

Deep Conditional Transformation Models

Philipp F. M. Baumann (✉)¹[0000-0001-8066-1615], Torsten Hothorn²[0000-0001-8301-0471], and David Rügamer³[0000-0002-8772-9202]

¹ KOF Swiss Economic Institute, ETH Zurich, Zurich, Switzerland
baumann@kof.ethz.ch

² Epidemiology, Biostatistics and Prevention Institute, University of Zurich, Zurich, Switzerland

torsten.hothorn@uzh.ch

³ Department of Statistics, LMU Munich, Munich, Germany
david.ruegamer@stat.uni-muenchen.de

Abstract. Learning the cumulative distribution function (CDF) of an outcome variable conditional on a set of features remains challenging, especially in high-dimensional settings. Conditional transformation models provide a semi-parametric approach that allows to model a large class of conditional CDFs without an explicit parametric distribution assumption and with only a few parameters. Existing estimation approaches within this class are, however, either limited in their complexity and applicability to unstructured data sources such as images or text, lack interpretability, or are restricted to certain types of outcomes. We close this gap by introducing the class of deep conditional transformation models which unifies existing approaches and allows to learn both interpretable (non-)linear model terms and more complex neural network predictors in one holistic framework. To this end we propose a novel network architecture, provide details on different model definitions and derive suitable constraints as well as network regularization terms. We demonstrate the efficacy of our approach through numerical experiments and applications.

Keywords: Transformation Models · Distributional Regression · Normalizing Flows · Deep Learning · Semi-Structured Regression.

1 Introduction

Recent discussions on the quantification of uncertainty have emphasized that a distinction between aleatoric and epistemic uncertainty is useful in classical machine learning [38,18]. Moreover, this distinction was also advocated in the deep learning literature [20,4]. While epistemic uncertainty describes the uncertainty of the model and can be accounted for in a Bayesian neural network, aleatoric uncertainty [12] can be captured by modeling an outcome probability distribution that has a stochastic dependence on features (i.e., conditional on features). Apart from non-parametric estimation procedures, four fundamental approaches in statistics exist that allow to model the stochastic dependence between features and the outcome distribution [14]. First, parametric models where additive functions of

the features describe the location, scale and shape (LSS) of the distribution [34] or where these features are used in heteroscedastic Bayesian additive regression tree ensembles [31]. Second, quantile regression models [23,28,1] that directly model the conditional quantiles with a linear or non-linear dependence on feature values. Third, distribution regression and transformation models [9,35,3,43,27] that have response-varying effects on the probit, logit or complementary log-log scale. Finally, hazard regression [25] which estimates a non-proportional hazard function conditional on feature values. Parallel to this, various approaches in machine learning and deep learning have been evolved to model the outcome distribution through input features. A prominent example is normalizing flows (see, e.g. [41,33,30,22]), used to learn a complex distribution of an outcome based on feature values. Normalizing flows start with a simple base distribution F_Z and transform F_Z to a more complex target distribution using a bijective transformation of the random variable coming from the base distribution. The most recent advances utilize monotonic polynomials [19,32] or splines [29,5,6] to learn this transformation. As pointed out recently by several authors [21,39], normalizing flows are conceptually very similar to transformation models. However, normalizing flows, in contrast to transformation models, usually combine a series of transformations with limited interpretability of the influence of single features on the distribution of the target variable. In this work, we instead focus on conditional transformation approaches that potentially yield as much flexibility as normalizing flows, but instead of defining a generative approach, strive to build interpretable regression models without too restrictive parametric distribution assumptions.

1.1 Transformation Models

The origin of transformation models (TM) can be traced back to [2] studying a parametric approach to transform the variable of interest Y prior to the model estimation in order to meet a certain distribution assumption of the model. Many prominent statistical models, such as the Cox proportional hazards model or the proportional odds model for ordered outcomes, can be understood as transformation models. Estimating transformation models using a neural network has been proposed by [39]. However, [39] only focus on a smaller subclass of transformation models, we call (linear) shift transformation models and on models that are not interpretable in nature. Recently, fully parameterized transformation models have been proposed [15,16] which employ likelihood-based learning to estimate the cumulative distribution function F_Y of Y via estimation of the corresponding transformation of Y . The main assumption of TM is that Y follows a known, log-concave error distribution F_Z after some monotonic transformation h . CTMs specify this transformation function conditional on a set of features \mathbf{x} :

$$\mathbb{P}(Y \leq y|\mathbf{x}) = F_{Y|\mathbf{x}}(y) = F_Z(h(y|\mathbf{x})). \quad (1)$$

The transformation function h can be decomposed as $h(y|\mathbf{x}) := h_1 + h_2$, where h_1 and h_2 can have different data dependencies as explained in the following.

When h_1 depends on y as well as \mathbf{x} , we call the CTM an *interacting CTM*. When h_1 depends on y only, we call the model a *shift CTM*, with shift term h_2 . When h_2 is omitted in an interacting CTM, we call the CTM a *distributional CTM*. In general, the bijective function $h(y|\mathbf{x})$ is unknown a priori and needs to be learned from the data. [16] study the likelihood of this transformation function and propose an estimator for the most likely transformation. [16] specify the transformation function through a flexible basis function approach, which, in the unconditional case $h(y)$ (without feature dependency), is given by $h(y) = \mathbf{a}(y)^\top \boldsymbol{\vartheta}$ where $\mathbf{a}(y)$ is a matrix of evaluated basis functions and $\boldsymbol{\vartheta}$ a vector of basis coefficients which can be estimated by maximum likelihood. For continuous Y Bernstein polynomials [8] with higher order M provide a more flexible but still computationally attractive choice for \mathbf{a} . That is,

$$\mathbf{a}(y)^\top \boldsymbol{\vartheta} = \frac{1}{(M+1)} \sum_{m=0}^M \vartheta_m f_{Be(m+1, M-m+1)}(\tilde{y}) \quad (2)$$

where $f_{Be(m, M)}$ is the probability density function of a Beta distribution with parameters m, M and a normalized outcome $\tilde{y} := \frac{y-l}{u-l} \in [0, 1]$ with $u > l$ and $u, l \in \mathbb{R}$. In order to guarantee monotonicity of the estimate of $F_{Y|\mathbf{x}}$, strict monotonicity of $\mathbf{a}(y)^\top \boldsymbol{\vartheta}$ is required. This can be achieved by restricting $\vartheta_m > \vartheta_{m-1}$ for $m = 1, \dots, M+1$. When choosing M , the interplay with F_Z should be considered. For example, when $F_Z = \Phi$, the standard Gaussian distribution function, and $M = 1$, then \hat{F}_Y will also belong to the family of Gaussian distributions functions. Further, when we choose $M = n - 1$ with n being the number of independent observations, then \hat{F}_Y is the non-parametric maximum likelihood estimator which converges to F_Y by the Glivenko-Cantelli lemma [13]. As a result, for small M the choice of F_Z will be decisive, while TMs will approximate the empirical cumulative distribution function well when M is large independent of the choice of F_Z . Different choices for F_Z have been considered in the literature (see, e.g., [16]), such as the standard Gaussian distribution function (Φ), the standard logistic distribution function (F_L) or the minimum extreme value distribution function (F_{MEV}).

In CTMs with structured additive predictors (STAP), features considered in h_1 and h_2 enter through various functional forms and are combined as an additive sum. The STAP is given by

$$\eta_{struc} = s_1(\mathbf{x}) + \dots + s_k(\mathbf{x}) \quad (3)$$

with s_1, \dots, s_k being partial effects of one or more features in \mathbf{x} . Common choices include linear effects $\mathbf{x}^\top \mathbf{w}$ with regression coefficient \mathbf{w} and non-linear effects based on spline basis representation, spatial effects, varying coefficients, linear and non-linear interaction effects or individual-specific random effects [7]. Structured additive models have been proposed in many forms, for example in additive (mixed) models where $\mathbb{E}(Y|\mathbf{x}) = \eta_{struc}$.

1.2 Related Work and Our Contribution

The most recent advances in transformation models [17,21,14] learn the transformation functions h_1 and h_2 separately, using, e.g., a model-based boosting algorithm with pre-specified base learners [14]. Very recent neural network-based approaches allow for the joint estimation of both transformation functions, but do either not yield interpretable models [39] or are restricted to STAP with ordinal outcomes [24].

Our framework combines the existing frameworks and thereby extends approaches for continuous outcomes to transformation models able to 1) learn more flexible and complex specifications of h_1 and h_2 simultaneously 2) learn the CDF without the necessity of specifying the (type of) feature contribution a priori, 3) retain the interpretability of the structured additive predictor in h_1 and h_2 4) estimate structured effects in high-dimensional settings due to the specification of the model class within a neural network 5) incorporate unstructured data source such as texts or images.

2 Model and Network Definition

We now formally introduce the deep conditional transformation model (DCTM), explain its network architecture and provide details about different model definitions, penalization and model tuning.

2.1 Model Definition

Following [14], we do not make any explicit parameterized distribution assumption about Y , but instead assume

$$\mathbb{P}(Y \leq y|\mathbf{x}) = F_Z(h(y|\mathbf{x})) \quad (4)$$

with error distribution $F_Z : \mathbb{R} \mapsto [0, 1]$, an a priori known CDF that represents the data generating process of the transformed outcome $h(Y|\mathbf{x})$ conditional on some features $\mathbf{x} \in \chi$. For tabular data, we assume $\mathbf{x} \in \mathbb{R}^p$. For unstructured data sources such as images, \mathbf{x} may also include multidimensional inputs. Let f_Z further be the corresponding probability density function of F_Z . We model this transformation function conditional on some predictors \mathbf{x} by $h(y|\mathbf{x}) = h_1 + h_2 = \mathbf{a}(y)^\top \boldsymbol{\vartheta}(\mathbf{x}) + \beta(\mathbf{x})$, where $\mathbf{a}(y)$ is a (pre-defined) basis function $\mathbf{a} : \mathcal{E} \mapsto \mathbb{R}^{M+1}$ with \mathcal{E} the sample space and $\boldsymbol{\vartheta} : \chi_\vartheta \mapsto \mathbb{R}^{M+1}$ a conditional parameter function defined on $\chi_\vartheta \subseteq \chi$. $\boldsymbol{\vartheta}$ is parameterized through structured predictors such as splines, unstructured predictors such as a deep neural network, or the combination of both and $\beta(\mathbf{x})$ is a feature dependent distribution shift. More specifically, we model $\boldsymbol{\vartheta}(\mathbf{x})$ by the following additive predictor:

$$\boldsymbol{\vartheta}(\mathbf{x}) = \sum_{j=1}^J \boldsymbol{\Gamma}_j \mathbf{b}_j(\mathbf{x}), \quad (5)$$

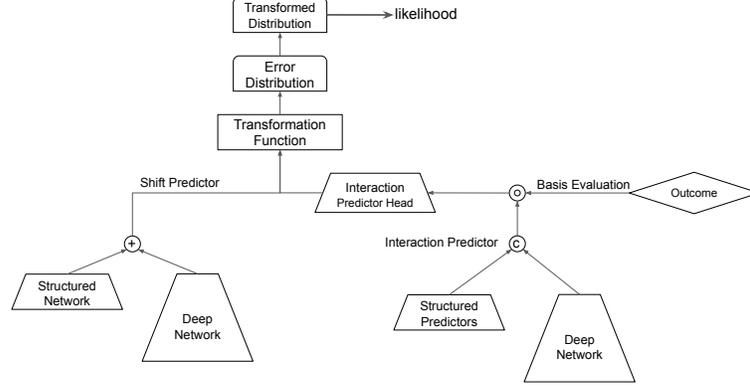


Fig. 1. Architecture of a deep conditional transformation model. Both the shift and interaction predictor can potentially be defined by a structured network including linear terms, (penalized) splines or other structured additive regression terms and deep neural network defined by an arbitrary network structure. While the shift predictor ($\mathcal{C}\Psi$) is a sum of both subnetwork predictions, the interaction predictor ($\mathcal{A} \odot \mathcal{B}$) is only multiplied with a final 1-hidden unit fully-connected layer (network head, $\text{vec}(\Gamma)$) after the structured predictors and latent features of the deep neural network are combined with the basis evaluated outcome. The shift and interaction network part together define the transformation function, which transforms the error distribution and yields the final likelihood used as loss function.

with $\Gamma_j \in \mathbb{R}^{(M+1) \times O_j}$, $O_j \geq 1$, being joint coefficient matrices for the basis functions in \mathbf{a} and the chosen predictor terms $\mathbf{b}_j : \chi_{b_j} \mapsto \mathbb{R}^{O_j}$, $\chi_{b_j} \subseteq \chi$. We allow for various predictor terms including an intercept (or bias term), linear effects $\mathbf{b}_j(\mathbf{x}) = x_{k_j}$ for some $k_j \in \{1, \dots, p\}$, structured non-linear terms $\mathbf{b}_j(\mathbf{x}) = G(x_{k_j})$ with some basis function $G : \mathbb{R} \mapsto \mathbb{R}^q$, $q \geq 1$ such as a B-spline basis, bivariate non-linear terms $\mathbf{b}_j(\mathbf{x}) = G'(x_{k_j}, x_{k_{j'}})$ using a tensor-product basis $G' : \mathbb{R} \times \mathbb{R} \mapsto \mathbb{R}^{q'}$, $q' \geq 1$ or neural network predictors $\mathbf{b}_j(\mathbf{x}) = d(\mathbf{x}_{k_j})$, which define an arbitrary (deep) neural network that takes (potentially multidimensional) features $\mathbf{x}_{k_j} \in \chi$. The network will be used to learn latent features representing the unstructured data source. These features are then combined as a linear combination when multiplied with Γ_j . The same types of predictors can also be defined for the shift term $\beta(\mathbf{x}) = \sum_{j=1}^J \mathbf{c}_j(\mathbf{x})^\top \psi_j$, which we also defined as an additive predictor of features, basis functions or deep neural networks times their (final) weighting ψ_j .

The final model output for the transformation of y is then given by

$$\mathbf{a}(y)^\top \boldsymbol{\vartheta}(\mathbf{x}) = \mathbf{a}(y)^\top \Gamma \mathbf{B}, \quad (6)$$

with $\Gamma = (\Gamma_1, \dots, \Gamma_J) \in \mathbb{R}^{(M+1) \times P}$, $P = \sum_{j=1}^J O_j$ the stacked coefficient matrix combining all Γ_j s and $\mathbf{B} \in \mathbb{R}^P$ a stacked vector of the predictor terms $\mathbf{b}_j(\mathbf{x})$ s. Based on model assumption (4) we can define the loss function based on the

change of variable theorem

$$f_Y(y|\mathbf{x}) = f_Z(h(y|\mathbf{x})) \cdot \left| \frac{\partial h(y|\mathbf{x})}{\partial y} \right|$$

as

$$\begin{aligned} \ell(h(y|\mathbf{x})) &= -\log f_Y(y|\boldsymbol{\vartheta}(\mathbf{x}), \boldsymbol{\beta}(\mathbf{x})) \\ &= -\log f_Z(\mathbf{a}(y)^\top \boldsymbol{\vartheta}(\mathbf{x}) + \beta(\mathbf{x})) - \log[\mathbf{a}'(y)^\top \boldsymbol{\vartheta}(\mathbf{x})] \end{aligned} \quad (7)$$

with $\mathbf{a}'(y) = \partial \mathbf{a}(y) / \partial y$.

For n observations $(y_i, \mathbf{x}_i), i = 1, \dots, n$, we can represent (6) as

$$(\mathcal{A} \odot \mathcal{B}) \text{vec}(\boldsymbol{\Gamma}^\top) \quad (8)$$

with $\mathcal{A} = (\mathbf{a}(y_1), \dots, \mathbf{a}(y_n))^\top \in \mathbb{R}^{n \times (M+1)}$, $\mathcal{B} = (\mathbf{B}_1, \dots, \mathbf{B}_n)^\top \in \mathbb{R}^{n \times P}$, vectorization operator $\text{vec}(\cdot)$ and the row-wise tensor product (also known as transpose Kathri-Rao product) operator \odot . Similar, the distribution shift can be written in matrix form as $\mathcal{C}\boldsymbol{\Psi}$ with $\mathcal{C} \in \mathbb{R}^{n \times Q}$ consisting of the stacked $\mathbf{c}_j(\mathbf{x})$ s and $\boldsymbol{\Psi} = (\boldsymbol{\psi}_1^\top, \dots, \boldsymbol{\psi}_J^\top)^\top \in \mathbb{R}^Q$ the stacked vector of all shift term coefficients. A schematic representation of an exemplary DCTM is given in Figure 2.

2.2 Network Definition

Our network consists of two main parts: a feature transforming network (FTN) part, converting $\mathbf{X} = (\mathbf{x}_1^\top, \dots, \mathbf{x}_n^\top)^\top \in \mathbb{R}^{n \times p}$ to \mathcal{B} and an outcome transforming network (OTN) part, transforming $\mathbf{y} = (y_1, \dots, y_n)^\top \in \mathbb{R}^n$ to $h(\mathbf{y}|\mathbf{X}) \in \mathbb{R}^n$. In the OTN part the matrix $\boldsymbol{\Gamma}$ is learned, while the FTN part only contains additional parameters to be learned by the network if some feature(s) are defined using a deep neural network. In other words, if only structured linear effects or basis function transformations are used in the FTN part, $\boldsymbol{\Gamma}$ contains all trainable parameters. Figure 1 visualizes an exemplary architecture.

After the features are processed in the FTN part, the final transformed outcome is modeled using a conventional fully-connected layer with input $\mathcal{A} \odot \mathcal{B}$, one hidden unit with linear activation function and weights corresponding to $\text{vec}(\boldsymbol{\Gamma})$. The deep conditional transformation model as visualized in Figure 1 can also be defined with one common network which is split into one part that is added to the shift predictor and one part that is used in the interaction predictor.

2.3 Penalization

L_1 -, L_2 -penalties can be incorporated in both the FTN and OTN part by adding corresponding penalty terms to the loss function. We further use smoothing penalties for structured non-linear terms by regularizing the respective entries in $\boldsymbol{\Psi}$ and $\boldsymbol{\Gamma}$ to avoid overfitting and easier interpretation. Having two smoothing directions, the penalty for $\boldsymbol{\Gamma}$ is constructed using a Kronecker sum of individual marginal penalties for anisotropic smoothing

$$\mathbf{D}_\Gamma = \lambda_a \mathbf{D}_a \oplus \lambda_b \mathbf{D}_b,$$

where the involved tuning parameters λ_a, λ_b and penalty matrices $\mathbf{D}_a, \mathbf{D}_b$ correspond to the direction of y and the features \mathbf{x} , respectively. Note, however, that for $\mathbf{\Gamma}$, the direction of y usually does not require additional smoothing as it is already regularized through the monotonicity constraint [16]. The corresponding penalty therefore reduces to

$$\mathbf{D}_{\Gamma} = \mathbf{I}_P \otimes (\lambda_b \mathbf{D}_b) \quad (9)$$

with the diagonal matrix \mathbf{I}_P of size P . These penalties are added to the negative log-likelihood defined by (7), e.g.,

$$\ell_{pen} = \ell(h(y|\mathbf{x})) + \text{vec}(\mathbf{\Gamma})^\top \mathbf{D}_{\Gamma} \text{vec}(\mathbf{\Gamma})$$

for a model with penalized structured effects only in \mathcal{B} . As done in [37] we use the Demmler-Reinsch orthogonalization to relate each tuning parameter for smoothing penalties to its respective degrees-of-freedom, which allows a more intuitive setting of parameters and, in particular, allows to define equal amount of penalization for different smooth terms. Leaving the least flexible smoothing term unpenalized and adjusting all others to have the same amount of flexibility works well in practice.

2.4 Bijectivity and Monotonicity Constraints

To ensure bijectivity of the transformation of each y_i , we use Bernstein polynomials for \mathcal{A} and constraint the coefficients in $\mathbf{\Gamma}$ to be monotonically increasing in each column. The monotonicity of the coefficients in $\mathbf{\Gamma}$ can be implemented in several ways, e.g., using the approach by [11] or [39] on a column-basis. Note that this constraint directly yields monotonically increasing transformation functions if $P = 1$, i.e., if no or only one feature is used for h_1 . If $P > 1$, we can ensure monotonicity of h_1 by using predictor terms in \mathcal{B} that are non-negative. A corresponding proof can be found in the Supplementary Material (Lemma 1)⁴. Intuitively the restriction can be seen as an implicit positivity assumption on the learned standard deviation of the error distribution F_Z as described in the next section using the example of a normal distribution. Although non-negativity of predictor terms is not very restrictive, e.g., allowing for positive linear features, basis functions with positive domain such as B-splines or deep networks with positivity in the learned latent features (e.g., based on a ReLU activation function), the restriction can be lifted completely by simply adding a positive constant to \mathcal{B} .

2.5 Interpretability and Identifiability Constraints

Several choices for M and F_Z will allow for particular interpretation of the coefficients learned in $\mathbf{\Psi}$ and $\mathbf{\Gamma}$. When choosing $F_Z = \Phi$ and $M = 1$, the DCTM effectively learns an additive regression model with Gaussian error distribution,

⁴ <https://github.com/PFMB/DCTMs>

i.e., $Y|\mathbf{x} \sim N(\tilde{\beta}(\mathbf{x}), \sigma_s^2)$. The unstandardized structured additive effects in $\tilde{\beta}(\mathbf{x})$ can then be divided by σ_s yielding $\beta(\mathbf{x})$. Therefore $\beta(\mathbf{x})$ can be interpreted as shifting effects of normalized features on the transformed response $\mathbb{E}(h_1(y)|\mathbf{x})$. For $M > 1$, features in $\beta(\mathbf{x})$ will also affect higher moments of $Y|\mathbf{x}$ through a non-linear h_1 , leading to a far more flexible modeling of $F_{Y|\mathbf{x}}$. Smooth monotonously increasing estimates for $\beta(\mathbf{x})$ then allow to infer that a rising \mathbf{x} leads to rising moments of $Y|\mathbf{x}$ independent of the choice for F_Z . Choosing $F_Z = F_{MEV}$ or $F_Z = F_L$ allows $\beta(\mathbf{x})$ to be interpreted as additive changes on the log-hazard ratio or on the log-odds ratio, respectively. The weights in \mathbf{T} determine the effect of \mathbf{x} on $F_{Y|\mathbf{x}}$ as well as whether $F_{Y|\mathbf{x}}$ varies with the values of y yielding a response-varying distribution [3] or not. In general, structured effects in \mathbf{T}

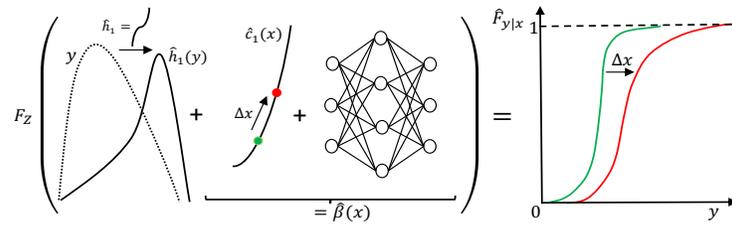


Fig. 2. Schematic representation of an exemplary DCTM with a learned transformation \hat{h}_1 for the outcome y . The shift term $\hat{\beta}(x)$ is composed of an estimated smooth term $\hat{c}_1(x) = c_1(x)\hat{\psi}_1$ for x and a neural network predictor. An increase in x is indicated by Δx with corresponding effect on $\hat{F}_{Y|\mathbf{x}}$ through $\hat{h}_2 = \hat{\beta}(x)$ on the right hand side of the equation.

are coefficients of the tensor product $\mathcal{A} \odot \mathcal{B}$ and can, e.g., be interpreted by 2-dimensional contour or surface plots (see, e.g., Figure 1 in the Supplementary Material).

In order to ensure identifiability and thus interpretability of structured effects in h_1 and h_2 , several model definitions require the additional specifications of constraints. If certain features in \mathcal{B} or \mathcal{C} are modeled by both a flexible neural network predictor $d(\mathbf{x})$ and structured effects $s(\mathbf{x})$, the subnetwork $d(\mathbf{x})$ can easily assimilate effects $s(\mathbf{x})$ is supposed to model. In this case, identifiability can be ensured by an orthogonalization cell [37], projecting the learned effects of $d(\mathbf{x})$ in the orthogonal complement of the space spanned by features modeled in $s(\mathbf{x})$. Further, when more than one smooth effect or deep neural network is incorporated in either \mathcal{B} or \mathcal{C} , these terms can only be learned up to an additive constants. To solve this identifiability issue we re-parameterize the terms and learn these effects with a sum-to-zero constraint. As a result, corresponding effects can only be interpreted on a relative scale. Note that this is a limitation of additive models per se, not our framework.

3 Numerical Experiments

We now demonstrate the efficacy of our proposed framework for the case of a shift CTM, a distributional CTM and an interacting CTM based on a general data generating process (DGP).

Data Generating Process The data for the numerical experiments were generated according to $g(y) = \eta(\mathbf{x}) + \epsilon(\mathbf{x})$ where $g : \mathbb{R}^n \mapsto \mathbb{R}^n$ is bijective and differentiable, $\eta(\mathbf{x})$ is specified as in (3) and $\epsilon \sim F_Z$ with F_Z being the error distribution. We choose $\epsilon(\mathbf{x}) \sim N(0, \sigma^2(\mathbf{x}))$ where $\sigma^2(\mathbf{x}) \in \mathbb{R}^+$ is specified as in (3) so that we can rewrite the model as

$$F_Z \left(\frac{g(y) - \eta(\mathbf{x})}{\sigma(\mathbf{x})} \right) = F_Z (h_1 + h_2). \quad (10)$$

From (1) and our model definition, (10) can be derived by defining h_1 as $g(y)\sigma^{-1}(\mathbf{x})$ and h_2 as $-\eta(\mathbf{x})\sigma^{-1}(\mathbf{x})$. We finally generate y according to $g^{-1}(\eta(\mathbf{x}) + \epsilon(\mathbf{x}))$ with $\epsilon(\mathbf{x}) \sim N(0, \sigma^2(\mathbf{x}))$. We consider different specification for g , η , σ and the order of the Bernstein polynomial M for different samples sizes n (see Supplementary Material).

Evaluation To assess the estimation performance, we compute the relative integrated mean squared error (RIMSE) between \hat{h}_1 , evaluated on a fine grid of $(y_i, \sigma(\mathbf{x}_i))$ value pairs, with the true functional form of h_1 as defined by the data generating process. For the estimation performance of h_2 , we evaluate the corresponding additive predictor by calculating the mean squared error (MSE) between estimated and true linear coefficients for linear feature effects and the RIMSE between estimated and true smooth non-linear functions for non-linear functional effects. We compare the estimation against transformation boosting machines (TBM) [14] that also allow to specify structured additive predictors. Note, however, that TBMs only implement either the shift (TBM-Shift) or distributional CTM (TBM-Distribution), but do not allow for the specification of an interacting CTM with structured predictors, a novelty of our approach. In particular, only the TBM-Shift comes with an appropriate model specification such that it can be used for comparison in some of the DGP defined here.

Results We first discuss the 4 out of 10 specifications of the true DGP where h_1 is not learned through features and thus allows for a direct comparison of TBM-Shift and DCTMs. For h_1 , we find that, independent of the size of the data set and the order of the Bernstein polynomial, DCTMs provide a viable alternative to TBM-Shift, given the overlap between the (RI)MSE distributions and the fact that the structured effects in DCTMs are not tuned extensively in these comparisons. For h_2 , DCTMs outperform TBM-Shift in all 16 configurations for M/n among the 4 DGPs depicted in Figure 3 when taking the mean or the median across the 20 replications. The simulation results for the 6 remaining DGPs can be found in the Supplementary Material. For h_1 and h_2 , the results for the majority of specifications reveal that DCTMs benefit from lower order Bernstein polynomials

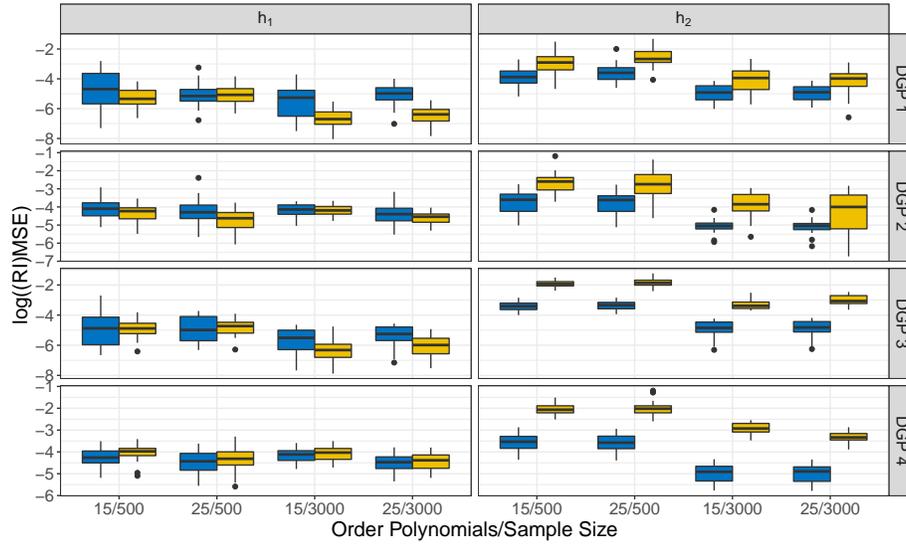


Fig. 3. Comparison of the logarithmic (RI)MSEs between TBM-Shift (yellow) and DCTM (blue) for different data generating processes (DGP in rows) as well as different orders of the Bernstein polynomial and the sample size (M/n on the x-axis) for 20 runs. The specification of the DGPs can be found in the Supplementary Material and for this figure are based on σ_1 with alternating $g \in \{g_1, g_2\}$ and $\eta \in \{\eta_1, \eta_2\}$.

independent of the sample size. When only unstructured model components were specified, DCTM’s estimation of h_1 benefits from Bernstein polynomial with higher order. This holds regardless of g . Figure 1 in the Supplementary Material exemplary depicts the estimation performance of DCTMs for one DGP setting.

4 Application

We now demonstrate the application of DCTMs by applying the approach to a movie reviews and a face data set.

4.1 Movie Reviews

The Kaggle movies data set consists of $n = 4442$ observations. Our goal is to predict the movies’ revenue based on their production budget, popularity, release date, runtime and genre(s). Figure 1 in the Supplementary Material depicts the revenue for different genres. We deliberately do not log-transform the response, but let the transformation network convert a standard normal distribution (our error distribution) to fit to the given data.

Model Description First, we define a DCTM solely based on a structured additive predictor (i.e. no deep neural net predictor) as a baseline model which we refer to

as the “Structured Model”. The structured additive predictor includes the binary effect for each element of a set of 20 available genres (x_0) as well as smooth effects (encoded as a univariate thin-plate regression splines [42]) for the popularity score (x_1), for the difference of release date and a chosen date in the future in days (x_2), for the production budget in US dollars (x_3) and the run time in days (x_4):

$$\sum_{r=1}^{20} \beta_r I(r \in x_{0,i}) + s_1(x_{1,i}) + s_2(x_{2,i}) + s_3(x_{3,i}) + s_4(x_{4,i}). \quad (11)$$

This linear predictor (11) is used to define the structured component in the shift term $\beta(\mathbf{x})$. For the interaction term, the STAP consists of all the genre effects and the resulting design matrix \mathbf{B} is then combined with the basis of a Bernstein polynomial \mathcal{A} of order $M = 25$. We compare this model with three deep conditional transformation models that use additional textual information of each movie by defining a “Deep Shift Model”, a “Deep Interaction Model” and a “Deep Combination Model”. The three models all include a deep neural network as input either in the shift term, in the interaction term or as input for both model parts, respectively. As deep neural network we use an embedding layer of dimension 300 for 10000 unique words and combine the learned outputs by flattening the representations and adding a fully-connected layer with 1 unit for the shift term and/or 1 units for the interaction term on top. As base distribution we use a logistic distribution, i.e., $F_Z(h) = F_L(h) = (1 + \exp(-h))^{-1}$.

Comparisons We use 20% of the training data as validation for early stopping and define the degrees-of-freedom for all non-linear structured effects using the strategy described in Section 2.3. We compare our approach again with the shift and distributional TBM (TBM-Shift and TBM-Distribution, respectively) as state-of-the-art baseline. We run both models with the predictor specification given in (11). For TBM, we employ a 10-fold bootstrap to find the optimal stopping iteration by choosing the minimum out-of-sample risks averaged over all folds. Finally we evaluate the performance on the test data for both algorithms.

Results The non-linear estimations of all models show a similar trend for the four structured predictors. Figure 4 depicts an example for the estimated partial effects in the h_2 term of each model. The resulting effects in Figure 4 can be interpreted as functional log-odds ratios due to the choice $F_Z = F_L$. For example, the log-odds for higher revenue linearly increase before the effect stagnates for three of the four model at a level greater than 150 million USD. Table 4.1 shows (Movie Reviews column) the mean predicted log-scores [10], i.e., the average log-likelihood of the estimated distribution of each model when trained on 80% of the data (with 20% of the training data used as validation data) and evaluated on the remaining 20% test data. Results suggest that deep extensions with movie descriptions as additional predictor added to the baseline model can improve over the TBM, but do not achieve as good prediction results as the purely structured DCTM model in this case. Given the small amount of data, this result is not surprising and showcases a scenario, where the potential of the structured model part outweighs the information of a non-tabular data source. The flexibility of our

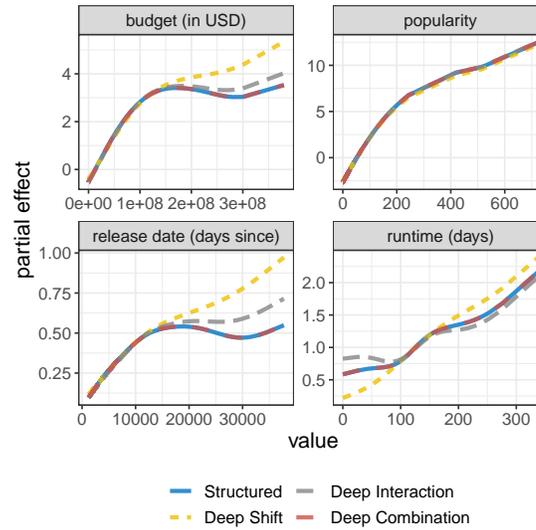


Fig. 4. Estimated non-linear partial effect of the 4 available numerical features for h_2 (in each sub-plot) based on the four different DCTM models (colors).

approach in this case allows to seamlessly traverse different model complexities and offers a trade-off between complexity and interpretability.

Table 1. Average result (standard deviation in brackets) over different training/test-splits on the movie reviews (left) and UTKFace data set. Values correspond to negative predicted log-scores (PLS; smaller is better) for each model with best score in bold.

	Model	Movie Reviews	UTKFace
DCTM	Structured	19.26 (0.18)	3.98 (0.02)
	Deep Shift	19.32 (0.20)	3.81 (0.52)
	Deep Interaction	19.69 (0.22)	3.79 (0.21)
	Deep Combination	19.67 (0.19)	3.37 (0.09)
TBM	Shift	23.31 (0.83)	4.25 (0.02)
	Distribution	22.38 (0.31)	4.28 (0.03)

4.2 UTKFace

The UTKFace dataset is a publicly available image dataset with $n = 23708$ images and additional tabular features (age, gender, ethnicity and collection date). We use this data set to investigate DCTMs in a multimodal data setting.

Model Description Our goal is to learn the age of people depicted in the images using both, the cropped images and the four tabular features. As in the previous

section we fit the four different DCTM models, all with the same structured additive predictor (here effect for race, gender and a smooth effect for the collection date) and add a deep neural network predictor to the h_1 (Deep Interaction), h_2 (Deep Shift), to both (Deep Combination) or only fit the structured model without any information of the faces (Structured). The architecture for the faces consists of three CNN blocks (see Supplementary Material for details) followed by flattening operation, a fully-connected layer with 128 units with ReLU activation, batch normalization and a dropout rate of 0.5. Depending on the model, the final layer either consists of 1 hidden unit (Deep Shift, Deep Interaction) or 2 hidden units (Deep Combination).

Comparisons The baseline model is a two-stage approach that first extracts latent features from the images using a pre-trained VGG-16 [40] and then uses these features together with the original tabular features in a TBM-Shift/-Distribution model to fit a classical structured additive transformation model. We again compare the 4 DCTM models and 2 baseline models using the PLS on 30% test data and report model uncertainties by repeating the data splitting and model fitting 4 times. For the DCTMs we use early stopping based on 20% of the train set used for validation. For TBM models we search for the best stopping iteration using a 3-fold cross-validation. The results in Table 4.1 (UTKFace column) suggest that our end-to-end approach works better than the baseline approach and that the DCTM benefits from a combined learning of h_1 and h_2 through the images.

4.3 Benchmark Study

We finally investigate the performance of our approach by comparing its density estimation on four UCI benchmark data sets (Airfoil, Boston, Diabetes, Forest Fire) against parametric alternatives. We use a deep distributional regression approach (DR) [37], a Gaussian process (GP) and a GP calibrated with an isotonic regression (IR) [26]. We adapt the same architecture as in DR to specifically examine the effect of the proposed transformation. To further investigate the impact of the polynomials' order M (i.e., flexibility of the transformation vs. risk of overfitting), we run the DCTM model with $M \in \{1, 16, 32, 64\}$ (DCTM- M). We also include a normalizing flow baseline with a varying number of radial flows \mathcal{M} (NF- \mathcal{M} ; [36]). This serves as a reference for a model with more than one transformation and thus potentially more expressiveness at the expense of the feature-outcome relationship being not interpretable. Details for hyperparameter specification can be found in the Supplementary Material. Results (Table 2) indicate that our approach performs similar to alternative methods. For two data sets, the greater flexibility of the transformation yields superior performance compared to methods without transformation (DR, GP, IR), suggesting that the transition from a pure parametric approach to a more flexible transformation model can be beneficial. For the other two data sets, DCTM's performance is one standard deviation apart from the best performing model. For Airfoil the

Table 2. Comparison of neg. PLS (with standard deviation in brackets) of different methods (rows; best-performing model in bold, second best underlined) on four different UCI repository datasets (columns) based on 20 different initializations of the algorithms.

	Airfoil	Boston	Diabetes	Forest F.
DR	3.11 (0.02)	3.07 (0.11)	5.33 (0.00)	1.75 (0.01)
GP	3.17 (6.82)	2.79 (2.05)	5.35 (5.76)	1.75 (7.09)
IR	3.29 (1.86)	3.36 (5.19)	5.71 (2.97)	1.00 (1.94)
DCTM-1	3.07 (0.01)	2.97 (0.03)	5.44 (0.02)	1.83 (0.02)
DCTM-16	3.07 (0.02)	2.76 (0.02)	<u>5.34</u> (0.01)	1.30 (0.12)
DCTM-32	3.08 (0.02)	<u>2.71</u> (0.03)	5.39 (0.02)	<u>1.08</u> (0.15)
DCTM-64	3.08 (0.03)	2.66 (0.05)	5.37 (0.01)	1.41 (1.03)
NF-1	3.04 (0.22)	2.98 (0.20)	5.59 (0.10)	1.77 (0.02)
NF-3	2.88 (0.17)	2.76 (0.14)	5.54 (0.11)	1.76 (0.02)
NF-5	<u>2.90</u> (0.18)	2.81 (0.20)	5.47 (0.09)	1.77 (0.12)

even greater flexibility of a chain of transformations (NF- \mathcal{M} in comparison to DCTM- M) improves upon the result of DCTMs.

5 Conclusion and Outlook

We introduced the class of deep conditional transformation models which unifies existing fitting approaches for transformation models with both interpretable (non-)linear model terms and more complex predictors in one holistic neural network. A novel network architecture together with suitable constraints and network regularization terms is introduced to implement our model class. Numerical experiments and applications demonstrate the efficacy and competitiveness of our approach.

Acknowledgements

This work has been partly funded by SNF grant 200021-184603 from the Swiss National Science Foundation (Torsten Hothorn) and the German Federal Ministry of Education and Research (BMBF) under Grant No. 01IS18036A (David Rügamer).

References

1. Athey, S., Tibshirani, J., Wager, S., et al.: Generalized random forests. *The Annals of Statistics* **47**(2), 1148–1178 (2019)
2. Box, G.E., Cox, D.R.: An analysis of transformations. *Journal of the Royal Statistical Society: Series B (Methodological)* **26**(2), 211–243 (1964)
3. Chernozhukov, V., Fernández-Val, I., Melly, B.: Inference on counterfactual distributions. *Econometrica* **81**(6), 2205–2268 (2013)

4. Depeweg, S., Hernandez-Lobato, J.M., Doshi-Velez, F., Udluft, S.: Decomposition of uncertainty in bayesian deep learning for efficient and risk-sensitive learning. In: International Conference on Machine Learning. pp. 1184–1193. PMLR (2018)
5. Durkan, C., Bekasov, A., Murray, I., Papamakarios, G.: Cubic-spline flows. arXiv preprint arXiv:1906.02145 (2019)
6. Durkan, C., Bekasov, A., Murray, I., Papamakarios, G.: Neural spline flows. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 32. Curran Associates, Inc. (2019)
7. Fahrmeir, L., Kneib, T., Lang, S., Marx, B.: Regression: Models, Methods and Applications. Springer Berlin Heidelberg (2013)
8. Farouki, R.T.: The Bernstein polynomial basis: A centennial retrospective. Computer Aided Geometric Design **29**(6), 379–419 (2012)
9. Foresi, S., Peracchi, F.: The conditional distribution of excess returns: An empirical analysis. Journal of the American Statistical Association **90**(430), 451–466 (1995)
10. Gelfand, A.E., Dey, D.K.: Bayesian model choice: asymptotics and exact calculations. Journal of the Royal Statistical Society: Series B (Methodological) **56**(3), 501–514 (1994)
11. Gupta, M., Cotter, A., Pfeifer, J., Voevodski, K., Canini, K., Mangylov, A., Moczydlowski, W., van Esbroeck, A.: Monotonic calibrated interpolated look-up tables. Journal of Machine Learning Research **17**(109), 1–47 (2016)
12. Hora, S.C.: Aleatory and epistemic uncertainty in probability elicitation with an example from hazardous waste management. Reliability Engineering & System Safety **54**(2-3), 217–223 (1996)
13. Hothorn, T.: Most likely transformations: The mlt package. Journal of Statistical Software, Articles **92**(1), 1–68 (2020)
14. Hothorn, T.: Transformation boosting machines. Statistics and Computing **30**(1), 141–152 (2020)
15. Hothorn, T., Kneib, T., Bühlmann, P.: Conditional transformation models. Journal of the Royal Statistical Society: Series B: Statistical Methodology pp. 3–27 (2014)
16. Hothorn, T., Möst, L., Bühlmann, P.: Most likely transformations. Scandinavian Journal of Statistics **45**(1), 110–134 (2018)
17. Hothorn, T., Zeileis, A.: Predictive distribution modeling using transformation forests. Journal of Computational and Graphical Statistics pp. 1–16 (2021)
18. Hüllermeier, E., Waegeman, W.: Aleatoric and epistemic uncertainty in machine learning: A tutorial introduction. arXiv preprint arXiv:1910.09457 (2019)
19. Jaini, P., Selby, K.A., Yu, Y.: Sum-of-squares polynomial flow. CoRR (2019)
20. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? In: Advances in neural information processing systems. pp. 5574–5584 (2017)
21. Klein, N., Hothorn, T., Kneib, T.: Multivariate conditional transformation models. arXiv preprint arXiv:1906.03151 (2019)
22. Kobyzev, I., Prince, S., Brubaker, M.: Normalizing flows: An introduction and review of current methods. IEEE Transactions on Pattern Analysis and Machine Intelligence p. 1–1 (2020). <https://doi.org/10.1109/tpami.2020.2992934>
23. Koenker, R.: Quantile Regression, vol. Economic Society Monographs. Cambridge University Press (2005)
24. Kook, L., Herzog, L., Hothorn, T., Dürr, O., Sick, B.: Ordinal neural network transformation models: Deep and interpretable regression models for ordinal outcomes. arXiv preprint arXiv:2010.08376 (2020)

25. Kooperberg, C., Stone, C.J., Truong, Y.K.: Hazard regression. *Journal of the American Statistical Association* **90**(429), 78–94 (1995)
26. Kuleshov, V., Fenner, N., Ermon, S.: Accurate uncertainties for deep learning using calibrated regression. *Proceedings of the 35th International Conference on Machine Learning* **80**, 2796–2804 (2018)
27. Leorato, S., Peracchi, F.: Comparing Distribution and Quantile Regression. EIEF Working Papers Series 1511, Einaudi Institute for Economics and Finance (EIEF) (2015)
28. Meinshausen, N.: Quantile regression forests. *Journal of Machine Learning Research* **7**(Jun), 983–999 (2006)
29. Müller, T., McWilliams, B., Rousselle, F., Gross, M., Novák, J.: Neural importance sampling (2019)
30. Papamakarios, G., Nalisnick, E., Rezende, D.J., Mohamed, S., Lakshminarayanan, B.: Normalizing flows for probabilistic modeling and inference (2019)
31. Pratola, M., Chipman, H., George, E.I., McCulloch, R.: Heteroscedastic BART via multiplicative regression trees. *Journal of Computational and Graphical Statistics* (2019)
32. Ramasinghe, S., Fernando, K., Khan, S., Barnes, N.: Robust normalizing flows using bernstein-type polynomials (2021)
33. Rezende, D., Mohamed, S.: Variational inference with normalizing flows. *Proceedings of Machine Learning Research* **37**, 1530–1538 (2015)
34. Rigby, R.A., Stasinopoulos, D.M.: Generalized additive models for location, scale and shape. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **54**(3), 507–554 (2005)
35. Rothe, C., Wied, D.: Misspecification testing in a class of conditional distributional models. *Journal of the American Statistical Association* **108**(501), 314–324 (2013)
36. Rothfuss, J., Ferreira, F., Boehm, S., Walther, S., Ulrich, M., Asfour, T., Krause, A.: Noise regularization for conditional density estimation (2020)
37. Rügamer, D., Kolb, C., Klein, N.: A Unified Network Architecture for Semi-Structured Deep Distributional Regression. *arXiv preprint arXiv:2002.05777* (2020)
38. Senge, R., Bösner, S., Dembczyński, K., Haasenritter, J., Hirsch, O., Donner-Banzhoff, N., Hüllermeier, E.: Reliable classification: Learning classifiers that distinguish aleatoric and epistemic uncertainty. *Information Sciences* **255**, 16–29 (2014)
39. Sick, B., Hothorn, T., Dürr, O.: Deep transformation models: Tackling complex regression problems with neural network based transformation models. *arXiv preprint arXiv:2004.00464* (2020)
40. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: Bengio, Y., LeCun, Y. (eds.) *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings* (2015)
41. Tabak, E.G., Turner, C.V.: A family of nonparametric density estimation algorithms. *Communications on Pure and Applied Mathematics* **66**(2), 145–164 (2013)
42. Wood, S.N.: Thin plate regression splines. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **65**(1), 95–114 (2003)
43. Wu, C.O., Tian, X.: Nonparametric estimation of conditional distributions and rank-tracking probabilities with time-varying transformation models in longitudinal studies. *Journal of the American Statistical Association* **108**(503), 971–982 (2013)