



City Research Online

City, University of London Institutional Repository

Citation: Merz, M., Richman, R., Tsanakas, A. ORCID: 0000-0003-4552-5532 and Wüthrich, M. V. (2021). Interpreting Deep Learning Models with Marginal Attribution by Conditioning on Quantiles. .

This is the preprint version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/25819/>

Link to published version:

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Interpreting Deep Learning Models with Marginal Attribution by Conditioning on Quantiles

Michael Merz* Ronald Richman†‡ Andreas Tsanakas§ Mario V. Wüthrich¶

Version of March 22, 2021

Abstract

A vastly growing literature on explaining deep learning models has emerged. This paper contributes to that literature by introducing a global gradient-based model-agnostic method, which we call Marginal Attribution by Conditioning on Quantiles (MACQ). Our approach is based on analyzing the marginal attribution of predictions (outputs) to individual features (inputs). Specifically, we consider variable importance by fixing (global) output levels and, thus, explain how features marginally contribute across different regions of the prediction space. Hence, MACQ can be seen as a marginal attribution counterpart to approaches such as accumulated local effects (ALE), which study the sensitivities of outputs by perturbing inputs. Furthermore, MACQ allows us to separate marginal attribution of individual features from interaction effect, and visually illustrate the 3-way relationship between marginal attribution, output level, and feature value.

Keywords. explainable AI (XAI), model-agnostic tools, deep learning, attribution, accumulated local effects (ALE), partial dependence plot (PDP), locally interpretable model-agnostic explanation (LIME), variable importance, post-hoc analysis.

1 Introduction

Deep learning models are typically trained to provide an optimal predictive performance. Interpreting and explaining the results of deep learning models has, until recently, only played a subordinate role. With growing complexity of deep learning models, the need and requirement of being able to explain deep learning solutions has become increasingly important. This applies to almost all fields of their applications: deep learning findings in medical fields and health care need to make sense to patients, loan and mortgage evaluations and credit approvals need to be understandable to customers, insurance pricing must be explained to insurance policyholders, business processes and decisions need to be transparent to regulators, autonomous robotic tools need to comply with safety standards according to admission offices and governments, etc. These needs are even reinforced by the requirements of being able to prove that deep learning

*University of Hamburg, Faculty of Business Administration, michael.merz@uni-hamburg.de

†QED Actuaries & Consultants, ronald.richman@qedact.com

‡University of the Witwatersrand, Johannesburg

§The Business School (formerly Cass), City, University of London, A.Tsanakas.1@city.ac.uk

¶RiskLab, Department of Mathematics, ETH Zurich, mario.wuethrich@math.ethz.ch

solutions do not discriminate w.r.t. protected features and are in line with data protection regulation. Thus, there is substantial social and political pressure to be able to explain, illustrate and verify deep learning solutions, in order to provide reassurance that these work properly. Recent research focuses on different methods of explaining deep learning decision making; an overview is given in [Samek and Müller 2019]. Some of these methods provide a post-hoc analysis which aims at understanding global model behavior, explaining individual outcomes and learned representations. Often this is done by explaining representative examples. We are going to discuss some of these post-hoc analysis methods in the literature overview presented in the next section. Other methods aim at a wider interdisciplinary approach by more broadly examining how decision making is done in a social context, see e.g. [Miller 2019]. All these approaches have in common that they try to “open up the black-box” to make decision making explainable to stakeholders.

Our paper contributes to this literature. Our main contribution is that we provide a novel gradient-based model-agnostic tool that is motivated by analyzing marginal contributions to deep learning decisions in the spirit of salience methods, as described in [Ancona et al. 2019]. Salience methods are local model-agnostic tools that attribute marginal effects on outputs to different inputs. Motivated by sensitivity analysis tools in risk measurement, we aggregate local marginal attributions to obtain a global picture at a given quantile level of the output variable. We call this method Marginal Attribution by Conditioning on Quantiles (MACQ). It describes a global variable importance measure that varies with the output level. The aggregation of local marginal effects is justified by the fact that this aggregation can be seen as a directional derivative of a distortion risk measure, see [Hong 2009] and Proposition 1 in [Tsanakas and Millossovich 2015]. As second contribution, we extend this view by including higher order derivatives beyond linear marginal contributions. This additional step can be seen in the context of deep Taylor decompositions (DTD), similar to [Montavon et al. 2017]. A difficulty in Taylor decompositions is that they depend on a reference point. By rearranging the terms and by taking advantage of our quantile view, we determine an optimal global reference point that allows us to quantify both variable importance and interaction strength in our MACQ approach. The third contribution is that we provide graphic tools that provide a 3-way relationship between (i) marginal attribution, (ii) response/output level and (iii) feature value.

Organization. In the next section we give a literature overview that embeds our MACQ method into the present toolbox of model explainability. This literature overview is also used to introduce the relevant notation. In Section 3 we present our main idea of aggregating local marginal attributions to a quantile sensitivity analysis. Section 4 presents a higher order expansion which grounds a study of interaction strength. Section 5 discusses the choice of the reference point. An extended example is presented in Section 6. Finally, in Section 7 we state brief conclusions.

2 Literature overview

We give a brief summary of recent developments in post-hoc interpretability and explainability tools for deep learning models. This summary also serves to introduce the relevant notation for this paper. Assume the following regression function is smooth (in fact, we are only going to use twice differentiable in our setting)

$$\mu : \mathbb{R}^q \rightarrow \mathbb{R}, \quad \mathbf{x} \mapsto \mu(\mathbf{x}), \quad (2.1)$$

with feature $\mathbf{x} = (x_1, \dots, x_q)^\top \in \mathbb{R}^q$. This regression function is assumed to describe the systematic effects of features on the random variable Y via the (conditional) expectation

$$\mathbb{E}[Y|\mathbf{x}] = \mu(\mathbf{x}).$$

We assume smoothness of regression function (2.1) because our model-agnostic proposal will be gradient-based. In our example in Section 6, we will use a deep feed-forward neural network on tabular input data, having the hyperbolic tangent as activation function. This gives us a smooth regression function and formal derivation can be done in standard software such as TensorFlow/Keras and PyTorch.

2.1 Model-agnostic tools

Recent literature aims understanding such regression functions (2.1) coming from deep learning models. One approach is to analyze marginal plots. We select one component x_j of \mathbf{x} and study the function

$$x_j \in \mathbb{R} \mapsto \mu(x_j, \mathbf{x}_{\setminus j}),$$

where $\mathbf{x}_{\setminus j}$ denotes the remaining components of \mathbf{x} which are kept fixed. This is the method of individual conditional expectation (ICE) of [Goldstein et al. 2015]. If we have thousands or millions of instances (Y, \mathbf{x}) , it might be advantageous to study ICE profiles on an aggregated level. This is the proposal of [Friedman 2001] and [Zhao and Hastie 2021] called partial dependence plots (PDPs). We introduce the feature distribution P which describes the family of all (potential) features $\mathbf{X} \sim P$. The PDP profile of component $1 \leq j \leq q$ is defined by

$$x_j \mapsto \mathbb{E}_P [\mu(x_j, \mathbf{X}_{\setminus j})] = \int \mu(x_j, \mathbf{x}_{\setminus j}) dP(\mathbf{x}_{\setminus j}).$$

The critical point in this approach is that it does not reflect the (true) dependence structure between feature components X_j and $\mathbf{X}_{\setminus j}$, i.e., as described by feature distribution P . The method of accumulated local effects (ALEs) introduced by [Apley and Zhu 2020] aims at correctly incorporating the dependence structure in \mathbf{X} . The local effect of component x_j in individual feature \mathbf{x} is given by the partial derivative

$$\mu_j(\mathbf{x}) = \frac{\partial \mu(\mathbf{x})}{\partial x_j}. \quad (2.2)$$

The average local effect of component $1 \leq j \leq q$ is obtained by

$$x_j \mapsto \Delta_j(x_j) = \mathbb{E}_P [\mu_j(\mathbf{X}) | X_j = x_j] = \int \mu_j(x_j, \mathbf{x}_{\setminus j}) dP(\mathbf{x}_{\setminus j} | x_j), \quad (2.3)$$

where $P(\mathbf{x}_{\setminus j} | x_j)$ denotes the conditional distribution of $\mathbf{X}_{\setminus j}$, given $X_j = x_j$. ALEs integrate the average local effects $\Delta_j(\cdot)$ over their domain, thus, the ALE profile is defined by

$$x_j \mapsto \int_{x_{j_0}}^{x_j} \Delta_j(z_j) dz_j = \int_{x_{j_0}}^{x_j} \int \mu_j(z_j, \mathbf{x}_{\setminus j}) dP(\mathbf{x}_{\setminus j} | z_j) dz_j, \quad (2.4)$$

where x_{j_0} is a given initialization point. The main difference between PDPs and ALEs is that the latter correctly considers the dependence structure between X_j and $\mathbf{X}_{\setminus j}$.

Remark 2.1 • The main difference between PDPs and ALEs is that the latter correctly considers the dependence structure between X_j and $\mathbf{X}_{\setminus j}$. The two profiles coincide if X_j and $\mathbf{X}_{\setminus j}$ are independent under P .

- [Apley and Zhu 2020] provide a discretized version of the ALE profile that can also be applied to non-differentiable regression functions $\mu(\cdot)$. Basically, this can be received either by finite differences or by a local analysis in an environment of a selected feature value x_j .
- More generally the local effect (2.2) allows us to consider a 1st order Taylor expansion. Denote by $\nabla_{\mathbf{x}}\mu(\mathbf{x})$ the gradient of $\mu(\cdot)$ w.r.t. \mathbf{x} . We have

$$\mu(\mathbf{x} + \boldsymbol{\epsilon}) = \mu(\mathbf{x}) + (\nabla_{\mathbf{x}}\mu(\mathbf{x}))^\top \boldsymbol{\epsilon} + o(\|\boldsymbol{\epsilon}\|), \quad (2.5)$$

for $\boldsymbol{\epsilon} \in \mathbb{R}^q$ going to zero. This gives us a 1st order local approximation to $\mu(\cdot)$ in \mathbf{x} , which reflects the local (linear) behavior similar to the locally interpretable model-agnostic explanation (LIME) introduced by [Ribeiro et al. 2016]. That is, (2.5) fits a local linear regression model around $\mu(\mathbf{x})$ with regression parameters described by the components of the gradient $\nabla_{\mathbf{x}}\mu(\mathbf{x})$. LIME then uses regularization, e.g. LASSO, to select the most relevant feature components in the neighborhood of $\mu(\mathbf{x})$.

- More generally, (2.5) defines a local surrogate model that can be used for a local sensitivity analysis by perturbing \mathbf{x} within a small environment. White-box surrogate models are popular tools to explain complex regression functions, for instance, decision trees can be fit to network regression models for extracting the most relevant feature information.

2.2 Gradient based model-agnostic tools

Gradient-based model-agnostic tools can be used to attribute outputs to (feature) inputs. Attribution denotes the process of assigning a relevance index to input components, in order to explain a certain output, see [Efron 2020]. [Ancona et al. 2019] provide a nice overview of gradient-based attribution methods. In formula (2.2) of the previous subsection we have met a first attribution method which gives the sensitivity of the output $\mu(\mathbf{x})$ as a function of the input \mathbf{x} . The marginal attribution we are going to present considers the contribution to a given output in the spirit of salience methods.

Marginal attribution is obtained by considering the directional derivative w.r.t. the features

$$x_j \mapsto x_j \mu_j(\mathbf{x}) = x_j \frac{\partial \mu(\mathbf{x})}{\partial x_j}. \quad (2.6)$$

This has first been discussed in the machine learning community by [Shrikumar et al. 2016] who observed that this can make attribution more concise; these directional derivatives have been coined Gradient*Input in the machine learning literature, see Ancona [Ancona et al. 2019]. Mathematically speaking, these marginal attributions can be understood as individual contributions to a certain value in a Taylor series sense (and relative to a reference point). Having a linear regression model $\mathbf{x} \mapsto \beta_0 + \sum_{j=1}^q \beta_j x_j$, the marginal attributions give an additive decomposition of the regression function, and β_j can be considered as the relevance index of component j . In non-linear regression models, such a linear decomposition only holds true very

locally, see (2.5), and other methods such as the Shapley value [Shapley 1953] are used to quantify non-linear effects and interaction effects, see [Lundberg and Lee 2017]. We also mention [Sundararajan et al. 2017], who consider integrated gradients

$$x_j \mapsto x_j \int_0^1 \mu_j(\mathbf{x}_0 + z(\mathbf{x} - \mathbf{x}_0)) dz, \quad (2.7)$$

for a given reference point \mathbf{x}_0 . This mitigates the problem of only being accurate locally. In practice, however, this is computationally demanding, similarly to the Shapley value.

There are other methods that are specific to deep networks. We mention layer-wise propagation (LRP) by [Binder et al. 2016] and DeepLIFT (Deep Learning Important FeaTures) by [Shrikumar et al. 2017]. These methods use a backward pass from the output to the input. In this backward pass a relevance index (budget) is locally redistributed (recursively from layer to layer), resulting in a relevance index on the inputs (for the given output). [Ancona et al. 2019] show in Propositions 1 and 2 that these two methods can be understood as an average over marginal attributions. We remark that these methods are mainly used for convolutional neural networks (CNNs), e.g., in image recognition, whereas our MACQ proposal is more suitable for tabular data because we require differentiability w.r.t. the inputs \mathbf{x} . CNNs architectures are often non-differentiable because of the use of max-pooling layers.

Our contribution builds on marginal attributions (2.6). Marginal attributions are, by definition, local explanations, and we are going to show how to integrate these local considerations into a global variable importance analysis. [Samek and Müller 2019] call such an aggregation of individual explanations a *global meta-explanation*. As a consequence, our MACQ approach is the marginal attribution counterpart to ALEs by fixing (global) output levels and describing how features marginally contribute to these levels, whereas ALEs rather study the sensitivities of the outputs by perturbing the inputs.

3 Marginal attribution by conditioning on quantiles

We consider regression model (2.1) from a marginal attribution point of view. Motivated by the risk sensitivity tools of [Hong 2009] and [Tsanakas and Millosovich 2015], we do not consider average local effects (2.3) conditioned on event $\{X_j = x_j\}$, but we would rather like to understand how feature components contribute to a certain response level $\mu(\mathbf{x})$. The former studies sensitivities of outputs $\mu(\mathbf{x})$ in inputs \mathbf{x} , whereas the latter considers marginal attribution of outputs $\mu(\mathbf{x})$ to inputs \mathbf{x} . This allows us to study how the response levels are composed in different regions of the decision space, as this is of intrinsic interest e.g. in financial applications. Select a quantile level $\alpha \in (0, 1)$, the α -quantile of $\mu(\mathbf{X})$ is given by

$$F_{\mu(\mathbf{X})}^{-1}(\alpha) = \inf \{y \in \mathbb{R}; F_{\mu(\mathbf{X})}(y) \geq \alpha\},$$

where $F_{\mu(\mathbf{X})}(y) = P[\mu(\mathbf{X}) \leq y]$ describes the distribution function of $\mu(\mathbf{X})$.

The *1st order attributions* to components $1 \leq j \leq q$ on quantile level α are defined by

$$S_j(\mu; \alpha) = \mathbb{E}_P \left[X_j \mu_j(\mathbf{X}) \mid \mu(\mathbf{X}) = F_{\mu(\mathbf{X})}^{-1}(\alpha) \right]. \quad (3.1)$$

These are the marginal attributions by conditioning on quantiles (MACQ).

[Tsanakas and Millossovich 2015] show that (3.1) naturally arises as sensitivities of distortion risk measures, and choosing the α -Dirac distortion we exactly receive (3.1), which corresponds to the sensitivities of the value-at-risk (VaR) risk measure on the given quantile level. Thus, the sensitivities of the VaR risk measure can be described by the average of the marginal attributions $X_j \mu_j(\mathbf{X})$, conditioned on being on the corresponding quantile level. The interested reader is referred to Appendix A for a more detailed description of distortion risk measures.

Alternatively, we can describe the 1st order attributions (3.1) by a 1st order Taylor expansion (2.5) in feature perturbation $\boldsymbol{\epsilon} = -\mathbf{x}$

$$\mu(\mathbf{0}) \approx \mu(\mathbf{x}) - (\nabla_{\mathbf{x}} \mu(\mathbf{x}))^\top \mathbf{x}. \quad (3.2)$$

This explains that the 1st order attributions (3.1) describe a 1st order Taylor approximation at the common *reference point* $\mathbf{0}$, and rearranging the terms we get the *1st order contributions* to a given response level

$$F_{\mu(\mathbf{X})}^{-1}(\alpha) = \mathbb{E}_P \left[\mu(\mathbf{X}) \mid \mu(\mathbf{X}) = F_{\mu(\mathbf{X})}^{-1}(\alpha) \right] \approx \mu(\mathbf{0}) + \sum_{j=1}^q S_j(\mu; \alpha). \quad (3.3)$$

Remark 3.1 • A 1st order Taylor expansion (2.5) gives a local model-agnostic description in the spirit of LIME. Explicit choice $\boldsymbol{\epsilon} = -\mathbf{x}$ provides (3.2), which can be viewed as a local description of $\mu(\mathbf{0})$ relative to \mathbf{x} . The 1st order contributions (3.3) combine all these local descriptions (3.1) w.r.t. a given quantile level to get the integrated MACQ view of $\mu(\mathbf{0})$, i.e.

$$\mu(\mathbf{0}) \approx \mathbb{E} \left[\mu(\mathbf{X}) - (\nabla_{\mathbf{x}} \mu(\mathbf{X}))^\top \mathbf{X} \mid \mu(\mathbf{X}) = F_{\mu(\mathbf{X})}^{-1}(\alpha) \right] = F_{\mu(\mathbf{X})}^{-1}(\alpha) - \sum_{j=1}^q S_j(\mu; \alpha).$$

This exactly corresponds to 1st order approximation (3.3). In the sequel it is less important that we can approximate $\mu(\mathbf{0})$ by this integrated view, but $\mu(\mathbf{0})$ plays the role of the *reference level* that calibrates our global meta-explanation. Thus, all explanations made are understood relative to this reference level $\mu(\mathbf{0})$.

- In (3.2)-(3.3) we implicitly assumed that $\mathbf{0}$ is a suitable reference point for calibrating our global meta-explanation. We further explore and improve this calibration in Section 5, below.
- Integrated gradients (2.7) integrate along a single path from a reference point \mathbf{x}_0 to \mathbf{x} to make the 1st order Taylor approximation precise. We exchange the roles of the points, here, and we approximate the reference point by aggregating over all local descriptions in features \mathbf{X} .
- 1st order contributions (3.3) provide a 3-way description of the regression function, namely, they combine (i) marginal attribution $S_j(\mu; \alpha)$ as a function of $1 \leq j \leq q$, (ii) response level $F_{\mu(\mathbf{X})}^{-1}(\alpha)$ as a function of α , and (iii) feature values x_j . In our application in Section 6 we will illustrate the data from these different angles, each having its importance in explaining the response.

- 1st order attribution (3.1) combines marginal attributions $X_j \mu_j(\mathbf{X})$ by focusing on a common quantile level. A similar approach could also be done for other model-agnostic tools, such as the Shapley value.

Example 3.2 (linear regression) A linear regression model considers regression function

$$\mathbf{x} \mapsto \mu(\mathbf{x}) = \beta_0 + \boldsymbol{\beta}^\top \mathbf{x}, \quad (3.4)$$

with bias/intercept $\beta_0 \in \mathbb{R}$ and regression parameter $\boldsymbol{\beta} \in \mathbb{R}^q$. The 1st order contributions (3.3) are for $\alpha \in (0, 1)$ given by

$$F_{\mu(\mathbf{X})}^{-1}(\alpha) = \mu(\mathbf{0}) + \sum_{j=1}^q S_j(\mu; \alpha) = \beta_0 + \sum_{j=1}^q \beta_j \mathbb{E}_P \left[X_j \mid \mu(\mathbf{X}) = F_{\mu(\mathbf{X})}^{-1}(\alpha) \right]. \quad (3.5)$$

Thus, we weight regression parameters β_j with the feature components X_j according to their contributions to quantile $F_{\mu(\mathbf{X})}^{-1}(\alpha)$; and the reference point $\mathbf{0}$ is given naturally providing initialization $\mu(\mathbf{0}) = \beta_0$.

This MACQ explanation (3.5) is rather different from the ALE profile(2.4). If we initialize $x_{j_0} = 0$ we receive ALE profile for the linear regression model

$$x_j \mapsto \int_0^{x_j} \Delta_j(z_j) dz_j = \beta_j x_j.$$

This is exactly the marginal attribution (2.6) of component j in the linear regression model and it explains the change of the linear regression function if we change feature component x_j , whereas (3.5) describes the contribution of each feature component to an expected response level $\mu(\mathbf{x})$. ■

In general, Taylor expansion (3.3) is accurate if the distance between $\mathbf{0}$ and \mathbf{X} is small enough for all relevant \mathbf{X} , and if the regression function can be well described around $\mu(\mathbf{X})$ by a linear function. The former requires that the reference point is chosen somewhere “in the middle” of the feature distribution P . The accuracy of the 1st order approximation is quantified by

$$\left| F_{\mu(\mathbf{X})}^{-1}(\alpha) - \mu(\mathbf{0}) - \sum_{j=1}^q S_j(\mu; \alpha) \right|. \quad (3.6)$$

Thus, we want (3.6) to be small uniformly in quantile level α , for the given reference point $\mathbf{0}$, as then the 1st order attributions give a good description on all quantile levels α . In the linear regression case this description is exact, see (3.5). In contrast to the Taylor decomposition in [Montavon et al. 2017], the quantiles $F_{\mu(\mathbf{X})}^{-1}(\alpha)$ give us a natural anchor point for determining a suitable reference point, which is also computationally feasible. This will be done in Section 5.

4 Interaction strength

[Friedman and Popescu 2008] and [Apley and Zhu 2020] have shown how higher order derivatives of $\mu(\cdot)$ allow us to study interaction strength in systematic effects. This requires the study of higher order Taylor expansions. The 2nd order Taylor expansion is given by

$$\mu(\mathbf{x} + \boldsymbol{\epsilon}) = \mu(\mathbf{x}) + (\nabla_{\mathbf{x}} \mu(\mathbf{x}))^\top \boldsymbol{\epsilon} + \frac{1}{2} \boldsymbol{\epsilon}^\top (\nabla_{\mathbf{x}}^2 \mu(\mathbf{x})) \boldsymbol{\epsilon} + o(\|\boldsymbol{\epsilon}\|^2), \quad (4.1)$$

where $\nabla_{\mathbf{x}}^2 \mu$ denotes the Hessian of μ w.r.t. \mathbf{x} . Setting $\boldsymbol{\epsilon} = -\mathbf{x}$ allows us, in complete analogy to (3.3), to study *2nd order contributions*

$$F_{\mu(\mathbf{X})}^{-1}(\alpha) \approx \mu(\mathbf{0}) + \sum_{j=1}^q S_j(\mu; \alpha) - \frac{1}{2} \sum_{j,k=1}^q T_{j,k}(\mu; \alpha), \quad (4.2)$$

with 2nd order attributions, for $1 \leq j, k \leq q$,

$$T_{j,k}(\mu; \alpha) = \mathbb{E}_P \left[X_j X_k \mu_{j,k}(\mathbf{X}) \mid \mu(\mathbf{X}) = F_{\mu(\mathbf{X})}^{-1}(\alpha) \right]. \quad (4.3)$$

Slightly rearranging the terms in (4.1) allows us to study individual feature contributions and interaction terms separately, that is,

$$F_{\mu(\mathbf{X})}^{-1}(\alpha) \approx \mu(\mathbf{0}) + \sum_{j=1}^q \left(S_j(\mu; \alpha) - \frac{1}{2} T_{j,j}(\mu; \alpha) \right) - \sum_{1 \leq j < k \leq q} T_{j,k}(\mu; \alpha). \quad (4.4)$$

The latter term quantifies all 2nd order contributions coming from interactions between X_j and X_k , $j \neq k$. We will show how interaction effects can be included in individual features' marginal attributions towards the end of Section 6.4.

Remark 4.1 The motivation for studying 1st order attributions (3.1) has been given in terms of the risk sensitivity tools of [Hong 2009] and [Tsanakas and Millosovich 2015]. These are obtained by calculating directional derivatives of distortion risk measures (using a Dirac distortion, see Appendix A). This argumentation does not carry forward to the 2nd order terms (4.3), as 2nd order directional derivatives of distortion risk measures turn out to be much more complicated, even in the linear case, see Property 1 in [Gourieroux et al. 2000].

5 Choice of reference point

To obtain sufficient accuracy in 1st and 2nd order approximations, respectively, the reference point should lie somewhere “in the middle” of the feature distribution P . We elaborate on this in this section. Typically, we want to get the following expression small, uniformly in $\alpha \in (0, 1)$,

$$\left| F_{\mu(\mathbf{X})}^{-1}(\alpha) - \mu(\mathbf{0}) - \sum_{j=1}^q S_j(\mu; \alpha) + \frac{1}{2} \sum_{j,k=1}^q T_{j,k}(\mu; \alpha) \right|.$$

This expression is for reference point $\mathbf{0}$. However, we can select any other reference point $\mathbf{a} \in \mathbb{R}^q$, by exploring the 2nd order Taylor expansion (4.1) for $\boldsymbol{\epsilon} = \mathbf{a} - \mathbf{x}$. This latter reference point then provides us with a 2nd order approximation

$$\begin{aligned} F_{\mu(\mathbf{X})}^{-1}(\alpha) &\approx \mu(\mathbf{a}) - \mathbb{E}_P \left[(\mathbf{a} - \mathbf{X})^\top \nabla_{\mathbf{x}} \mu(\mathbf{X}) \mid \mu(\mathbf{X}) = F_{\mu(\mathbf{X})}^{-1}(\alpha) \right] \\ &\quad - \frac{1}{2} \mathbb{E}_P \left[(\mathbf{a} - \mathbf{X})^\top (\nabla_{\mathbf{x}}^2 \mu(\mathbf{X})) (\mathbf{a} - \mathbf{X}) \mid \mu(\mathbf{X}) = F_{\mu(\mathbf{X})}^{-1}(\alpha) \right]. \end{aligned} \quad (5.1)$$

The same can be received by translating the distribution P of the features by setting $\mathbf{X}^{\mathbf{a}} = \mathbf{X} - \mathbf{a}$ and letting $\mu^{\mathbf{a}}(\cdot) = \mu(\mathbf{a} + \cdot)$. Approximation (5.1) motivates us to look for a reference point $\mathbf{a} \in \mathbb{R}^q$ which makes the 2nd order approximation as accurate as possible for “all” quantile levels.

Being a bit less ambitious, we select a discrete quantile grid $0 < \alpha_1 < \dots < \alpha_L < 1$ on which we would like to have a good approximation capacity. Define the events $\mathcal{A}_l = \{\mu(\mathbf{X}) = F_{\mu(\mathbf{X})}^{-1}(\alpha_l)\}$ for $1 \leq l \leq L$. Consider the objective function

$$G(\mathbf{a}; \mu) = \sum_{l=1}^L \left(F_{\mu(\mathbf{X})}^{-1}(\alpha_l) - \mu(\mathbf{a}) + \mathbb{E}_P \left[(\mathbf{a} - \mathbf{X})^\top \nabla_{\mathbf{x}} \mu(\mathbf{X}) \middle| \mathcal{A}_l \right] + \frac{1}{2} \mathbb{E}_P \left[(\mathbf{a} - \mathbf{X})^\top (\nabla_{\mathbf{x}}^2 \mu(\mathbf{X})) (\mathbf{a} - \mathbf{X})^\top \middle| \mathcal{A}_l \right] \right)^2. \quad (5.2)$$

Minimizing this objective function in \mathbf{a} gives an optimal reference point w.r.t. the quantile levels $(\alpha_l)_{1 \leq l \leq L}$. Unfortunately, $\mathbf{a} \mapsto G(\mathbf{a}; \mu)$ is not a convex function, and therefore numerical methods may only find local minima. These can be found by a plain vanilla gradient descent method. We calculate the gradient of G w.r.t. \mathbf{a}

$$\begin{aligned} \nabla_{\mathbf{a}} G(\mathbf{a}; \mu) &= 2 \sum_{l=1}^L \left(F_{\mu(\mathbf{X})}^{-1}(\alpha_l) - \mu(\mathbf{a}) + \mathbb{E}_P \left[(\mathbf{a} - \mathbf{X})^\top \nabla_{\mathbf{x}} \mu(\mathbf{X}) \middle| \mathcal{A}_l \right] + \frac{1}{2} \mathbb{E}_P \left[(\mathbf{a} - \mathbf{X})^\top (\nabla_{\mathbf{x}}^2 \mu(\mathbf{X})) (\mathbf{a} - \mathbf{X})^\top \middle| \mathcal{A}_l \right] \right) \\ &\quad \times \left(-\nabla_{\mathbf{a}} \mu(\mathbf{a}) + \mathbb{E}_P \left[\nabla_{\mathbf{x}} \mu(\mathbf{X}) \middle| \mathcal{A}_l \right] - \mathbb{E}_P \left[\mathbf{X}^\top \nabla_{\mathbf{x}}^2 \mu(\mathbf{X}) \middle| \mathcal{A}_l \right] + \frac{1}{2} \mathbf{a}^\top \mathbb{E}_P \left[\nabla_{\mathbf{x}}^2 \mu(\mathbf{X}) \middle| \mathcal{A}_l \right] \right). \end{aligned}$$

The gradient descent algorithm then provides for a tempered learning rate $\varepsilon_{t+1} > 0$ updates at algorithmic time t

$$\mathbf{a}^{(t)} \mapsto \mathbf{a}^{(t+1)} = \mathbf{a}^{(t)} - \varepsilon_{t+1} \nabla_{\mathbf{a}} G(\mathbf{a}^{(t)}; \mu). \quad (5.3)$$

This step-wise locally decreases the objective function G .

Remark 5.1 The above algorithm provides a global optimal reference point, thus, a calibration for a global 2nd order meta-explanation. In some cases this global calibration may not be satisfactory, in particular, if the reference point is far from the feature values $\mathbf{X} = \mathbf{x}$ that mainly describe a given quantile level $F_{\mu(\mathbf{X})}^{-1}(\alpha)$, i.e. through the corresponding conditional probability $P[\cdot | \mu(\mathbf{X}) = F_{\mu(\mathbf{X})}^{-1}(\alpha)]$. In that case, one may be interested in different local reference points that are optimal for certain quantile levels, say, between 95% and 99%. In some sense, this will provide a more “honest” description (4.2) because we do not try to simultaneously describe all quantile levels. The downside of multiple reference points is that we lose comparability of marginal effects across the whole decision space.

6 Example

6.1 Model choice and model fitting

We consider the bike rental example of [Fanaee-T and Gama 2014] which has also been studied in [Apley and Zhu 2020]. The data describes the bike sharing process over the years 2011 and

2012 of the Capital Bikesharing system in Washington DC. On an hourly time grid we have information about the proportion of casual bike rentals relative to all bike rentals of casual and registered users. This data is supported by explanatory variables such as weather conditions and seasonal variables. We provide a descriptive analysis of this data in Appendix B. On average 17% of all bike rentals are made by casual users and 83% are done by registered users. However, these proportions heavily fluctuate w.r.t. daytime, holidays, weather conditions, etc. This variability is illustrated in Figure 12 in Appendix B. We design a neural network regression function to forecast the proportion of casual rentals. We denote this response variable (proportion) by Y , and we denote the features (explanatory variables) by $\mathbf{x} \in \mathbb{R}^q$.

For our example we choose a fully-connected feed-forward neural network $\theta : \mathbb{R}^q \rightarrow \mathbb{R}$ of depth $d = 3$ having $(q_1, q_2, q_3) = (20, 15, 10)$ neurons in the three hidden layers. This provides us with network regression function

$$\mathbf{x} \in \mathbb{R}^q \mapsto \mu(\mathbf{x}) = \sigma(\theta(\mathbf{x})) \in (0, 1), \quad (6.1)$$

where σ is the sigmoid output activation, and $\mathbf{x} \mapsto \theta(\mathbf{x})$ models the canonical parameter of a logistic regression model. In order to have a smooth network regression function we choose the hyperbolic tangent as activation function in the three hidden layers. We have implemented this network in [TensorFlow 2015] and [Keras 2015], these allow us to directly formally calculate gradients and Hessians.

In all what follows we do not consider the attributions of the regression function $\mathbf{x} \mapsto \mu(\mathbf{x})$, but we directly focus on the corresponding attributions on the canonical scale $\mathbf{x} \mapsto \theta(\mathbf{x})$. This has the advantage that the results do not get distorted by the sigmoid output activation. Thus, we replace μ by θ in (4.4), resulting in the study of 2nd order contributions

$$F_{\theta(\mathbf{X})}^{-1}(\alpha) \approx \theta(\mathbf{0}) + \sum_{j=1}^q \left(S_j(\theta; \alpha) - \frac{1}{2} T_{j,j}(\theta; \alpha) \right) - \sum_{1 \leq j < k \leq q} T_{j,k}(\theta; \alpha). \quad (6.2)$$

The network architecture is fitted to the available data using early stopping to prevent from overfitting. Importantly, we do not say here anything about the quality of the predictive model, but we aim at understanding the fitted regression function $\mathbf{x} \mapsto \theta(\mathbf{x})$. This can be done regardless whether the chosen model is suitable for the predictive task at hand.

Figure 1 (lhs) shows the empirical density of the canonical parameters $\mathbf{x}_i \mapsto \theta(\mathbf{x}_i)$ of the fitted model over our data $1 \leq i \leq n$. We have negative skewness in this empirical density. A simple way of analyzing importance of feature components is to randomly permute one component x_j at a time across all records $1 \leq i \leq n$ and study the increase in objective function; this is the method of variable permutation importance introduced by [Breiman 2001]. We use as objective function the Bernoulli deviance loss which is proportional to the binary cross-entropy (also called log loss). Figure 1 (rhs) shows the variable permutation importances. There are three variables (hour, working day and temperature) that highly dominate the others. Note that variable permutation importance does not properly consider the dependence structure in \mathbf{X} , similarly to ICEs and PDPs, because permutation of x_j is done without impacting $\mathbf{x}_{\setminus j}$.

6.2 1st and 2nd order contributions

The accuracy of the 2nd order contributions (6.2) will depend on the choice of the reference point $\mathbf{a} \in \mathbb{R}^q$. For network gradient descent fitting we have normalized the feature components

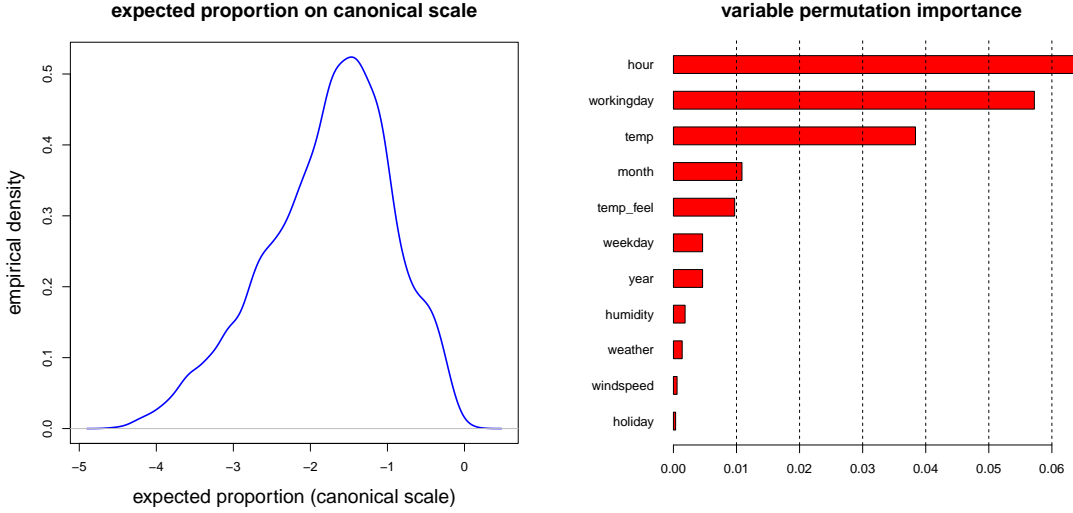


Figure 1: (lhs) Empirical density of canonical parameters $(\theta(\mathbf{x}_i))_{1 \leq i \leq n}$, (rhs) variable permutation importance.

to be centered and having unit variance, i.e. $\mathbb{E}_P[\mathbf{X}] = \mathbf{0}$ and $\text{Var}_P(X_j) = 1$ for all $1 \leq j \leq q$. This pre-processing is needed to efficiently apply stochastic gradient descent network fitting. We now translate these feature components by choosing a reference point \mathbf{a} such that the objective function $G(\mathbf{a}; \theta)$ is minimized, see (5.2). We use a plain vanilla gradient descent update (5.3) using a learning rate of $\varepsilon_{t+1} = 10^{-2} / \|\nabla_{\mathbf{a}} G(\mathbf{a}^{(t)}; \theta)\|$. For the quantile grid we choose $\alpha_l = l/100$ for $1 \leq l \leq L = 99$, thus, $\alpha \in \{1\%, \dots, 99\%\}$. The resulting decrease in objective function $G(\cdot; \theta)$ is plotted in Figure 2.




Figure 2: (lhs) Gradient descent for reference point \mathbf{a} , (rhs) 2nd order contributions (6.2).

Working with observed data, we need to discretize the MACQ analysis for quantile levels

$\{\theta(\mathbf{X}) = F_{\theta(\mathbf{X})}^{-1}(\alpha)\}$, $\alpha \in (0, 1)$. We do this on a discrete grid by using a local smoother of degree 2, in particular, we use the R function `locfit` with parameters `deg=2` and `alpha=0.1` (being the chosen bandwidth) for observations $x_{i,j}^{\mathbf{a}}\theta_j(\mathbf{x}_i)$ and $x_{i,j}^{\mathbf{a}}x_{i,k}^{\mathbf{a}}\theta_{j,k}(\mathbf{x}_i)$, $1 \leq i \leq n$, where we set $\mathbf{x}_i^{\mathbf{a}} = \mathbf{x}_i - \mathbf{a}$. We then fit the local smoother to these observations being ordered according to the ranks of $\theta(\mathbf{x}_i)$, to work with the corresponding empirical output quantiles. Thus, for instance, the \mathbf{a} -adjusted 1st order attributions $S_j(\theta; \alpha_l)$, $1 \leq l \leq L$, are estimated empirically by using the pseudo code

```
predict(locfit(x_{i,j}^{\mathbf{a}}\theta_j(\mathbf{x}_i) \sim \text{rank}(\theta(\mathbf{x}_i))/n, \text{alpha} = 0.1, \text{deg} = 2), \text{newdata} = \text{c}(1 : 99)/100).
```

Figure 2 (rhs) gives the results after optimizing for the reference point \mathbf{a} . The orange color shows the 1st order contributions $C_1 = \theta(\mathbf{a}) + \sum_{j=1}^q S_j(\theta; \alpha)$, the cyan line shows the 2nd order contributions without interaction terms $C_2 = \theta(\mathbf{a}) + \sum_{j=1}^q (S_j(\theta; \alpha) - \frac{1}{2}T_{j,j}(\theta; \alpha))$ and the red line shows the full 2nd order contributions $C_{2,2} = \theta(\mathbf{a}) + \sum_{j=1}^q (S_j(\theta; \alpha) - \frac{1}{2}T_{j,j}(\theta; \alpha)) - \sum_{1 \leq j < k \leq q} T_{j,k}(\theta; \alpha)$.

We observe from Figure 2 (rhs) that the full 2nd order contributions $C_{2,2}$ match the empirical quantiles (black dots) quite well which explains that there is a reference point \mathbf{a} that allows for suitable 2nd order approximations over the entire quantile set. The shaded cyan area between C_2 (cyan line) and $C_{2,2}$ (red line) shows the influence of the interaction terms $T_{j,k}(\theta; \alpha)$, $j \neq k$, which illustrates that this model undergoes substantial interactions, and a simple generalized additive model (GAM) will not be able to model this data accurately.




Figure 3: (lhs) attributions $S_j(\theta; \alpha) - \frac{1}{2}T_{j,j}(\theta; \alpha)$ excluding interaction terms, see (6.2), (rhs) attributions $S_j(\theta; \alpha) - \frac{1}{2}T_{j,j}(\theta; \alpha)$ for selected quantile levels $\alpha \in \{20\%, 40\%, 60\%, 80\%\}$.

In Figure 3 (lhs) we show the attributions $S_j(\theta; \alpha) - \frac{1}{2}T_{j,j}(\theta; \alpha)$, excluding interaction terms $T_{j,k}(\theta; \alpha)$, $j \neq k$, relative to the optimal reference point \mathbf{a} . These attributions show the differences relative to canonical parameter in the reference point $\theta(\mathbf{a})$; when aggregating over $1 \leq j \leq q$ this results in the cyan line of Figure 2 (rhs). Figure 3 (lhs) shows substantial sensitivities in the variables month, hour, working day and temperature. From this we conclude that these are the important variables in our regression model for differentiating the responses

Y w.r.t. available feature information \mathbf{x} . In contrast to the variable permutation importance plot of Figure 1 (rhs), this assessment correctly considers the dependence structure within the features \mathbf{X} . Moreover, this plot now allows us to analyze variable importance on different quantile levels by considering vertical slices in Figure 3 (lhs). We consider such vertical slices in Figure 3 (rhs) for four selected quantile levels $\alpha \in \{20\%, 40\%, 60\%, 80\%\}$. We observe that the variables `month`, `hour`, `workingday` and `temp` undergo the biggest changes when moving from small quantiles to big ones. The quantile level at 20% can be explained by the three features `temp`, `month` and `workingday`, whereas for the quantile level at 60% has `hour` (daytime) as an important variable, see Figure 3 (rhs). Note that this is not the full picture, yet, as we do not consider interactions in these vertical slices; the importance of interactions is indicated by the cyan shaded area in Figure 2 (rhs) for different quantile levels.

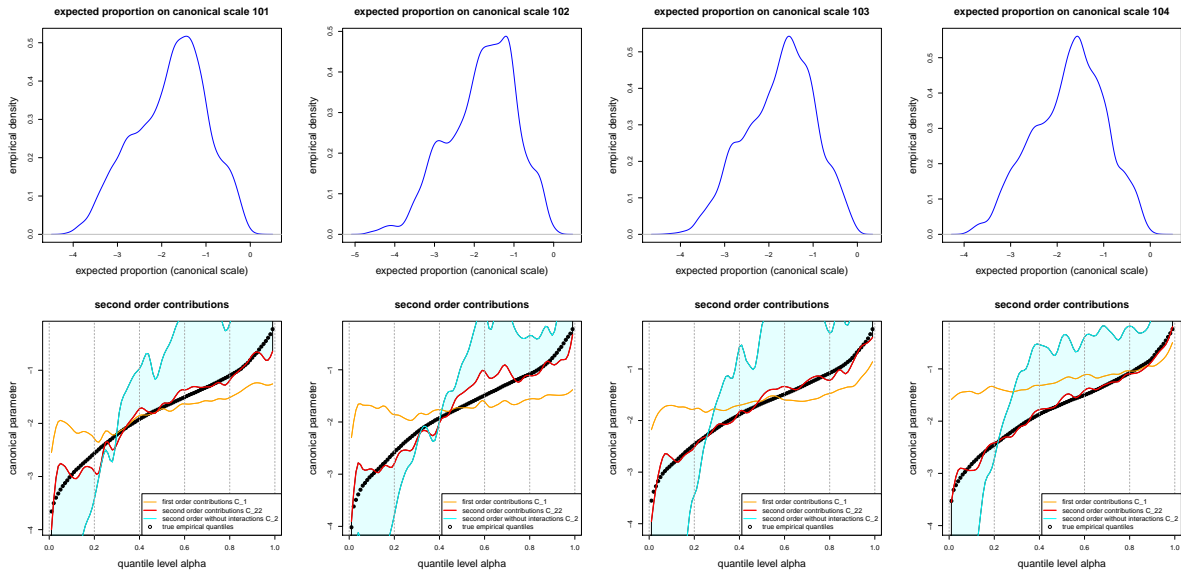


Figure 4: Robustness of 2nd order contributions across 4 different networks: (top row) empirical densities of canonical parameters $(\theta(\mathbf{x}_i))_{1 \leq i \leq n}$, (bottom row) 2nd order contributions (6.2).

In Figure 4 we analyze the robustness of the attribution results. We do this by considering different networks $\mathbf{x} \mapsto \theta(\mathbf{x})$ for predicting the response variable Y . Network regression models lack a certain degree of robustness as gradient descent network fitting explores different (local) minima of the objective function; note that, in general, neural network fitting is not a convex minimization problem. This issue of non-uniqueness of good predictive models has been widely discussed in the literature, and ensembling may be one solution to mitigate this problem, we refer to [Dietterich 2000a, Dietterich 2000b], [Zhou et al. 2002], [Zhou 2012] and [Richman and Wüthrich 2020]. The top row shows the empirical distributions of the canonical parameters $(\theta(\mathbf{x}_i))_{1 \leq i \leq n}$ for 4 different networks; we observe that there are some differences in these empirical densities. The bottom row shows the corresponding 2nd order contributions (6.2), split by 1st order contributions C_1 , 2nd order contributions without interactions C_2 and the full 2nd order contributions $C_{2,2}$. At this level, we judge the attributions made to be rather robust over the different models, the general shapes of these graphs being similar, and also the interaction terms $C_{2,2} - C_2$ showing a similar structure and magnitude across the 4 different

network models.




Figure 5: Choice of reference point \mathbf{a} across 4 different networks illustrated for all feature components $1 \leq j \leq q$.

From Figure 4 we also observe that the 1st order contributions C_1 intersect the quantiles $F_{\theta(\mathbf{X})}^{-1}(\alpha)$ at different levels for the 4 different calibrations. This indicates that the optimal reference point \mathbf{a} is chosen differently in the different networks. Figure 5 shows the chosen reference points \mathbf{a} in relation to the features $(\mathbf{x}_i)_{1 \leq i \leq n}$; as explained above, we have centered and normalized the feature components for gradient descent network fitting. The boxplots in Figure 5 show these centered and normalized features in comparison to the reference points of the 4 different networks. Some feature components have a very skewed distribution as can be seen from the thicker horizontal boxplot lines showing the median of each feature component $(x_{i,j})_{1 \leq i \leq n}$, $1 \leq j \leq q$. The reference point mostly lies within the interquartile range (IQR).

Remark 6.1 The feature components of \mathbf{x} need pre-processing in order to be suitable for gradient descent fitting. Continuous and binary variables have been centered and normalized so that their gradients live in a similar range. This makes gradient descent fitting more efficient because all partial derivatives of the gradient are directly comparable. Our example does not have categorical feature components. Categorical feature components can be treated in different ways. For our MACQ proposal we envisage two different treatments. Firstly, dummy coding could be used. This requires the choice of a reference level, and considers all other levels relative to this reference level. The resulting marginal attributions should then be interpreted as differences to the reference level. Secondly, one can use embedding layers for categorical variables, see [Bengio et al. 2003] and [Guo and Berkhahn 2016]. In that case the attribution analysis can directly be done on these learned embeddings of categorical levels, in complete analogy to the continuous variables.

6.3 Attribution to individual instances

Next, we focus on individual instances $\mathbf{x}_i^a = \mathbf{x}_i - \mathbf{a}$ and study individual marginal contributions $\omega_{i,j} = (x_{i,j} - a_j)\theta_j(\mathbf{x}_i) - (x_{i,j} - a_j)^2\theta_{j,j}(\mathbf{x}_i)/2$ to attribution $S_j(\theta; \alpha) - T_{j,j}(\theta; \alpha)/2$.

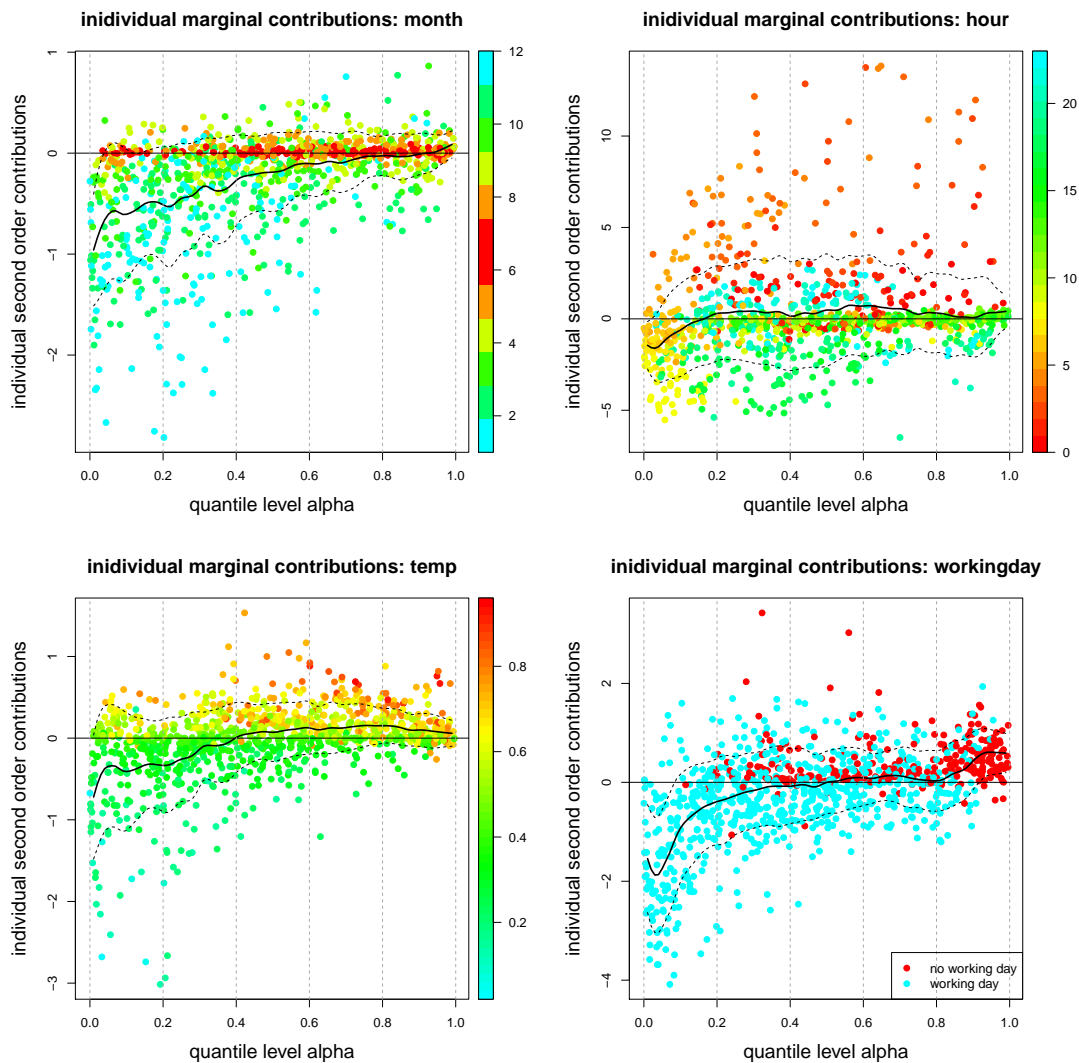


Figure 6: Individual marginal contributions $\omega_{i,j}$ of 1,000 randomly selected instances \mathbf{x}_i for (top-left) $j = \text{month}$, (top-right) $j = \text{hour}$, (bottom-left) $j = \text{temp}$ and (bottom-right) $j = \text{workingday}$; the black line shows attribution $S_j(\theta; \alpha) - T_{j,j}(\theta; \alpha)/2$ and the black dotted line gives one standard deviation; the y -scales differs in the plots and the colors illustrate the feature values x_j .

For Figure 6 we select at random 1,000 different instances, and plot their individual marginal contributions $\omega_{i,j}$ to the attributions $S_j(\theta; \alpha) - T_{j,j}(\theta; \alpha)/2$ (black solid line). The ordering on the x -axis for the selected instances \mathbf{x}_i is obtained by considering the empirical quantiles of the responses $\theta(\mathbf{x}_k)$ over all instances $1 \leq k \leq n$. We start with Figure 6 (bottom-right) which shows the binary variable `workingday`. This variable clearly differentiates low from high

quantiles $F_{\theta(\mathbf{X})}^{-1}(\alpha)$, showing that the casual rental proportion Y is in average bigger for non-working days (red dots). Moreover, for low quantiles levels the working day variable clearly lowers (expected response) $\theta(\mathbf{x})$ compared to the reference level $\theta(\mathbf{a})$, as the cyan dots are below the horizontal black line at 0 which corresponds to the reference level. In addition to the average attributions $S_j(\theta; \alpha) - T_{j,j}(\theta; \alpha)/2$ (black solid line), the plot is complemented by black dotted lines giving one (empirical) standard deviation

$$\text{Var}_P \left((X_j - a_j)\theta_j(\mathbf{X}) - (X_j - a_j)^2\theta_{j,j}(\mathbf{X})/2 \mid \theta(\mathbf{X}) = F_{\theta(\mathbf{X})}^{-1}(\alpha) \right)^{1/2}.$$

The sizes of these standard deviations quantify the heterogeneity in the individual marginal contributions $\omega_{i,j}$. This can either be because of heterogeneity of the portfolio $x_{i,j}$ on a certain quantile level, or because we have a rough regression surface implying heterogeneity in derivatives $\theta_j(\mathbf{x}_i)$ and $\theta_{j,j}(\mathbf{x}_i)$.

Next, we study the variable `temp` of Figure 6 (bottom-left). In this plot we see a clear positive dependence between quantile levels and temperature, showing that casual rentals are generally low for low temperatures, which can either be the calendar season or bad weather conditions. We have clearly more heterogeneity in features (and resulting derivatives $\theta_j(\mathbf{x}_i)$ and $\theta_{j,j}(\mathbf{x}_i)$) contributing to low quantile levels than to higher ones. The variable `temp` is highly correlated with calendar month, and the calendar month plot in Figure 6 (top-left) looks similar, saying that casual rental proportions Y are negatively impacted by winter seasons. There are some low proportions, though, also for summer months, these need to be explained by other variables, e.g., they may correspond to a rainy day or to a specific daytime. The interpretation of the variable `hour` in Figure 6 (top-right) is slightly more complicated since we do not have monotonicity of $\theta(\mathbf{x})$ in this variable, see also Figure 12. Nevertheless we also see a separation between working and leisure times (for the time-being ignoring interactions with holidays and weekends).




Figure 7: Individual marginal contributions $\omega_{i,j}$ of 1,000 randomly selected instances \mathbf{x}_i for (lhs) $j = \text{month}$, (middle) $j = \text{hour}$ and (rhs) $j = \text{temp}$; the black line shows the empirical average; the colors show the expected responses $\mu(\mathbf{x}_i) \in (0, 1)$ (casual rental proportions).

In Figure 6 we have plotted the individual marginal contributions $\omega_{i,j}$ on the y -axis against the quantiles $\alpha \in (0, 1)$ on the x -axis to explain how the features \mathbf{x}_i enter the quantile levels $F_{\theta(\mathbf{X})}^{-1}(\alpha)$. This is the 3-ways analysis mentioned above, where the third dimension is highlighted by using different colors in Figure 6. Alternatively, we can also try to understand how this third dimension

of different feature values x_j contributes to the individual marginal contributions $\omega_{i,j}$. Figure 7 plots the individual marginal contributions $\omega_{i,j}$ on the y -axis against the feature values x_j on the x -axis. The black line shows the averages of $\omega_{i,j}$ over all instances, and the colored dots show the 1,000 randomly selected instances \mathbf{x}_i with the colors illustrating the expected responses, i.e. the expected casual rental proportions $\mu(\mathbf{x}_i) = \sigma(\theta(\mathbf{x}_i)) \in (0, 1)$. The general shape of the black lines in these graphs reflects well the marginal empirical observations in Figure 12. However, the detailed structure slightly differs in these plots as they do not exactly show the same quantity, the latter shows a marginal empirical graph, whereas Figure 7 quantifies individual marginal contributions to expected responses $\theta(\mathbf{x})$ in an additive way (on the canonical scale). Figure 7 (rhs) shows a clear monotone plot which also results in a separation of the colors, whereas the colors in Figure 7 (lhs, middle) can only be fully understood by also studying contributions and interactions with other components $x_{i,k}$, $k \neq j$.

6.4 Interaction terms

There remains the analysis of the interaction terms $-T_{j,k}(\theta; \alpha)$, $j \neq k$, that account for the cyan shaded area in Figure 2 (rhs). These interaction terms are shown in Figure 8.




Figure 8: Off-diagonal terms $-T_{j,k}(\theta; \alpha)$ giving the interactions.

To not overload Figure 8 we only show the interaction terms $T_{j,k}$ for which $\max_{\alpha} |T_{j,k}(\theta; \alpha)| > 0.2$. We identify three major interaction terms: **workingday-hour**, **workingday-month** and **hour-month**. Of course, these interactions make perfect sense in describing the casual rental proportion. For small quantiles also interactions **temp-month** and **temp-hour** are important. Interestingly, we also find an interaction **workingday-year**: in the data there is a positive trend of registered rental bike users (in absolute terms) which interacts differently on working and non-working days because casual rentals are more frequent on non-working days. Identifying the importance of these interactions highlights that it will not be sufficient to work within a generalized linear model (GLM) or a generalized additive model (GAM) unless we add explicit interaction terms to them.




Figure 9: (lhs) 2nd order attributions $V_j(\theta; \alpha)$ including interaction terms, and (rhs) $V_j(\theta; \alpha)$ for selected quantile levels $\alpha \in \{20\%, 40\%, 60\%, 80\%\}$.

In the final step we combine the attributions $S_j(\theta; \alpha) - T_{j,j}(\theta; \alpha)/2$ with the interaction terms $T_{j,k}(\theta; \alpha)$, $k \leq j$. A natural way is to just allocate half of the interaction terms $T_{j,k}(\theta; \alpha)$ to each component j and k . This then provides allocated 2nd order attribution to components $1 \leq j \leq q$

$$V_j(\theta; \alpha) = S_j(\theta; \alpha) - T_{j,j}(\theta; \alpha)/2 - \sum_{j \neq k} T_{j,k}(\theta; \alpha)/2 = S_j(\theta; \alpha) - \sum_{k=1}^q T_{j,k}(\theta; \alpha)/2.$$

Adding the reference level $\theta(\mathbf{a})$, we again receive the full 2nd order contributions $C_{2,2} = \theta(\mathbf{a}) + \sum_{j=1}^q V_j(\theta; \alpha)$ illustrated by the red line in Figure 2 (rhs). In Figure 9 we provide these attributions $V_j(\theta; \alpha)$ for quantiles $\alpha \in (0, 1)$. These plots differ from Figure 3 only by the inclusion of the 2nd order off-diagonal (interaction) terms. Comparing the right-hand sides of these two plots we observe that firstly the level is shifted, which is explained by the shaded cyan area in Figure 2 (rhs). Secondly, interactions impact mainly the small quantiles in our example, this is clear from Figure 8 and, for instance, impacts the significance of **hour** on the 20% quantile level.

6.5 Scrolling through the network layers

Up to this point our MACQ analysis has been fully general, in the sense that it can be applied to any smooth deep learning model. In the last step of our analysis we specifically focus on the deep network introduced in Section 6.1, and we try to better understand how networks learn new representations through the network layers. A deep feed-forward neural network $\theta : \mathbb{R}^q \rightarrow \mathbb{R}$ is a composition of d hidden neural network layers $\mathbf{z}^{(k)} : \mathbb{R}^{q_{k-1}} \rightarrow \mathbb{R}^{q_k}$, $1 \leq k \leq d$; we initialize input dimension $q_0 = q$. Define the composition $\mathbf{x} \mapsto \mathbf{z}^{(d:1)}(\mathbf{x}) = (\mathbf{z}^{(d)} \circ \dots \circ \mathbf{z}^{(1)})(\mathbf{x})$ which maps input $\mathbf{x} \in \mathbb{R}^q$ to the last hidden network layer having dimension q_d . Network (6.1) with

logistic output can then be written as

$$\mathbf{x} \in \mathbb{R}^q \mapsto \mu(\mathbf{x}) = \sigma(\theta(\mathbf{x})) = \sigma\left(\beta_0 + \boldsymbol{\beta}^\top \mathbf{z}^{(d:1)}(\mathbf{x})\right),$$

with bias/intercept $\beta_0 \in \mathbb{R}$ and regression parameter/weight $\boldsymbol{\beta} \in \mathbb{R}^{q_d}$. This should be compared to linear regression (3.4).

Each hidden layer learns a new representation of the inputs \mathbf{x}_i , that is, the representations learned in layer k are given by $\mathbf{x}_i^{(k:1)} := (\mathbf{z}^{(k)} \circ \dots \circ \mathbf{z}^{(1)})(\mathbf{x}_i)$, for $1 \leq i \leq n$. We can view these learned representations as new inputs to the remaining network after hidden layer k

$$\mathbf{x} \in \mathbb{R}^{q_k} \mapsto \sigma\left(\beta_0 + \boldsymbol{\beta}^\top \mathbf{z}^{(d:k+1)}(\mathbf{x})\right) = \sigma\left(\beta_0 + \boldsymbol{\beta}^\top (\mathbf{z}^{(d)} \circ \dots \circ \mathbf{z}^{(k+1)})(\mathbf{x})\right).$$

In the following analysis we consider the instances $(Y_i, \mathbf{x}_i^{(k:1)})$ with these learned features $\mathbf{x}_i^{(k:1)}$ as inputs to the remaining network $\mathbf{z}^{(d:k+1)}$ after layer k , and we perform the same MACQ analysis in this reduced setup.




Figure 10: 2nd order contributions (6.2) of the (learned) representations: (lhs) original inputs \mathbf{x}_i , (middle) learned representations $\mathbf{x}_i^{(1:1)}$, and (rhs) learned representations $\mathbf{x}_i^{(2:1)}$.

Figure 10 provides the 2nd order contributions (6.2) of the original inputs (lhs), the learned representations $\mathbf{x}_i^{(1:1)}$ in the first hidden layer (middle), and the learned representations $\mathbf{x}_i^{(2:1)}$ in the second hidden layer (rhs) on the corresponding remaining networks $\mathbf{z}^{(3:k+1)}$. We interpret these MACQ results as follows. The first hidden layer (middle graph) has mainly a smoothing effect in recomposing the inputs \mathbf{x}_i suitably. The second layer takes care of the interaction effects diminishing the cyan shaded area in Figure 10 (rhs). Of course, this makes perfect sense as the output layer considers a linear function with weight $\boldsymbol{\beta} \in \mathbb{R}^{q_d}$ which no longer allows for interactions. Therefore, interactions need to be learned in the previous layers. The same applies to non-linear structures (on the canonical scale). This completes our example.

7 Conclusions

This manuscript proposes a novel gradient-based global model-agnostic tool that can be calculated efficiently for differentiable deep learning models and produces informative visualizations. This tool studies marginal attribution to feature components on a given response level. Marginal

attributions allow us to separate marginal effects of individual feature components from interaction effects, and they allow us to study resulting variable importance plots on different parts of the decision space characterized by different response levels. This variable importance is measured w.r.t. a reference point that calibrates the entire space for our explanation. Finding a good reference point has been efficiently performed by a simple gradient descent search. A main result of our model-agnostic tool is a 3-way relationship between marginal attribution, output level and feature value which can be illustrated in different ways. This extends response sensitivity analyses, such as accumulated local effects, by an additional marginal attribution view.

References

- [TensorFlow 2015] Abadi, M., et al. (2015) TensorFlow: large-scale machine learning on heterogeneous systems. <https://www.tensorflow.org/>
- [Acerbi 2002] Acerbi, C. (2002). Spectral measures of risk: a coherent representation of subjective risk aversion. *Journal of Banking and Finance* **7**, 1505-1518.
- [Ancona et al. 2019] Ancona, M., Ceolini, E., Öztireli, C., Gross, M. (2019). Gradient-based attribution methods. In: *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*. Samek, W., Montavon, G., Vedaldi, A., Hansen, L.K., Müller K.-R. (Eds.). Springer, Lecture Notes in Artificial Intelligence 11700, 168-191.
- [Apley and Zhu 2020] Apley, D.W., Zhu, J. (2020). Visualizing the effects of predictor variables in black box supervised learning models. *Journal of the Royal Statistical Society: Series B* **82/4**, 1059-1086
- [Bengio et al. 2003] Bengio Y., Ducharme R., Vincent P., Jauvin C. (2003). A neural probabilistic language model. *Journal of Machine Learning Research* **3/Feb**, 1137-1155.
- [Binder et al. 2016] Binder, A., Bach, S., Montavon, G., Müller K.-R., Samek, W. (2016). Layer-wise relevance propagation for deep neural network architectures. In: *Information Science and Applications (ICISA)*. Kim K., Joukov N. (Eds.). Springer, Lecture Notes in Electrical Engineering 376.
- [Breiman 2001] Breiman, L. (2001). Random forests. *Machine Learning* **45/1**, 5-32.
- [Keras 2015] Chollet, F., et al. (2015). Keras. <https://github.com/fchollet/keras>
- [Dietterich 2000a] Dietterich, T.G. (2000). An experimental comparison of three methods for constructing ensembles of decision trees: bagging, boosting, and randomization. *Machine Learning* **40/2**, 139-157.
- [Dietterich 2000b] Dietterich, T.G. (2000). Ensemble methods in machine learning. In: *Multiple Classifier Systems*, J. Kittel, F. Roli (eds.). Lecture Notes in Computer Science, 1857, Springer, 1-15.
- [Efron 2020] Efron, B. (2020). Prediction, estimation and attribution. *International Statistical Review* **88/S1**, S28-S59.
- [Fanaee-T and Gama 2014] Fanaee-T, H., Gama, J. (2014). Event labeling combining ensemble detectors and background knowledge. *Progress in Artificial Intelligence* **2**, 113-127.
- [Friedman 2001] Friedman, J.H. (2001). Greedy function approximation: a gradient boosting machine. *Annals of Statistics* **29/5**, 1189-1232.
- [Friedman and Popescu 2008] Friedman, J.H., Popescu, B.E. (2008). Predictive learning via rule ensembles. *Annals of Applied Statistics* **2/3**, 916-954.

- [Goldstein et al. 2015] Goldstein, A., Kapelner, A., Bleich, J., Pitkin, E. (2015). Peeking inside the black box: visualizing statistical learning with plots of individual conditional expectation. *Journal of Computational and Graphical Statistics* **24/1**, 44-65.
- [Gourieroux et al. 2000] Gourieroux, C., Laurent, J.P., Scaillet, O. (2000). Sensitivity analysis of values at risk. *Journal of Empirical Finance* **7**, 225-245.
- [Guo and Berkahn 2016] Guo, C., Berkahn, F. (2016). Entity embeddings of categorical variables. *arXiv:1604.06737*.
- [Hong 2009] Hong, L.J. (2009). Estimating quantile sensitivities. *Operations Research* **57/1**, 118-130.
- [Lundberg and Lee 2017] Lundberg, S.M., Lee, S.-I. (2017). A unified approach to interpreting model predictions. In: *Advances in Neural Information Processing Systems* **30**, Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.), 4765-74. Montreal: Curran Associates.
- [Miller 2019] Miller, T. (2019). Explanation in artificial intelligence: insights from social sciences. *Artificial Intelligence* **267**, 1-38.
- [Montavon et al. 2017] Montavon, G., Lapuschkin, S., Binder, A., Samek, W., Müller K.-R. (2017). Explaining nonlinear classification decisions with deep Taylor decomposition. *Pattern Recognition* **65**, 211-222.
- [Ribeiro et al. 2016] Ribeiro, M.T., Singh, S., Guestrin, C. (2016). “Why should I trust you?”: explaining the predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*. New York: Association for Computing Machinery, 1135-1144.
- [Richman and Wüthrich 2020] Richman, R., Wüthrich, M.V. (2020). Nagging predictors. *Risks* **8/3**, article 83.
- [Samek and Müller 2019] Samek, W., Müller K.-R. (2019). Toward explainable artificial intelligence. In: *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*. Samek, W., Montavon, G., Vedaldi, A., Hansen, L.K., Müller K.-R. (Eds.). Springer, Lecture Notes in Artificial Intelligence 11700, 5-23.
- [Shapley 1953] Shapley, L.S. (1953). A Value for n-Person Games. In: *Contributions to the Theory of Games (AM-28)*, Vol. II. Kuhn, H.W., Tucker, A.W. (eds.), Princeton University Press, 307-318.
- [Shrikumar et al. 2017] Shrikumar, A., Greenside, P., Kundaje, A. (2017). Learning important features through propagating activation differences. In: *Proceedings of the 34th International Conference on Machine Learning, Proceedings of Machine Learning Research, PMLR*. International Convention Centre, Sydney, Australia, **70**, 3145-3153.
- [Shrikumar et al. 2016] Shrikumar, A., Greenside, P., Shcherbina, A., Kundaje, A. (2016). Not just a black box: learning important features through propagating activation differences. *arXiv:1605.01713*.
- [Sundararajan et al. 2017] Sundararajan, M., Taly, A., Yan, Q. (2017). Axiomatic attribution for deep networks. In: *Proceedings of the 34th International Conference on Machine Learning, Proceedings of Machine Learning Research, PMLR*. International Convention Centre, Sydney, Australia, **70**, 3319-3328.
- [Tsanakas and Millossovich 2015] Tsanakas, A., Millossovich, P. (2015). Sensitivity analysis using risk measures. *Risk Analysis* **36/1**, 30-48.
- [Wang 1996] Wang, S. (1996). Premium calculation by transforming the layer premium density. *ASTIN Bulletin* **26/1**, 71-92.

- [Zhao and Hastie 2021] Zhao, Q., Hastie, T. (2021). Causal interpretations of black-box models. *Journal of Business & Economic Statistics* **39/1**, 272-281.
- [Zhou 2012] Zhou, Z.-H. (2012). *Ensemble Methods: Foundations and Algorithms*. Chapman & Hall/CRC.
- [Zhou et al. 2002] Zhou, Z.-H., Wu, J., Tang, W. (2002). Ensembling neural networks: many could be better than all. *Artificial Intelligence* **137/1-2**, 239-263.

A Sensitivities in distortion risk measures

The purpose of this appendix is to briefly explain distortion risk measures and how they relate to marginal attribution. For this discussion we impose stronger assumptions than we need above, i.e., these more restrictive assumptions are only made for the explanation here. Assume the expected response $\mu(\mathbf{X})$ has a continuous distribution function $F_{\mu(\mathbf{X})}$. It follows that $U_{\mu(\mathbf{X})} = F_{\mu(\mathbf{X})}(\mu(\mathbf{X}))$ is uniformly distributed on $[0, 1]$. Choose a density ζ on $[0, 1]$. We can interpret $\zeta(U_{\mu(\mathbf{X})})$ as a probability distortion (probability re-weighting scheme inducing a change of probability measure) because we have

$$\mathbb{E}_P [\zeta(U_{\mu(\mathbf{X})})] = \int_0^1 \zeta(u) du = 1.$$

The distorted expected response can then be defined by

$$\varrho(\mu(\mathbf{X}); \zeta) = \mathbb{E}_P [\mu(\mathbf{X})\zeta(U_{\mu(\mathbf{X})})].$$

The functional $\varrho(\mu(\mathbf{X}); \zeta)$ describes a *distortion risk measure*, see [Wang 1996] and [Acerci 2002]. It can be interpreted as a Radon–Nikodým derivative changed probability measure $dP_\zeta(\mathbf{X} = \mathbf{x}) = \zeta(U_{\mu(\mathbf{x})})dP(\mathbf{X} = \mathbf{x})$. We study the sensitivities of this distortion risk measure w.r.t. the components of \mathbf{X} . Assume that the following directional derivatives exist in zero for all $1 \leq j \leq q$

$$S_j(\mu; \zeta) = \frac{\partial}{\partial \varepsilon} \varrho \left(\mu \left((X_1, \dots, X_{j-1}, X_j(1 + \varepsilon), X_{j+1}, \dots, X_q)^\top \right); \zeta \right) \Big|_{\varepsilon=0}.$$

Then, $S_j(\mu; \zeta)$ can be interpreted as the sensitivity of $\mathbf{X} \mapsto \mu(\mathbf{X})$ in feature component X_j . [Hong 2009] and [Tsanakas and Millosovich 2015] prove under different sets of assumptions that these sensitivities satisfy

$$S_j(\mu; \zeta) = \mathbb{E}_P [X_j \mu_j(\mathbf{X}) \zeta(U_{\mu(\mathbf{X})})].$$

Observe that this exactly uses the marginal attribution (2.6). We still have the freedom of choosing the density ζ on $[0, 1]$. If we choose the uniform distribution $\zeta \equiv 1$ on $[0, 1]$ we receive the average expected response and its average marginal attribution

$$\varrho(\mu(\mathbf{X}); \zeta \equiv 1) = \mathbb{E}_P[\mu(\mathbf{X})] \quad \text{and} \quad S_j(\mu; \zeta \equiv 1) = \mathbb{E}_P[X_j \mu_j(\mathbf{X})].$$

If we choose for density ζ the Dirac measure δ_α in $\alpha \in (0, 1)$, which allocates probability weight 1 to α , this gives us the α -quantile

$$\varrho(\mu(\mathbf{X}); \zeta = \delta_\alpha) = F_{\mu(\mathbf{X})}^{-1}(\alpha).$$

For its sensitivities we receive for $1 \leq j \leq q$

$$S_j(\mu; \zeta = \delta_\alpha) = \mathbb{E}_P \left[X_j \mu_j(\mathbf{X}) \Big| \mu(\mathbf{X}) = F_{\mu(\mathbf{X})}^{-1}(\alpha) \right],$$

which exactly corresponds to 1st order attribution (3.1).

Remark. We could choose any other density ζ on $[0, 1]$ to obtain sensitivities of other distortion risk measures. Such other choices may also have interesting counterparts in interpreting smooth deep learning models, by reflecting attention to different areas of the prediction space.

B Descriptive analysis of bike rental example

In this appendix, we give a brief descriptive analysis of the data used that helps us to interpret the network regression models. The data comprises the number of casual and registered bike rentals every hour from 2011/01/01 until 2012/12/31. This data has originally been studied in [Fanaee-T and Gama 2014] and [Apley and Zhu 2020], and it can be downloaded from <https://archive.ics.uci.edu/ml/datasets/Bike+Sharing+Dataset>. Listing 1 gives a short excerpt of the data.

Listing 1: Excerpt of bike rental data.

```
1 'data.frame':  17379 obs. of  13 variables:
2 $ date      : Date, format: "2011-01-01" "2011-01-01" "2011-01-01" ...
3 $ year      : num  2011 2011 2011 2011 2011 ...
4 $ month     : int  1 1 1 1 1 1 1 1 1 ...
5 $ hour      : int  0 1 2 3 4 5 6 7 8 9 ...
6 $ weekday   : int  6 6 6 6 6 6 6 6 6 ...
7 $ holiday   : Factor w/ 2 levels "holiday","no-holiday": 2 2 2 2 2 2 2 2 2 ...
8 $ workingday: Factor w/ 2 levels "no-working","workingday": 1 1 1 1 1 1 1 1 1 ...
9 $ weather   : num  1 1 1 1 1 2 1 1 1 ...
10 $ temp      : num  0.24 0.22 0.22 0.24 0.24 0.24 0.22 0.2 0.24 0.32 ...
11 $ temp_feel : num  0.288 0.273 0.273 0.288 0.288 ...
12 $ humidity  : num  0.81 0.8 0.8 0.75 0.75 0.75 0.8 0.86 0.75 0.76 ...
13 $ windspeed : num  0 0 0 0 0.0896 0 0 0 0 ...
14 $ casual    : int  3 8 5 3 0 0 2 1 1 8 ...
15 $ registered: int  13 32 27 10 1 1 0 2 7 6 ...
16 $ count     : int  16 40 32 13 1 1 2 3 8 14 ...
```

As response variable we consider the proportion of casual rentals relative to all rentals, thus, we set response $Y = \text{casual}/\text{count} \in [0, 1]$ on an hourly grid over the entire observation period. These are $n = 17,379$ hours from 2011/01/01 until 2012/12/31, see line 1 of Listing 1. We note that $\text{count} \geq 1$ for all observations, which makes Y well-defined throughout the whole observation period. The goal is to predict this response variable Y based on available feature information \mathbf{x} which is provided on lines 3-13 of Listing 1. These are the `year`, `month` and `hour` of the observations Y . The `weekday` (with 0 for Sunday), `holiday` (yes/no for public holiday), `workingday` (yes/no, the former neither being a public holiday nor a weekend), `weather` (1,2 and 3 for clear, cloudy and rain/snow), temperature `temp`, the felt temperature `temp_feel`, `humidity` and `windspeed`. Note that all these features are continuous or binary, thus, we can directly use this feature encoding for regression modeling.

We illustrate this data. Figure 11 shows the observed responses $Y = \text{casual}/\text{count}$ over the entire observation period. In average the casual rentals make 17% of all rentals, and the empirical density of Y is strongly skewed.

In Figure 12 we provide the marginal observed responses for each label of all features. The top-left shows the average response for each calendar week from 2011/01/01 until 2012/12/31. This depicts a strong seasonal pattern of the casual rentals proportion. Moreover, daytime, weekdays, working days/holidays and weather conditions such as temperature is important information for predicting the proportion of casual rentals. Only wind speed does not seem to be very relevant. From the top-middle we also observe that the proportion of casual rentals slightly decreases over time which can be explained by increasing regular rental subscriptions from 2011 to 2012.

For many of the feature components it is clear that they are highly correlated. In Figure 13 we




Figure 11: (lhs) Histogram and (rhs) boxplot of (hourly) responses $Y = \text{casual}/\text{count} \in [0, 1]$ over the entire observation period; the orange line shows the empirical mean of 17%.

plot temperature, humidity and wind speed against calendar month (top row), daytime (middle row) and weather conditions (bottom row). These plots clearly show this dependence. Moreover, humidity is negatively correlated with wind speed and positively correlated with temperature (at least up to moderate temperatures).




Figure 12: Average response Y for each label of all features date (in weekly units), year, month, hour, weekday, holiday, workingday, weather, temp, temp_feel, humidity and windspeed.




Figure 13: Dependence between feature components: (top) temperature, humidity and wind speed against calendar month, (middle) temperature, humidity and wind speed against daytime, (bottom) temperature, humidity and wind speed against weather conditions.