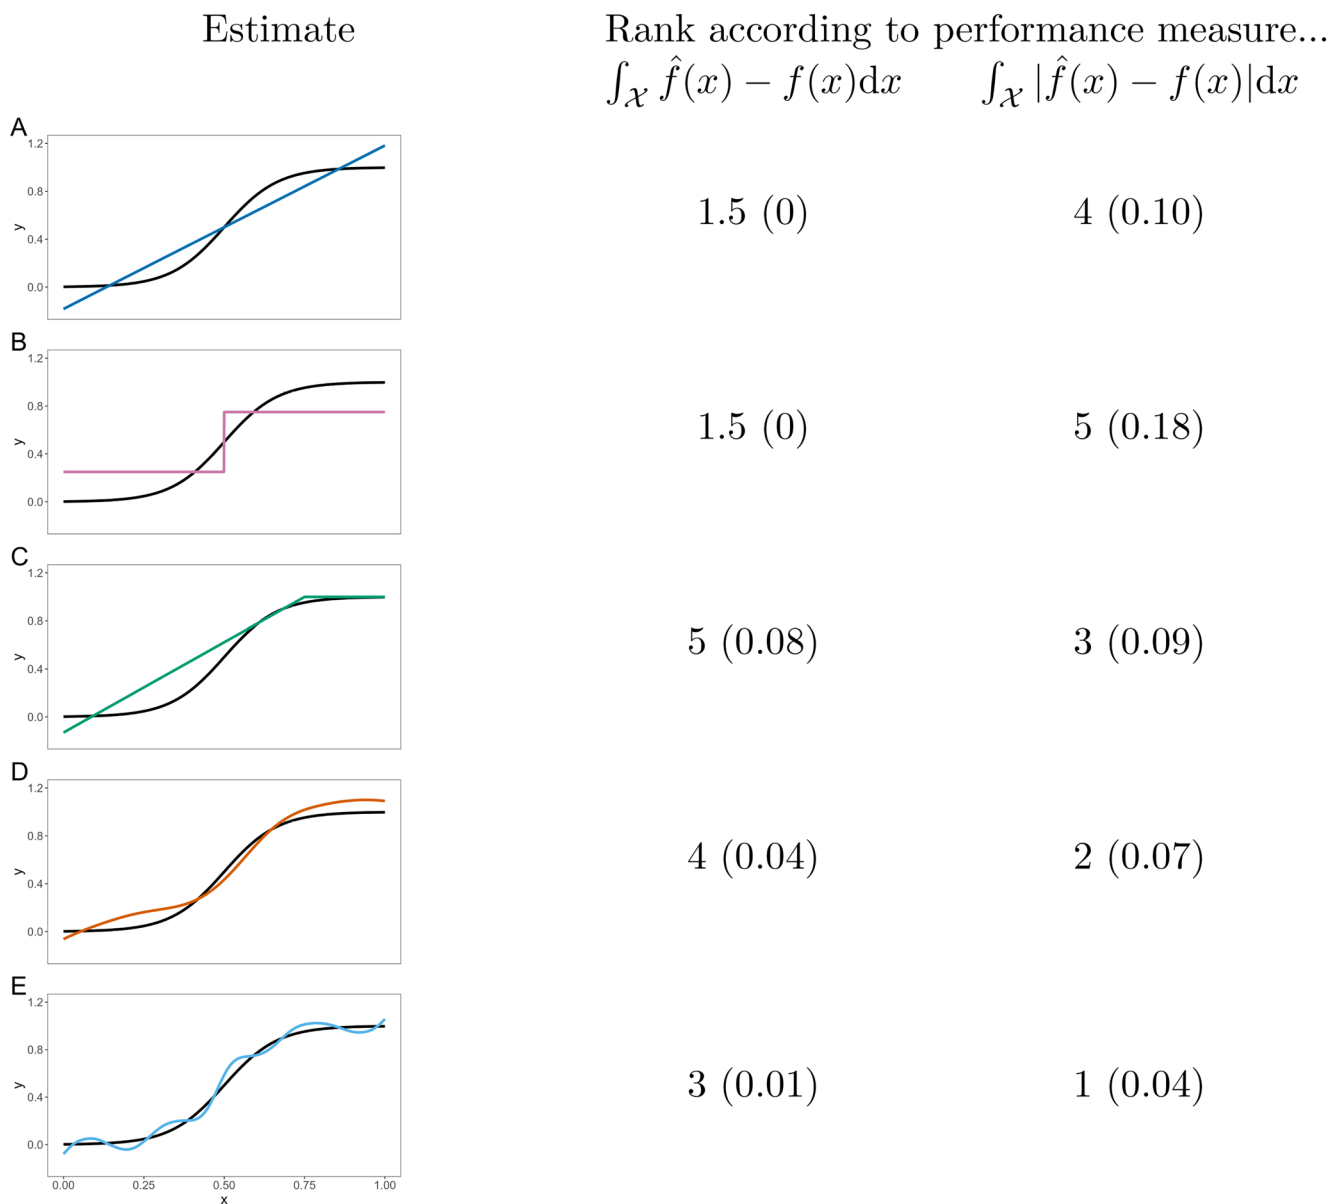
The univariable performance measures we consider do not directly cover multivariable non-linear effects, for example, when the effect of two variables  $X_1, X_2$  is given in a non-linear form

$f(X_1, X_2)$  that does not boil down to the transformation of a simple product variable  $V$ . However, our categorization of measures could be generalized to cover such cases. For example, the univariable measure

$$\int_{\mathcal{X}} |\hat{f}(x) - f(x)| dF_{\mathcal{X}}(x)$$

can be generalized to the following bivariable measure for two variables  $X_1, X_2$ :

$$\int_{\mathcal{X}_1 \times \mathcal{X}_2} |\hat{f}(x_1, x_2) - f(x_1, x_2)| dF_{(X_1, X_2)}(x)$$




**FIGURE 11** | Example: Comparing measures with difference loss versus absolute loss. The black curve is the true curve, the colored lines (a–e) are five hypothetical estimates. Next to each estimate, its rank among all five estimates according to the respective performance measure is shown. The values of the performance measures are shown in brackets, rounded to two decimal places.

with  $\mathcal{X}_1 \times \mathcal{X}_2 \subseteq \mathbb{R}^2$  denoting the space where  $(X_1, X_2)$  take values and  $F_{(X_1, X_2)}$  denoting the joint distribution of  $(X_1, X_2)$ . This integral can then be approximated with Riemann sums by partitioning  $\mathcal{X}_1 \times \mathcal{X}_2$  into sub-rectangles.

## 4.2 | Performance Measures for Predictive and Explanatory Modeling

As mentioned in the Introduction, we have considered performance measures in the context of descriptive modeling. Estimating non-linear functional forms can also be relevant for simulation studies on predictive or explanatory modeling, but in these contexts, the evaluation of the estimates might require somewhat different performance measures. In predictive modeling, one is typically less interested in what the functional

form of a predictor looks like exactly, only in whether modeling the functional form in a suitable (potentially non-linear) way can improve prediction. The prediction error is then the most relevant performance measure. In explanatory modeling, where one aims to estimate the causal effect of an explanatory variable, we can distinguish between two cases: whether the explanatory variable is binary or continuous. If the explanatory variable is binary, non-linear functional forms for continuous confounders could be estimated in order to circumvent misspecification of the confounder-outcome relation and to avoid residual confounding (Benedetti and Abrahamowicz 2004; Brenner and Blettner 1997). A too complex functional form may lead to increased variance of the intervention effect, while underfitting the confounder-outcome relation may lead to bias in the effect of the explanatory variable of interest (but probably a lower variance). Hence, the accuracy of the estimated effect



**FIGURE 12** | Example: Comparing measures with absolute loss versus squared loss. The black curve is the true curve, the colored lines (A–E) are five hypothetical estimates. Next to each estimate, its rank among all five estimates according to the respective performance measure is shown. The values of the performance measures are shown in brackets, rounded to two decimal places.

measured, for example, with the root mean squared error of the respective regression coefficient, is a relevant performance measure. If the explanatory variable is continuous (e.g., when estimating a dose–response relationship between an exposure and an outcome), non-linear functional forms could again be estimated for continuous confounders, but in addition, the dose–response relation for the exposure could also be modeled non-linearly (Desquilbet and Mariotti 2010). In that case, the performance measures in the present paper could be used to compare the estimated dose–response curve to the true curve. Exploring performance measures in the context of predictive and explanatory modeling in more detail goes beyond the scope of the present article but could be addressed in future research.

### 4.3 | Choosing a Suitable Subset of Performance Measures

Our categorization of performance measures includes a large number of measures. If we consider one option for the type of

aggregation = “quantile with respect to  $F_X$ ” (e.g., the median) and two options for the scope of aggregation ( $\mathcal{X}$  and  $[F_X^{-1}(l), F_X^{-1}(u)]$  for some specific quantiles  $l, u$ , for example  $l = 5\%$ ,  $u = 95\%$ ), there are 216 measures with localization = “range”. Additionally, for each point  $x^* \in \mathcal{X}$ , there are 12 measures with localization = “point”. Obviously, not all of these performance measures should be included in a simulation study which compares different methods for estimating non-linear associations. At least, not all measures can be reported in the main manuscript. The question then is, how researchers can choose a smaller, sufficient set of performance measures for a simulation study. A pre-selection of measures might possibly be based on the simulation settings: the “ground truth” curves might have certain properties that make some measures more relevant than others. Yet, the set of possible measures might still be rather large after this pre-selection. To further reduce the set of measures, one could look for groups of measures that give similar results, that is, that attribute similar relative performance to the methods. Such groups could be detected, for example, with a clustering algorithm. Then it would suffice to report only one measure per

group. Such a procedure would be somewhat more objective than a subjective choice of measures, in the sense that researchers who perform the same cluster analysis would obtain the same measures. We plan to explore this idea in future research.

### Author Contributions

**Theresa Ullmann:** conceptualization (equal), methodology (equal), software (lead), visualization (lead), writing – original draft (lead), writing – review and editing (equal). **Georg Heinze:** conceptualization (equal), methodology (equal), supervision (equal), writing – original draft (supporting), writing – review and editing (equal). **Michal Abrahamowicz:** writing – review and editing (equal). **Aris Perperoglou:** writing – review and editing (equal). **Willi Sauerbrei:** funding acquisition (supporting), writing – review and editing (equal). **Matthias Schmid:** writing – review and editing (equal). **Daniela Dunkler:** conceptualization (equal), funding acquisition (lead), methodology (equal), supervision (equal), visualization (supporting), writing – original draft (supporting), writing – review and editing (equal).

### Acknowledgments

We would like to thank the members of Topic Group (TG) 2 and the Publications Panel of the STRengthening Analytical Thinking for Observational Studies (STRATOS) initiative for helpful comments. The STRATOS initiative is an international consortium of researchers with a background in biostatistical and epidemiological methods dedicated to offering guidance on designing and analyzing observational studies for both specialist and non-specialist audiences. N/A

At the time of submission, STRATOS TG2 consisted of the following members (in alphabetical order): Michal Abrahamowicz, Harald Binder, Daniela Dunkler, Frank Harrell, Georg Heinze, Marc Henrion, Michael Kammer, Mariana Nold, Aris Perperoglou, Willi Sauerbrei, and Matthias Schmid. The group is co-chaired by Georg Heinze, Aris Perperoglou, and Willi Sauerbrei. Theresa Ullmann is an early-career adjunct member of STRATOS. Michal Abrahamowicz, Harald Binder, Daniela Dunkler, Frank Harrell, Georg Heinze, Marc Henrion, Michael Kammer, Mariana Nold, Aris Perperoglou, Willi Sauerbrei, Matthias Schmid.

### Conflicts of Interest

The authors declare no conflicts of interest.

### Data Availability Statement

All data and code are openly available on GitHub at <https://github.com/thullmann/performance-meas>.

### Related WIREs Articles

[Splines, knots, and penalties](#)

[Smoothing using fractional polynomials: an alternative to polynomials and splines in applied research](#)

[A survey of smoothing techniques based on a backfitting algorithm in estimation of semiparametric additive models](#)

[B-splines](#)

### References

Benedetti, A., and M. Abrahamowicz. 2004. “Using Generalized Additive Models to Reduce Residual Confounding.” *Statistics in Medicine* 23, no. 24: 3781–3801.

Binder, H., W. Sauerbrei, and P. Royston. 2011. *Multivariable Model-Building With Continuous Covariates: 1. Performance Measures and Simulation Design (Tech. Rep.)*. Universität Freiburg.

Binder, H., W. Sauerbrei, and P. Royston. 2013. “Comparison Between Splines and Fractional Polynomials for Multivariable Model Building With Continuous Covariates: A Simulation Study With Continuous Response.” *Statistics in Medicine* 32, no. 13: 2262–2277.

Boulesteix, A.-L., M. Baillie, D. Edelmann, L. Held, T. P. Morris, and W. Sauerbrei. 2024. “Editorial for the Special Collection “Towards Neutral Comparison Studies in Methodological Research.”” *Biometrical Journal* 66, no. 2: 2400031.

Boulesteix, A.-L., S. Lauer, and M. J. Eugster. 2013. “A Plea for Neutral Comparison Studies in Computational Sciences.” *PLoS One* 8, no. 4: e61562.

Boulesteix, A.-L., R. Wilson, and A. Hapfelmeier. 2017. “Towards Evidence-Based Computational Statistics: Lessons From Clinical Research on the Role and Design of Real-Data Benchmark Studies.” *BMC Medical Research Methodology* 17: 138.

Brenner, H., and M. Blettner. 1997. “Controlling for Continuous Confounders in Epidemiologic Research.” *Epidemiology* 8, no. 4: 429–434.

Buchholz, A., W. Sauerbrei, and P. Royston. 2014. “A Measure for Assessing Functions of Time-Varying Effects in Survival Analysis.” *Open Journal of Statistics* 4, no. 11: 977–998.

Carlin, J. B., and M. Moreno-Betancur. 2025. “On the Uses and Abuses of Regression Models: A Call for Reform of Statistical Practice and Teaching.” *Statistics in Medicine* 44, no. 13–14: e10244.

CDC & NCHS. 2024. *National Health and Nutrition Examination Survey Data*. U.S. Department of Health, Human Services, Centers for Disease Control, and Prevention. <https://www.cdc.gov/nchs/nhanes/>.

De Boor, C. 1978. *A Practical Guide to Splines*. Springer.

Desquilbet, L., and F. Mariotti. 2010. “Dose-Response Analyses Using Restricted Cubic Spline Functions in Public Health Research.” *Statistics in Medicine* 29, no. 9: 1037–1057.

Eilers, P. H., and B. D. Marx. 1996. “Flexible Smoothing With B-Splines and Penalties.” *Statistical Science* 11, no. 2: 89–121.

Govindarajulu, U. S., D. Spiegelman, S. W. Thurston, B. Ganguli, and E. A. Eisen. 2007. “Comparing Smoothing Techniques in Cox Models for Exposure–Response Relationships.” *Statistics in Medicine* 26, no. 20: 3735–3752.

Green, P. J., and B. W. Silverman. 1993. *Nonparametric Regression and Generalized Linear Models: A Roughness Penalty Approach*. Hall/CRC.

Harrell, F. E. 2015. *Regression Modeling Strategies: With Applications to Linear Models, Logistic and Ordinal Regression, and Survival Analysis*. Springer International Publishing.

Hirji, K. F., A. A. Tsiatis, and C. R. Mehta. 1989. “Median Unbiased Estimation for Binary Data.” *American Statistician* 43, no. 1: 7–11.

Perperoglou, A., W. Sauerbrei, M. Abrahamowicz, and M. Schmid. 2019. “A Review of Spline Function Procedures in R.” *BMC Medical Research Methodology* 19: 46.

Royston, P., and W. Sauerbrei. 2008. *Multivariable Model-Building: A Pragmatic Approach to Regression Analysis Based on Fractional Polynomials for Modelling Continuous Variables*. John Wiley & Sons.

Royston, P., and W. Sauerbrei. 2009. “Bootstrap Assessment of the Stability of Multivariable Models.” *Stata Journal* 9, no. 4: 547–570.

Shmueli, G. 2010. “To Explain or to Predict?” *Statistical Science* 25, no. 3: 289–310.

Stone, C. J., and C.-Y. Koo. 1985. “Additive Splines in Statistics.” In *Proceedings of the Statistical Computing Section ASA*, 45–48. American Statistical Association.

Strasak, A. M., N. Umlauf, R. M. Pfeiffer, and S. Lang. 2011. “Comparing Penalized Splines and Fractional Polynomials for Flexible Modelling of

the Effects of Continuous Predictor Variables.” *Computational Statistics & Data Analysis* 55, no. 4: 1540–1551.

Wood, S. 2003. “Thin Plate Regression Splines.” *Journal of the Royal Statistical Society. Series B, Statistical Methodology* 65, no. 1: 95–114.

Wood, S. 2017. *Generalized Additive Models: An Introduction With R*. 2nd ed. Hall/CRC.

### Supporting Information

Additional supporting information can be found online in the Supporting Information section. **Data S1:** Supporting Information.