Multidimensional Shape Constraints

Maya R. Gupta [*1] **Erez Louidor** [*1] **Olexander Mangylov** [*1] **Nobuyuki Morioka** [*1] **Taman Narayan** [*1] **Sen Zhao** [*1]

Abstract

We propose new multi-input shape constraints across four intuitive categories: complements, diminishers, dominance, and unimodality constraints. We show these shape constraints can be checked and even enforced when training machine-learned models for linear models, generalized additive models, and the nonlinear function class of multi-layer lattice models. Real-world experiments illustrate how the different shape constraints can be used to increase explainability and improve regularization, especially for non-IID train-test distribution shift.

1. Introduction

Shape constraints are a classic way to characterize a function by whether its shape obeys certain properties (see e.g., Barlow et al. (1972); Groeneboom & Jongbloed (2014); Chetverikov et al. (2018)). The most popular shape constraint in machine learning is monotonicity (see, e.g., Archer & Wang (1993); Sill (1998); Howard & Jebara (2007); Minin et al. (2010); Gupta et al. (2016)). For example, suppose one is building a model to predict the price of a house. Then one might expect the price to be a monotonically increasing function with respect to square meters, holding fixed any value of the other inputs such as location.

Shape constraints are useful in machine learning because they aid interpretability: they provide ways to describe the function in terms of everyday relationships between its inputs and outputs that hold everywhere. Shape constraints are also semantically-meaningful regularizers that can improve generalization, especially when the distribution shifts between the train and test data (Canini et al., 2016; You et al., 2017). Shape constraints can also be imposed to ensure some notions of AI safety, societal norms and deontological ethics (Wang & Gupta, 2020).

[*]Equal contribution [1]Google Research, Mountain View, California, USA. Correspondence to: Sen Zhao <senzhao@google.com>.

Proceedings of the 37^{th} International Conference on Machine Learning, Vienna, Austria, PMLR 119, 2020.

Shape constraints historically have been defined in terms of how the function responds to changes in a single input, e.g. how house size affects its price, with modelers often applying many such constraints simultaneously on different inputs. Cotter et al. (2019a) extended this idea and proposed two new shape constraints that are defined on *pairs* of features that are complements with one another, such as a measurement and its precision. For example, a model that predicts whether a user will click on a web link might use as features a: the past click-through-rate (CTR) for that link, and b: the number of impressions that a was calculated from. Intuitively, the model output should be more sensitive to the past CTR a if the number of past impressions b used to calculate that CTR was higher, since the model can then trust the CTR to be more accurate. Cotter et al. (2019a) captured this type of feature interaction with two different mathematical formulations, which they termed *Edgeworth* and *trapezoid* shape constraints, and showed these two-feature shape constraints are broadly applicable whenever models include a pair of inputs where input a is a measurement correlated with the label y, and input b estimates how much one can trust input a.

Inspired by that work, here we address the broader question, "What multidimensional shape constraints can be specified that would be useful to machine-learning practitioners?" We identified four useful categories of multidimensional shape constraints: complements, diminishers, dominance, and unimodality. We review and propose new shape constraints spanning these four categories that we think are most valuable to practitioners, pictured in Fig. 1. We show that these shape constraints can be expressed as linear inequality constraints for some function classes including flexible lattice models, and thus can be efficiently checked for trained models. Moreover, we show that machine-learned models from these function classes can be *trained* to respect these shape constraints by minimizing empirical risk with the appropriate linear inequality constraints on the model parameters. We give intuitive examples of usage, and present experimental results on benchmark and real problems illustrating the use and effectiveness of multi-input shape constraints for interpretability and regularization of nonlinear models.

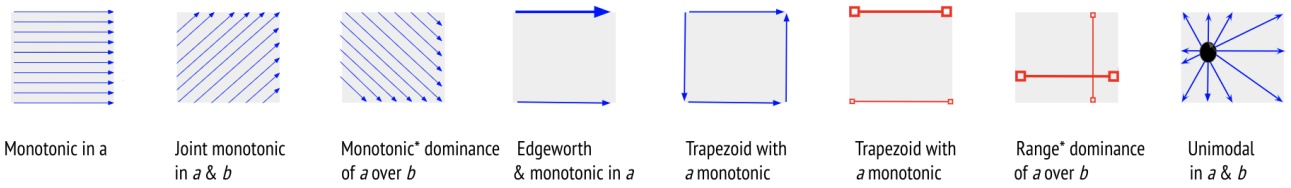

Figure 1. Illustration of prior and proposed shape constraints that hold for each a-b slice of the feature space, where a is the horizontal feature and b is the vertical feature. Arrows (in blue) denote the function can only increase in that direction. The * on monotonic dominance and range dominance is to remind readers there are additionally monotonicity constraints in a and b that are not shown; for example, $f(a, b) = 5a + 3b$ is monotonically dominant in a over b, but also increasing in a and b. A bigger arrow denotes a steeper increase than a smaller arrow. The red lines with open-box-ends denote the range of the outputs over the line, with bigger open-boxes denoting a larger output range. As shown, the trapezoid constraint imposes a *twist* shape that can equivalently be expressed in terms of arrows or ranges, where for trapezoid the larger output range must also be a superset of the smaller output ranges.

2. Shape Constraint Properties

We considered many shape constraint definitions during our investigation, but we propose only the ones we think will be useful, based on the following criteria.

Intuitive: Can a data scientist easily understand the meaning of the constraint and apply it appropriately to regularize a model? Can an end user easily understand the meaning of the constraint making it helpful in explaining a model?

Unit-sensitivity: Is the constraint sensitive or dependent on how we measure or count inputs a and b? Unit sensitivity can be a good thing if it enables us to capture more domain knowledge, but can be a bad thing if it makes the constraint too fragile or too difficult to specify.

Composability: Deep models are formed by composing layers. If a shape constraint holds for $f(x)$, what needs to be true about $g : \mathbb{R} \to \mathbb{R}$ for the constraint to hold for the composition $g(f(x))$? For example, for monotonicity it is sufficient (though not necessary) for each layer to be monotonic with respect to an input a for the multi-layer model to be monotonic in a, making it possible to construct monotonic deep models (You et al., 2017).

Verification and Training: To be useful for aiding in the explanation of a machine-learned model, it must be computationally tractable to check if a machine-learned model satisfies the shape constraint. This requires that the shape constraint can be efficiently expressed in terms of the parameters of the model. Further, we would like to be able to *train* models that satisfy shape constraints. Thus, it is helpful if the shape constraint can be expressed as a set of *linear* inequality constraints on the parameters, as is true for monotonicity and diminishing returns constraints for lattice functions (You et al., 2017; Gupta et al., 2018), because then training can be done using empirical risk minimization *subject to* the necessary linear inequalities.

No Distribution Dependence: Shape constraints describe the *shape* of a function and thus their definitions do not depend on data, in contrast to *data-dependent constraints* which are defined in terms of expectations on sample distributions (see e.g., Mann & McCallum (2007); Zafar et al. (2017); Cotter et al. (2019b)).

Consider the goal of a "dominance" shape constraint that feature a is more important than feature b. There are already many feature importance metrics, and most of these can be readily checked given a trained model to interpret if feature A is more important than feature B, but they are not expressible as shape constraints. For example, one could try to capture prior knowledge that the model f should be more sensitive to feature A than feature B by requiring that the variance of $f(A, \bar{B}, X)$ is bigger than the variance of $f(\bar{A}, B, X)$, where $\bar{A}, \bar{B}$ denote the expectations of those inputs, and the variance is taken over the joint distribution $P_{A,B,X}$. That might be a useful constraint, but whether it holds depends on $P_{A,B,X}$, and thus is not a shape constraint. Similarly, measuring feature importance by the effect on validation accuracy if you drop that feature is data-dependent. In addition, for most prior feature importance metrics, it is difficult to train a model that guarantees that a specified feature a is more important than a specified feature b. We restrict our attention to *shape constraints* in this paper.

3. Notation and Preliminaries

For $n \in \mathbb{N}$ we denote by $[n]$ the set $\{1, 2, \ldots, n\}$. For $\mathbf{x} \in \mathbb{R}^D$ let $\mathbf{x}[d]$ denote the dth entry of $\mathbf{x}$; if bounded we denote its bounds by $[\ell_d, u_d]$, and let $L_d(\mathbf{x})$ and $U_d(\mathbf{x})$ be vectors obtained from $\mathbf{x}$ by replacing its dth entry by ℓ_d or u_d, respectively. Let $\mathbf{e}_d \in [0, 1]^D$ denote the one-hot vector where $\mathbf{e}_d[j] = 1$ if $j = d$ and $\mathbf{e}_d[j] = 0$ for $j \neq d$.

We consider functions $f : \mathbb{R}^D \to \mathbb{R}$ and consider two features indexed by $a, b \in [D]$ as the features involved in the shape constraints. For notational simplicity, some of our shape constraint definitions will assume f is continuous and differentiable. However, all of the shape constraints can be applied to non-differentiable functions by

changing derivatives to differences in the standard manner. For example, if f is increasing in input a, it means that $f(x[a]+\epsilon) \geq f(x[a])$ for all $x[a]$ and $\epsilon \geq 0$, and if f is differentiable we write that the slope of f in the direction $x[a]$ is nonnegative: $\partial f/\partial x[a] \geq 0$ for all x.

We will analyze the proposed shape constraints for the following function classes, which have all been previously shown to be particularly amenable to shape constraints. The proposed shape constraints might also be verifiable or trainable for other function classes, such as neural networks and decision trees, but we leave that as an open question.

Linear: $f(\mathbf{x}) = \alpha_0 + \sum_{d=1}^{D} \alpha_d \mathbf{x}[d]$.

Piecewise Linear Function (PLF): Given a set of $K+1$ knot-value pairs $\{(\xi_k, \beta_k)\}_{k=0}^{K} \subseteq \mathbb{R}^2$, where $\xi_0 < \ldots < \xi_K$, we denote by $\text{PLF}(x; \{(\xi_k, \beta_k)\})$ the univariate piecewise linear function interpolating the $K+1$ knot-value pairs. For $k \in [K]$, let γ_k denote the slope of the kth linear segment: $(\beta_k - \beta_{k-1})/(\xi_k - \xi_{k-1})$.

Generalized Additive Model (GAM): $f(\mathbf{x}) = \sum_{d=1}^{D} f_d(\mathbf{x}[d])$, where $f_d : \mathbb{R} \to \mathbb{R}$ for $d \in [D]$. A GAM-PLF is a special case where each $f_d(x)$ is $\text{PLF}(x; \{(\xi_{d,k}, \beta_{d,k})\})$, with slopes $\gamma_{d,k}$.

Lattice: A multidimensional interpolated look-up table, such that the function parameters are the function values sampled on a regular grid (Garcia & Gupta, 2009; Garcia et al., 2012). See Fig. 5 in the Appendix for an example. For $x \in \mathbb{R}^1$, a lattice is simply a PLF with uniform knots. With enough look-up table parameters, one can fit arbitrary bounded continuous functions. The look-up table structure is helpful for imposing shape constraints (Gupta et al., 2016; 2018; Cotter et al., 2019a).

A D-dimensional lattice of size $\mathbf{V} \in \mathbb{N}^D$ consists of a regular D-dimensional grid of look-up table vertices $\mathcal{M}_{\mathbf{V}} = \{0, 1, \ldots, \mathbf{V}[1]-1\} \times \ldots \times \{0, 1, \ldots, \mathbf{V}[D]-1\}$. Thus, $\mathbf{V}[d]$ is the number of vertices in the lattice in the dth dimension, and the grid has $\prod_{d=1}^{D} \mathbf{V}[d]$ vertices. We assume the input vector $\mathbf{x} \in \mathbb{R}^D$ has been bounded (clipped if necessary), shifted and scaled such that each $\mathbf{x}[d]$ lies in $[0, \mathbf{V}[d]-1]$. We define the cell of $\mathbf{x}$ to be the set of its 2^D neighboring grid vertices given by $\mathcal{N}(\mathbf{x}) = \{\lfloor \mathbf{x}[1] \rfloor, \lfloor \mathbf{x}[1] \rfloor + 1\} \times \ldots \times \{\lfloor \mathbf{x}[D] \rfloor, \lfloor \mathbf{x}[D] \rfloor + 1\}$.

For each vertex $\mathbf{v} \in \mathcal{M}_{\mathbf{V}}$, there is a corresponding look-up table parameter $\theta_{\mathbf{v}} \in \mathbb{R}$. The lattice function is produced by *interpolating* the grid's parameters over each cell. While there are many possible interpolation operators, here we consider only the popular multilinear interpolation:

$$f(\mathbf{x}) = \sum_{\mathbf{v} \in \mathcal{N}(\mathbf{x})} \theta_{\mathbf{v}} \Phi_{\mathbf{v}}(\mathbf{x}), \tag{1}$$

where $\Phi_{\mathbf{v}}(\mathbf{x})$ is the linear interpolation weight on vertex $\mathbf{v}$ given by:

$$\Phi_{\mathbf{v}}(\mathbf{x}) = \prod_{d=1}^{D} \Big(1 + (\mathbf{x}[d] - \mathbf{v}[d])(-1)^{I_{\mathbf{v}[d] = \lfloor \mathbf{x}[d] \rfloor}}\Big), \tag{2}$$

and I is the standard indicator function.

Calibrated Lattice: A generalization of a GAM-PLF, where instead of simply summing PLF's, the D PLF's enter a second layer that is a lattice (or ensemble of lattices) that captures nonlinear feature interactions (Gupta et al., 2016). The first layer of PLFs are called *calibrators* and are often *capped* such that $\beta_0 = 0$ and $\beta_K = 1$ to control the domain for the second layer. Ensembles of calibrated lattices perform similarly to random forests (Canini et al., 2016), and calibrator and lattice layers can be cascaded into *deep lattice networks* (You et al., 2017) that perform similarly to DNN's (You et al., 2017; Gupta et al., 2018; Cotter et al., 2019a).

Ensemble: An ensemble assumes the existence of T base models $\{f_t(x)\}$ where each $f_t : \mathbb{R}^D \to \mathbb{R}$. The base models may ignore some of the D inputs, essentially acting on only a subset of the features. The ensemble outputs the sum $f(x) = \sum_{t=1}^{T} f_t(x)$.

4. Dominance Shape Constraints

We propose new shape constraints to capture the prior knowledge or a policy that feature a should be more important than feature b. For example, in time series modeling, we often believe that recent information should be more important than past information at predicting future values. Or if a model is trained to predict CTR for a web link from feature a, the past CTR on that web link, and feature b, the past mean CTR for the whole website, then one might want to constrain the model to be more sensitive to a than b.

One option would be to require the model to always respond more strongly to changes in input a than to changes in input b, that is, for a differentiable model, require $\left|\frac{\partial f(x)}{\partial \mathbf{x}[a]}\right| \geq \left|\frac{\partial f(x)}{\partial \mathbf{x}[b]}\right|$ for all x. This constraint is easy to verify or guarantee for linear models, as it holds if the coefficient on feature a has a larger magnitude than the coefficient on feature b: $|\alpha_a| \geq |\alpha_b|$. For more flexible functions, we can say more if we *also* require the model to be monotonic with respect to both features a and b:

Monotonic Dominance: $\frac{\partial f(x)}{\partial \mathbf{x}[a]} \geq \frac{\partial f(x)}{\partial \mathbf{x}[b]} \geq 0$

This constraint can be expressed as a set of linear inequality constraints for GAM and lattice models too, as detailed in Table 1 with proofs in the Appendix. See Fig. 1 and Fig. 2 for illustrations.

Dominance and monotonic dominance are sensitive to the

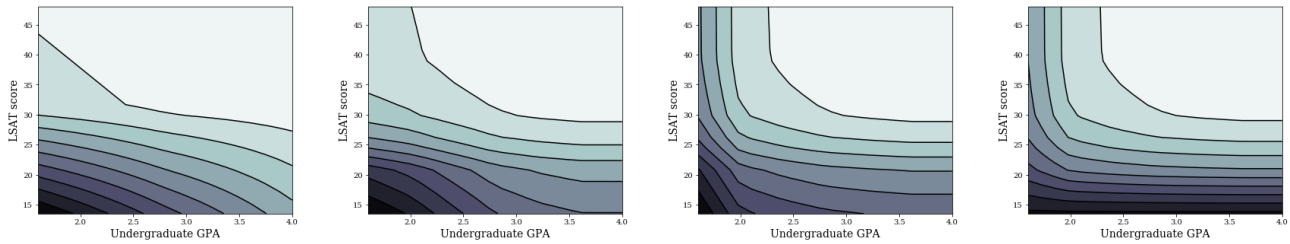

Figure 2. Contour plots illustrating the different functions produced by training a calibrated lattice model with different 2D shape constraints. All were trained on the same dataset to predict the probability that a person passes the bar given their GPA (horizontal axis) and LSAT test score (vertical axis) (Wightman, 1998) (see Appendix for details). A lighter contour color indicates a higher probability. All four plots were constrained to be monotonic in both GPA and LSAT (Wang & Gupta, 2020). **Left:** Trained with scaled monotonic dominance of LSAT over GPA, consistent with folklore that LSAT is the single most important factor (US News, 2018). **Middle Left:** Trained with range dominance of LSAT over GPA. **Middle Right:** Trained with Edgeworth complements constraint. **Far Right:** Trained with the trapezoid constraint that higher GPA means that the model should be more sensitive to the LSAT score.

units in which a and b are defined. For example, suppose one wishes to predict the increase in COVID cases after a large live event given a, the number of people who show up, and b, the number of bathroom stalls available per person at the venue. We expect the cases to be a monotonically increasing function of attendees a, and a monotonically decreasing function of the bathroom density b. One might wish to check or impose a dominance constraint that attendees a is more important than bathroom density b. However, clearly one-to-one is not the right trade-off between attendees-and-bathroom-density, so we also propose a scaled variant:

Scaled Monotonic Dominance: a monotonically dominates b with respect to scale $C \geq 0$ if $\frac{\partial f(x)}{\partial \mathbf{x}[a]} \geq C \frac{\partial f(x)}{\partial \mathbf{x}[b]} \geq 0$.

Given a trained model, we can check if that model satisfies scaled monotonic dominance for any C. Further, if we wish to train a model that satisfies scaled monotonic dominance, we simply need to architect the model with a first-layer that multiplies input a by $1/C$, and then enforce monotonic dominance on the rest of the layers. See Fig. 2 for an illustration.

For GAMs and two-layer calibrated lattice models, satisfying monotonic dominance requires the strong requirement that the slope of the 1-d transform for feature a must be steeper *everywhere* than the slope of the 1-d transform on feature b *anywhere*; that is quite restrictive on the 1-d feature transformations for a and b. For example, the model could not learn to transform a and b into an approximation of $\log a$ and $\log b$ using PLF's. To allow the model to learn flexible feature transformations for a and b, but still capture an intuition that feature a should be more important than feature b, we propose a mathematically more relaxed dominance shape constraint:

Range Dominance: For any input x the range of possible outputs $f(x)$ must be bigger if one varies input a than if one varies input b.

Returning to the example of predicting COVID cases from a live event, bound the domain of a to the number of people allowed at the venue (the capacity) $a_{\max}$, then range dominance of a over b would require that for any choice of $a_0 \in [1, a_{\max}]$ and any bathroom density $b_0 \in [0, 1]$, keeping the bathroom density b fixed at b_0 but evaluating the model for a set to the minimum number of attendees or the maximum number should change the predicted cases f more than keeping the number of attendees a fixed at a_0 and ranging the bathroom density from $b = 0$ to $b = 1$.

See Fig. 1 and Fig. 2 for illustrations. In Table 1 we give the mathematical definition and the sufficient conditions to achieve range dominance for different function classes (proofs in the appendix). Range dominance enables more flexible two-layer modeling than monotonic dominance. Both range dominance and monotonic dominance are not symmetric but are transitive.

5. Complements Shape Constraints

We first review and expand on two notions of complements from Cotter et al. (2019a), and then propose a third shape constraint that also captures complementarity. Table 6 in the Appendix summarizes definitions and properties. Proofs for all statements are in the Appendix.

In economics, *Edgeworth complementarity* says that a feature a and feature b are *complements* if the marginal value of feature a increases for larger values of feature b (Amir, 2005). For example, having more books a in a community's library is more valuable if the community literacy rate b is higher. Cotter et al. (2019a) proposed codifying this as the *Edgeworth shape constraint*. For a

Name	Monotonic Dominance	Range Dominance
Definition	$\frac{\partial f(\mathbf{x})}{\partial \mathbf{x}[a]} \geq \frac{\partial f(\mathbf{x})}{\partial \mathbf{x}[b]} \geq 0$ for all $\mathbf{x}$	$\frac{\partial f(x)}{\partial \mathbf{x}[a]} \geq 0$, $\frac{\partial f(x)}{\partial \mathbf{x}[b]} \geq 0$, and $f(U_a(\mathbf{x})) - f(L_a(\mathbf{x})) \geq f(U_b(\mathbf{x})) - f(L_b(\mathbf{x}))$
Linear	$\alpha_a \geq \alpha_b \geq 0$	$\alpha_a, \alpha_b \geq 0$ & $\alpha_a(u_a - \ell_a) \geq \alpha_b(u_b - \ell_b)$
GAM-PLF	$\gamma_{a,i} \geq \gamma_{b,j} \geq 0, \forall i \in [K_a], j \in [K_b]$	$\gamma_{a,i} \geq 0, \gamma_{b,j} \geq 0, \forall i \in [K_a], j \in [K_b]$ and $\beta_{a,K_a} - \beta_{a,0} \geq \beta_{b,K_b} - \beta_{b,0}$
Lattice	$\theta_{\mathbf{v}+\mathbf{e}_a+\mathbf{e}_b} \geq \theta_{\mathbf{v}+\mathbf{e}_a}, \theta_{\mathbf{v}+\mathbf{e}_a+\mathbf{e}_b} \geq \theta_{\mathbf{v}+\mathbf{e}_b}$, $\theta_{\mathbf{v}+\mathbf{e}_a} \geq \theta_{\mathbf{v}}, \theta_{\mathbf{v}+\mathbf{e}_b} \geq \theta_{\mathbf{v}}$, $\theta_{\mathbf{v}+\mathbf{e}_b} \leq \frac{\theta_{\mathbf{v}}+\theta_{\mathbf{v}+\mathbf{e}_a+\mathbf{e}_b}}{2} \leq \theta_{\mathbf{v}+\mathbf{e}_a}$ $\forall \mathbf{v} \in \mathcal{M}_{\mathbf{V}}, \mathbf{v}[a] \leq \mathbf{V}[a] - 2, \mathbf{v}[b] \leq \mathbf{V}[b] - 2$	$\theta_{\mathbf{v}+\mathbf{e}_a+\mathbf{e}_b} \geq \theta_{\mathbf{v}+\mathbf{e}_a}, \theta_{\mathbf{v}+\mathbf{e}_a+\mathbf{e}_b} \geq \theta_{\mathbf{v}+\mathbf{e}_b}$, $\theta_{\mathbf{v}+\mathbf{e}_a} \geq \theta_{\mathbf{v}}, \theta_{\mathbf{v}+\mathbf{e}_b} \geq \theta_{\mathbf{v}}$, $\theta_{U_a(\mathbf{v})} - \theta_{L_a(\mathbf{v})} \geq \theta_{U_b(\mathbf{v})} - \theta_{L_b(\mathbf{v})}, \forall \mathbf{v} \in \mathcal{M}_{\mathbf{V}}$
Cal. Lattice	Lattice constraints & GAM-PLF constraints	Lattice constraints & capped increasing calibrators
Composition	$\sum_t f_t(x)$ holds, and $g(f(A, B))$ holds if g is increasing	$\sum_t f_t(x)$ holds, and $g(f(A, B))$ holds if g is increasing and affine

Table 1. Dominance definitions, sufficient conditions, and properties for a dominates b. Proofs in the Appendix.

continuous twice-differentiable function:

Edgeworth: $\frac{\partial}{\partial \mathbf{x}[b]} \left(\frac{\partial f(x)}{\partial \mathbf{x}[a]} \right) \geq 0.$

Note that one can satisfy the Edgeworth shape constraint without f being monotonic in either feature a or b. For example, suppose a model predicts how fun a vacation will be, and that feature a represents the prevalence of mosquitos, and feature b represents packing bug spray. The more bugs there are, the more value there is in packing more bug spray. However, neither bugs nor carrying bug spray is a positive feature for a vacation.

Economists might say that Edgeworth is equivalent to *supermodularity* if it holds for the entire input space (Amir, 2005; Topkis, 1978), though the machine learning literature tends to use the term *supermodular* only for Boolean inputs. There is prior work in learning functions guaranteed to be supermodular (or submodular) on Boolean inputs by learning good weights on submodular component functions (Tschiatschek et al., 2014; Dolhansky & Bilmes, 2016). By Schwarz's Theorem, the Edgeworth shape constraint is symmetric, but it is not transitive.

The major downside to Edgeworth constraints is that they are difficult to guarantee for deep models because monotone transformations of Edgeworth functions are not necessarily still Edgeworth (Amir, 2005; Cotter et al., 2019a). So Cotter et al. (2019a) proposed a second more robust formalization of the notion of complements that they called a *trapezoid* shape constraint. The trapezoid constraint differs from Edgeworth in that its definition is *asymmetric* on a and b, and assumes that f is monotonic with respect to a. Trapezoid requires that for larger values of b, the *range* of possible outputs if you vary a must get bigger, forming a trapezoid of possible outputs (see Fig. 1). Trapezoid is useful when feature a is a measurement correlated with the training label, and b captures the precision or trustworthiness of the measurement a, and thus the higher the precision b the more the model output should vary as a varies. For continuous differentiable functions, trapezoid is equivalent to:

Trapezoid: $\partial f(\mathbf{x})/\partial \mathbf{x}[a] \geq 0$ and $\partial f(L_a(\mathbf{x}))/\partial \mathbf{x}[b] \leq 0$ and $\partial f(U_a(\mathbf{x}))/\partial \mathbf{x}[b] \geq 0$.

Trapezoid is nicer than Edgeworth for deep models because if f is trapezoid on $a - b$, then $g(f)$ is also trapezoid on $a - b$ if g is monotonically increasing (Cotter et al., 2019a).

We propose a third shape constraint that captures a different sense of complementarity: let *joint monotonicity* require the model to be monotonically increasing along the diagonal of every $a - b$ slice of the feature space, as illustrated in Fig. 1. For differentiable functions:

Joint Monotonicity: $\frac{\partial f(x)}{\partial \mathbf{x}[a]} + \frac{\partial f(x)}{\partial \mathbf{x}[b]} \geq 0.$

Joint monotonicity is weaker than requiring the function to be monotonic with respect to each of the constrained features individually. For example, suppose $f = 5ab - 2a + 9b$ models the profit of a hotel given a hotel beds and b hotel guests. Then the profit is monotonically increasing w.r.t. to guests b, but not w.r.t. beds a, but it is *jointly monotonic* along the diagonal of +1 bed for every +1 guest.

The following conditions are sufficient to make a model jointly monotonic on a and b (proofs in the Appendix; see Table 6 in the Appendix to compare and contrast these conditions with Edgeworth and trapezoid conditions):

Linear: $\alpha_a + \alpha_b \geq 0$.

GAM-PLF: $\gamma_{a,i} + \gamma_{b,j} \geq 0 \, \forall i \in [K_a], j \in [K_b]$.

Lattice: $\theta_{\mathbf{v}} \leq \frac{\theta_{\mathbf{v}+\mathbf{e}_a}+\theta_{\mathbf{v}+\mathbf{e}_b}}{2} \leq \theta_{\mathbf{v}+\mathbf{e}_a+\mathbf{e}_b}, \forall \mathbf{v} \in \mathcal{M}_{\mathbf{V}}$, $\mathbf{v}[a] \leq \mathbf{V}[a] - 2, \mathbf{v}[b] \leq \mathbf{V}[b] - 2$.

Calibrated Lattice: The lattice layer must satisfy the lattice conditions, and the calibrators for a and b must be increasing and affine with the same slope.

Composition: If f is jointly monotonic and g is increasing, then $g(f)$ is jointly monotonic. If each f_t is jointly monotonic, then so is an ensemble $\sum_t f_t$.

Mathematically, joint monotonicity is analogous to the monotonic dominance constraint we propose in Section 4, in that both constraints require the function to be monotonic along a diagonal direction of every $a - b$ subspace.

Joint monotonicity is unit-sensitive. For example, consider a model that predicts sale price of a house based on the size of the house (a) and the number of bedrooms (b), as well as other features. For a fixed house size a, having more bedrooms is not necessarily good for the sale price as the bedrooms become too small. But if we increase a by some size for each bedroom we add, we may be confident that *a larger house with more bedrooms will sell for more*. That is, we believe there exists some ray in the $a - b$ space along which the model should be monotonic. In such cases, we propose that a relaxed version of the joint monotonicity constraint may be more appropriate:

Scaled Joint Monotonicity: a is jointly monotonic with b with respect to scale C if $C\frac{\partial f(x)}{\partial \mathbf{x}[b]} + \frac{\partial f(x)}{\partial \mathbf{x}[a]} \geq 0$.

This constraint can be achieved, e.g., with a two-layer model by imposing the joint monotonicity constraints on the second layer, and by allowing the first layer to *only* calibrate a and b, where their calibrators are affine transformations. The slopes of these transformations can additionally be constrained to be sensible.

Two or more of these complements shape constraints may be useful at once. For example, a firm may believe that its profit satisfies scaled joint monotonicity in how many new machines it buys (a) and how much training it pays for (b); and Edgeworth in that the more training they buy, the more value they will get out of their machines.

These complements shape constraints can be extended to address larger sets of complements features, as occurs in such diverse settings as multi-item auctions (Roth, 2002) and manufacturing (Milgrom et al., 1991).

6. Diminisher Shape Constraints

We say that a feature b *diminishes* feature a if the model cares less about feature a if the value for feature b is bigger. Classic examples are substitutable features, like predicting the amount of time a person spends watching videos given the Boolean features a indicating if the person has a subscription to Netflix, and b indicating if the person has a subscription to HBO.

Diminishers also occur in utility models where there is a finite budget. For example, suppose a model predicts the happiness of customers if two products A and B are put on sale for a and b off respectively. Some customers cannot afford to buy both A and B. So the value of either being on sale is of less value given that the other is also on sale.

Diminishers do not need to be substitutes. For example, economists believe many humans value an absolute discount of a less if the sales price b is higher (Husemann-Kopetzky, 2018; Tversky & Kahneman, 1981). Or, in ranking coffee shops, the farther a coffee shop is from you, the less you tend to care about its star rating.

Mathematically, diminishers are simply complements in reverse, so we propose modeling them with analogous shape constraints. For example, *reverse Edgeworth*, which is equivalent to submodularity for Boolean inputs, and which for continuous differentiable functions can be expressed:

Reverse Edgeworth: $\frac{\partial}{\partial \mathbf{x}[b]}\left(\frac{\partial f(x)}{\partial \mathbf{x}[a]}\right) \leq 0$.

7. Unimodality Shape Constraints

The last category of shape constraints we consider are functions where, for any fixed value of the other features, the function is *unimodal* over some subset of features such that there exists a global minimizer and the function is increasing along any ray starting at the global minimizer:

Unimodal: Let $\mathcal{X} \subset \mathbb{R}^D$ be a convex set. If there exists some $x^* \in \mathcal{X}$ such that $f(x^* + \epsilon v) \geq f(x^* + \delta v)$ for any $\epsilon \geq \delta \geq 0$ and $v \in \mathcal{R}^D$ such that $x^* + \epsilon v \in \mathcal{X}$ and $x^* + \delta v \in \mathcal{X}$, then f is unimodal on $\mathcal{X}$.

Examples of unimodal functions are the optimal location on the soccer field for a free kick, and progesterone levels over the menstrual cycle (Dunson, 2005).

Stout (2008) fits *1-dimensional* unimodal functions using

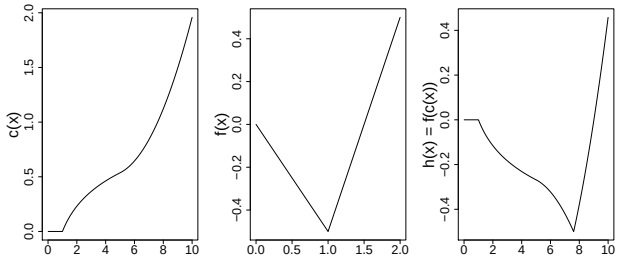

Figure 3. Example two-layer unimodal function. **Left:** The first layer is a 1D monontonically increasing PLF $c(x)$. **Middle:** The second layer is a 1D unimodal PLF $f(x)$ with its minimizer fixed at 1.0. **Right:** The composition $h(x) = f(c(x))$ is unimodal with an arbitrary minimizer (here, roughly 7.8)

two separate isotonic regressions. Similarly, Köllmann et al. (2014) models *1-dimensional* unimodal functions using Bernstein–Schoenberg splines, which can also be trained with linear inequality constraints. Dunson (2005) considered this problem under the name *umbrella-order restriction* for the means of Bayesian models. All of these strategies require either already knowing the minimizer, or require consideration of all possible candidates for the minimizer.

A special case of unimodality is jointly convex functions, which can be achieved by summing jointly convex basis functions (Kim et al., 2004; Magnani & Boyd, 2009). DNN's with ReLU activations can be constrained to be jointly convex over a subset of features (Dugas et al., 2009; Amos et al., 2017). GAMs have also been constrained for convexity (Pya & Wood, 2015; Chen & Samworth, 2016), as have calibrated lattice models (Gupta et al., 2018).

Here we show that one can verify or train *multi-dimensional* unimodal functions, without requiring *convexity*, by using two-layer models.

First, consider the simple 1D case. We propose a two-layer model, where each layer is a PLF. For the second layer PLF, we impose unimodality with a pre-fixed minimizer x^* at the center knot $k = K/2 + 1$ for odd K, then constrain the PLF to be increasing to the right of x^* and decreasing to the left of x^*. Then let the first layer PLF be an arbitrary monotonic 1D function, which nonlinearly stretches the input domain of the second layer, so that the x^* of the second layer can be achieved by any raw input value to the first layer PLF. See Fig. 3 for an example.

For a *multidimensional* lattice function, one can check if it is unimodal by finding its minimal look-up table parameter value x^* and checking if for every look-up table vertex v, the directional derivative in direction $v - x^*$ is non-negative. The training strategy is similar to the 1D case: make the first layer monotonically increasing calibrators to give the model flexibility to place the global minimizer, and make the second layer a lattice layer whose look-up table must have an odd number of knots in each dimension, and constrain the center knot to be the global minimizer x^*. Then, sufficient (but not necessary) constraints on the second layer are that *every* edge in the lattice needs to be increasing in the direction away from the center. Full details, proofs and illustrations can be found in the Appendix.

8. Training With Multi-d Shape Constraints

Let $\{x_i, y_i\}$ be a train set with $i = 1, \ldots, n$ training example pairs where $x_i \in \mathbb{R}^D$ and $y_i \in \mathbb{R}$ for regression problems or $y_i \in [-1, 1]$ for classification problems. Let ℓ be a loss of interest, e.g. squared error or logistic loss. Let θ denote the parameters of any $f \in \mathcal{F}$, where $\mathcal{F}$ is a function class where the desired shape constraints can be expressed as linear inequalities on the parameters θ (as we have shown is the case for the proposed shape constraints for linear, GAM-PLF, lattice, calibrated lattice, ensembles, and $g(f())$ multi-layer models). Collect the linear inequalities corresponding to the set of desired shape constraints into the matrix equation $S^T\theta \geq 0$. Then train by minimizing the empirical risk subject to the linear inequality constraints:

$$\arg\min_{\theta} \sum_{i=1}^{n} l((f(x_i; \theta), y_i) \quad \text{such that } S^T\theta \geq 0. \tag{3}$$

To solve (3), we use projected stochastic gradient descent in TensorFlow, and the TensorFlow Lattice 2.0 library's PLF layers and lattice layers. After each minibatch, we project the model parameters toward the constraint set using ten steps of Dykstra's algorithm, with a longer final projection to ensure constraint satisfaction to within numerical precision.

Open-source code has been pushed to the TensorFlow Lattice 2.0 library and can be downloaded at github.com/tensorflow/lattice. Train time with 2D shape constraints was 10-20% longer than train time without 2D shape constraints (data in the Appendix).

9. Experiments

We present experiments on public and proprietary real-world problems illustrating training with the proposed shape constraints, including the first public experimental evidence with Edgeworth constraints (Cotter et al. (2019a) only presented experiments with trapezoid constraints).

For all experiments, we used the default ADAM stepsize of .001, ran the optimization of (3) until train loss converged, and used squared error as the training loss l in (3). All models that use PLFs use 10 keys $K_d = 10$ for each PLF, fixed before training to the train data quantiles. For each metric, we report a 95% margin of error, computed under a Gaus-

sian assumption as plus-or-minus 1.96 times the estimated standard error.

9.1. Weekly Sales Transactions (Regression)

We compare models that forecast next week's sales based on sales in the past weeks (Kaggle, 2020c). The dataset contains purchases of 811 products with a feature for the normalized transactions for each of the last 52 weeks. We train lattice models using the last $K \in \{2, \ldots, 10\}$ weeks of transactions as features to predict the transactions of the most recent week.

Results in Fig. 4 are averaged over 100 random 80-20 train/test splits. The figure shows that the unconstrained calibrated lattice model overfits to the data as we use a longer transaction history as features, reflected by the fact that the Train MSE improves yet the Test MSE deteriorates. Imposing monotonicity constraints on past history (i.e. any increase in prior sales should only ever increase our prediction) helps alleviate the issue, and imposing dominance constraints that encode our intuition that recent weeks should matter more than distant weeks helps even more.

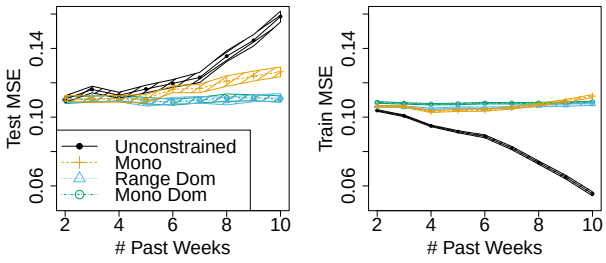

Figure 4. The Test and Train MSE of lattice models with different constraints on the Weekly Sales Transactions data. The x-axis shows what happens if you include more past weeks as features in the model. The shaded area shows the 95% confidence region of the performance.

9.2. Mount Rainier Climbing (Regression)

In this experiment, we compare the performance of different models in predicting the daily success probability of climbing Mount Rainier (Kaggle, 2020a) for 464 days based on five features: temperature, battery voltage, relative humidity, solar radiation and the month of the climbing day. On days people didn't climb (mostly during winter seasons), we labeled as having success probability zero.

The results, shown in Table 2, are averaged over 100 random 80-20 train/test splits. The calibrated lattice model does better than the GAM-PLF, suggesting that feature interactions are important. Preliminary analysis of the marginals of the data suggested to us that climbing success might be unimodal in humidity and month, consistent with a story that average-humidity summer days are the best time to climb Mt. Rainier. In fact, imposing 2D unimodality on humidity-month slices lowered Test MSE.

Table 2. Results of different models on the Mount Rainer Climbing dataset from Kaggle.

Model	Train MSE	Test MSE
GAM-PLF	0.029 ± 0.0004	0.032 ± 0.0018
Calib. Lattice	0.022 ± 0.0004	0.030 ± 0.0019
Above + Unimodal	0.023 ± 0.0004	0.027 ± 0.0018

9.3. Play Store App Installs Dataset (Regression)

The Google Play Store Apps dataset (Kaggle, 2020b) contains various pieces of information about individual apps. We trained a calibrated lattice model on $D = 5$ features: log real-valued number of reviews, average rating, price, download size, and categorical content rating. The label was bucketed number of times the app has been installed: we used the log of the lower bound of each interval as our label, and filtered for apps that have been installed at least once, leaving $N = 10,285$ examples.

We ran two sets of experiments. The first set used a non-IID train/test split based on a 6th piece of information, the app category: we trained on the most common category "Family"(18 percent of samples), and tested on the other app categories (82 percent of samples). The second set of experiments used 80/20 train/test IID random splits, and used the app category as a 6th feature.

We imposed three common sense monotonicity constraints that installs are increasing number of reviews, increasing in average rating, and decreasing in price. We tested an Edgeworth constraint that number of reviews and average rating are complements, with the motivation that lots of reviews are a better predictor of installs if the ratings are good, and high ratings are a more useful and reliable predictor of installs if there are sufficiently many reviews. We also tested range dominance on different feature interactions to better understand the underlying data, and found that the best dominance constraints were that number of reviews dominates price, and price dominates average rating, which are the ones reported in Tables 3 and 4. We also report imposing all the shape constraints of the above rows.

For the non-IID experiments, Table 3 shows that the constrained models never hurt Test MSE, and the dominance constraints improved the Test MSE. Further, one can explain what the constrained models are doing in terms of their shape constraints, which makes them easier to understand and debug. As expected, train MSE's are a little higher for the constrained models.

Table 4 shows the results averaged over 100 random IID

Table 3. Results on the Play Store dataset from Kaggle for 100 retrains of a fixed non-IID train/test split and $D = 5$ features (randomness due to initialization and optimization).

Model	Train MSE	Test MSE
Unconstrained	1.161 ± 0.0003	1.279 ± 0.0028
Mono.	1.198 ± 0.0009	1.276 ± 0.0028
Mono. + Edge.	1.199 ± 0.0009	1.279 ± 0.0032
Mono. + Dom.	1.189 ± 0.0003	1.250 ± 0.0020
All Constraints	1.189 ± 0.0003	1.252 ± 0.0020

80/20 train/test splits with $D = 6$ features. Because the train/test is IID, and because there is more train data, there is less need to regularize, and the Test MSE's are all statistically indistinguishable from the unconstrained model. Still, the constrained models can be explained in terms of their shape constraint properties, which also make these models more predictable and thus easier to debug.

Table 4. Results on the Play Store dataset from Kaggle for 100 different IID 80/20 train/test splits and $D = 6$ features.

Model	Train MSE	Test MSE
Unconstrained	1.232 ± 0.0028	1.234 ± 0.0112
Mono.	1.230 ± 0.0031	1.236 ± 0.0119
Mono. + Edge.	1.233 ± 0.0025	1.229 ± 0.0091
Mono. + Dom.	1.236 ± 0.0029	1.237 ± 0.0097
All Constraints	1.236 ± 0.0028	1.246 ± 0.0093

9.4. User Intent Prediction (Classification)

We predict the user's intent given a query between two classes; the data is Google proprietary. We use a non-IID split: we train on 300K labeled samples from the U.S., and test on 350K from other countries. Of the $D = 19$ features, based on domain expertise, a priori 11 features were constrained to be monotonic, 4 feature pairs were chosen for range dominance, 11 feature pairs were chosen as complements and 3 feature pairs were chosen as diminishers, for which we applied either Edgeworth/reverse Edgeworth or trapezoid/reverse trapezoid constraints. The model was an ensemble of 50 calibrated lattices with each base model seeing 6-10 of the $D = 19$ possible features (Canini et al., 2016). We did not try joint monotonicity or monotonic dominance for this problem as they force the calibrators to be linear, which we knew would reduce the model flexibility too much given the large train set. All models were trained for 100 epochs; longer training and smaller lattices show similar trends but with worse Test MSE - see the Appendix.

While this is a classification problem, we trained with squared-error loss to produce stable prediction scores; Table 5 shows the Train and Test MSE averaged over 5 runs. As expected, Train MSE increases slightly as more constraints are added, but the monotonic models with dominance, Edgeworth, and both dominance and Edgeworth constraints have slightly better Test MSE compared to the monotonic only baseline. In contrast, trapezoid appears to hurt performance by Test MSE a little. The main advantage is that applying these constraints guarantees the model is behaving in reasonable and explainable ways, and its greater predictability makes debuggging model errors easier.

Table 5. Results on User Intent Prediction for 5 retrains of a fixed non-IID train/test split (randonmess due to initialization and optimization).

Model	Train MSE	Test MSE
Mono.	0.666 ± 0.0001	0.752 ± 0.0002
Mono. + Dom.	0.669 ± 0.0001	0.751 ± 0.0002
Mono. + Edge.	0.670 ± 0.0001	0.751 ± 0.0001
Mono. + Edge. + Dom.	0.673 ± 0.0001	0.750 ± 0.0002
Mono. + Trap.	0.686 ± 0.0001	0.764 ± 0.0002
Mono. + Trap. + Dom.	0.695 ± 0.0002	0.766 ± 0.0001

10. Conclusions

We compared and contrasted new and recent definitions for multidimensional shape constraints. Shape constraints play two key roles for interpretability: if applied, we can explain the feature interactions, and we can test different 2D shape constraints and see the effect to understand whether certain interactions fit or fight the data. Experimentally, we showed applying relevant shape constraints can regularize models and improve test metrics, particularly in non-IID settings.

We found that a key differentiator between the definitions is how easy the different constraints are to apply to two-layer models for flexible modeling. The proposed joint monotonicity and monotonic dominance satisfy intellectually for their simplicity: constrain the model to be increasing along a direction in the 2D a-b subspace (see Fig. 1). However, we found these two constraints difficult to use in practice because they overconstrained the calibrator layer. We found Edgeworth to be the most broadly applicable and useful 2D shape constraint for use with ensembles of two-layer calibrated lattices (which have similar flexibility as random forests (Canini et al., 2016)), and easy to explain to non-experts. We saw trapezoid performed worse than Edgeworth, probably due to its greater restrictions on the calibrator layer. We found range dominance to also be useful and easy to apply. The unimodality shape constraints were helpful, and are promising for training functions that one wants to minimize (as in Amos et al. (2017)), but more research is needed.

References

Amir, R. Supermodularity and complementarity in economics: An elementary survey. *Southern Economic Journal*, 71(3):636–660, 2005.

Amos, B., Xu, L., and Kolter, J. Z. Input convex neural networks. *ICML*, 2017.

Archer, N. P. and Wang, S. Application of the back propagation neural network algorithm with monotonicity constraints for two-group classification problems. *Decision Sciences*, 24(1):60–75, 1993.

Barlow, R. E., Bartholomew, D. J., and Bremner, J. M. *Statistical inference under order restrictions; the theory and application of isotonic regression*. Wiley, 1972.

Canini, K., Cotter, A., Fard, M. M., Gupta, M. R., and Pfeifer, J. Fast and flexible monotonic functions with ensembles of lattices. *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.

Chen, Y. and Samworth, R. J. Generalized additive and index models with shape constraints. *Journal Royal Statistical Society B*, 2016.

Chetverikov, D., Santos, A., and Shaikh, A. M. The econometrics of shape restrictions. *Annual Review of Economics*, 2018.

Cotter, A., Gupta, M. R., Jiang, H., Louidor, E., Muller, J., Narayan, T., Wang, S., and Zhu, T. Shape constraints for set functions. *ICML*, 2019a.

Cotter, A., Jiang, H., Wang, S., Narayan, T., Gupta, M. R., You, S., and Sridharan, K. Optimization with non-differentiable constraints with applications to fairness, recall, churn, and other goals. *JMLR*, 2019b.

Dolhansky, B. and Bilmes, J. Deep submodular functions: Definitions and learning. *NeurIPS*, 2016.

Dugas, C., Bengio, Y., Bélisle, F., Nadeau, C., and Garcia, R. Incorporating functional knowledge in neural networks. *JMLR*, 2009.

Dunson, D. A transformation approach for incorporating monotone or unimodal constraints. *Biostatistics*, 2005.

Garcia, E. K. and Gupta, M. R. Lattice regression. *NeurIPS*, 2009.

Garcia, E. K., Arora, R., and Gupta, M. R. Optimized regression for efficient function evaluation. *IEEE Trans. Image Processing*, 21(9):4128–4140, September 2012.

Groeneboom, P. and Jongbloed, G. *Nonparametric estimation under shape constraints*. Cambridge Press, New York, USA, 2014.

Gupta, M. R., Cotter, A., Pfeifer, J., Voevodski, K., Canini, K., Mangylov, A., Moczydlowski, W., and Esbroeck, A. V. Monotonic calibrated interpolated look-up tables. *Journal of Machine Learning Research*, 17(109):1–47, 2016. URL http://jmlr.org/papers/v17/15-243.html.

Gupta, M. R., Bahri, D., Cotter, A., and Canini, K. Diminishing returns shape constraints for interpretability and regularization. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

Howard, A. and Jebara, T. Learning monotonic transformations for classification. *Advances in Neural Information Processing Systems (NeurIPS)*, 2007.

Husemann-Kopetzky, M. *Handbook on the Psychology of Pricing*. Pricing School Press, New York, USA, 2018.

Kaggle. Kaggle Mount Rainier Weather and Climbing Data, 2020a. URL https://www.kaggle.com/codersree/mount-rainier-weather-and-climbing-data.

Kaggle. Kaggle Google Play Store Apps Data, 2020b. URL https://www.kaggle.com/lava18/google-play-store-apps.

Kaggle. Kaggle Weekly Sales Transactions, 2020c. URL https://www.kaggle.com/crawford/weekly-sales-transactions.

Kim, J., Lee, J., Vandenberghe, L., and Yang, C. Techniques for improving the accuracy of geometric-programming based analog circuit design optimization. *Proc. IEEE International Conference on Computer-aided Design*, 2004.

Köllmann, C., Bornkamp, B., and Ickstadt, K. Unimodal regression using Bernstein–Schoenberg splines and penalties. *Biometrics*, 70(4), 2014.

Magnani, A. and Boyd, S. P. Convex piecewise-linear fitting. *Optimization and Engineering*, 2009.

Mann, G. S. and McCallum, A. Simple, robust, scalable semi-supervised learning with expectation regularization. *ICML*, 2007.

Milgrom, P., Qian, Y., and Roberts, J. Complementarities, momentum, and the evolution of modern manufacturing. *The American Economic Review*, 1991.

Minin, A., Velikova, M., Lang, B., and Daniels, H. Comparison of universal approximators incorporating partial monotonicity by structure. *Neural Networks*, 23(4):471–475, 2010.

Pya, N. and Wood, S. N. Shape constrained additive models. *Statistics and Computing*, 2015.

Roth, A. E. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70(4):1341–1378, 2002.

Sill, J. Monotonic networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 1998.

Stout, Q. Unimodal regression via prefix isotonic regression. *Computational Statistics and Data Analysis*, 2008.

Topkis, D. Minimizing a submodular function on a lattice. *Operations Research*, 78(2):302–321, 1978.

Tschiatschek, S., Iyer, R., Wei, H., and Bilmes, J. Learning mixtures of submodular functions for image collection summarization. *NeurIPS*, 2014.

Tversky, A. and Kahneman, D. The framing of decisions and psychology of choice. *Science*, 1981.

US News. 5 traits that help people get into top law schools, 2018. URL https://www.usnews.com.

Wang, S. and Gupta, M. R. Deontological ethics by monotonicity shape constraints. In *AIStats*, 2020.

Wightman, L. LSAC national longitudinal bar passage study. *Law School Admission Council*, 1998.

You, S., Canini, K., Ding, D., Pfeifer, J., and Gupta, M. R. Deep lattice networks and partial monotonic functions. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

Zafar, M. B., Valera, I., Rodriguez, M. G., and Gummadi, K. P. Fairness constraints: Mechanisms for fair classification. In *AIStats*, 2017.

11. Appendix: Comparison of Complements Shape Constraints

In Table 6, we compare and contrast the Edgeworth and trapezoid constraints proposed in Cotter et al. (2019a) with the proposed joint monotonicity constraint that captures a different notion of complementarity. Proofs for all the results in the Table are given in the next section.

12. Appendix: Proofs

12.1. Notation

Throughout this section, for a function $f : \mathbb{R}^D \to \mathbb{R}$ we denote its partial derivative with respect to the ith input variable by $\partial_i f$. For a univariate function f we denote its derivative by f'. For a D-dimensional lattice of size $\mathbf{V}$ we denote the domain of its function by $\overline{\mathcal{M}_\mathbf{V}} = [0, \mathbf{V}[1]-1] \times \ldots \times [0, \mathbf{V}[D]-1]$.

12.2. Tools

To prove our results on lattices we'll need the following lemma, which gives a formula for the partial derivative of a lattice function.

Lemma 1. *Let $f : \mathbb{R}^D \to \mathbb{R}$ be the function of a lattice with dimension D, size vector $\mathbf{V}$ and vertex values $\{\theta_\mathbf{v}\}_{\mathbf{v}\in\mathcal{M}_\mathbf{V}}$. Then for all $d \in [D]$, and $\mathbf{x} \in \mathcal{M}_\mathbf{v}$ with $\mathbf{x}[d] \notin \mathbb{Z}$ (i.e. $\mathbf{x}$ does not lie on the boundary of two adjacent lattice cells in the dth direction)*

$$\partial_d f(\mathbf{x}) = \sum_{\mathbf{v}\in\mathcal{N}(\mathbf{x})} \Phi_\mathbf{v}(\mathbf{x})(\theta_{\lceil\mathbf{v}\rceil_{d,\mathbf{x}}} - \theta_{\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}}),$$

where $\lceil\mathbf{v}\rceil_{d,\mathbf{x}}$ is $\mathbf{v}+\mathbf{e}_d$, if $\mathbf{v}[d]=\lfloor\mathbf{x}[d]\rfloor$, or $\mathbf{v}$, otherwise, and $\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}=\lceil\mathbf{v}\rceil_{d,\mathbf{x}}-\mathbf{e}_d$.

Proof. Let $\mathbf{x}$ satisfy the requirements of the lemma. By (1), $\partial_d f(\mathbf{x})=\sum_{\mathbf{v}\in\mathcal{N}(\mathbf{x})} \theta_\mathbf{v}\partial_d\Phi_\mathbf{v}(\mathbf{x})$. Denoting by $\lambda(v,x)=1+(x-v)(-1)^{I_{v=\lfloor x\rfloor}}$, for $x\in\mathbb{R}$ and $v\in\mathbb{N}$, we get

$$\begin{aligned}\partial_d f(\mathbf{x}) &= \sum_{\mathbf{v}\in\mathcal{N}(\mathbf{x})} \theta_\mathbf{v}\partial_d \prod_{i=1}^{D}\lambda(\mathbf{v}[i],\mathbf{x}[i]) \\ &= \sum_{\mathbf{v}\in\mathcal{N}(\mathbf{x})} \theta_\mathbf{v}(-1)^{I_{\mathbf{v}[d]=\lfloor\mathbf{x}[d]\rfloor}} \prod_{i\neq d}\lambda(\mathbf{v}[i],\mathbf{x}[i]),\end{aligned}$$

, where we used the fact that for $x \in \mathbb{R}\setminus\mathbb{Z}$, $\partial\lambda/\partial x = (-1)^{I_{v=\lfloor x\rfloor}}$. Partitioning the set $\mathcal{N}(\mathbf{x})$ of size 2^D into the 2^{D-1} pairs $\{(\mathbf{v},\lceil\mathbf{v}\rceil_{d,\mathbf{x}}) : \mathbf{v}\in\mathcal{N}(\mathbf{x}), \mathbf{v} = \lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}\}$, we may regroup the summands to obtain

$$\partial_d f(\mathbf{x}) = \sum_{\substack{\mathbf{v}\in\mathcal{N}(\mathbf{x})\\ \mathbf{v}=\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}}} \left(\theta_{\lceil\mathbf{v}\rceil_{d,\mathbf{x}}} - \theta_{\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}}\right) \prod_{i\neq d}\lambda(\mathbf{v}[i],\mathbf{x}[i]) \quad (4)$$

Now, observe that $1 = \lambda(\lfloor\mathbf{x}[d]\rfloor,\mathbf{x}[d])+\lambda(\lfloor\mathbf{x}[d]\rfloor+1,\mathbf{x}[d])$. Thus, for $\mathbf{v}\in\mathcal{N}(\mathbf{x})$ with $\mathbf{v} = \lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}$, it holds that

$$\begin{aligned}\prod_{i\neq d}\lambda(\mathbf{v}[i],\mathbf{x}[i]) &= \big(\lambda(\lfloor\mathbf{x}[d]\rfloor,\mathbf{x}[d]) + \lambda(\lfloor\mathbf{x}[d]\rfloor+1,\mathbf{x}[d])\big)\cdot \prod_{i\neq d}\lambda(\mathbf{v}[i],\mathbf{x}[i]) \\ &= \Phi_\mathbf{v}(\mathbf{x}) + \Phi_{\lceil\mathbf{v}\rceil_{d,\mathbf{x}}}(\mathbf{x}). \quad (5)\end{aligned}$$

Substituting (5) into (4), we get

$$\begin{aligned}\partial_d f(\mathbf{x}) &= \sum_{\substack{\mathbf{v}\in\mathcal{N}(\mathbf{x})\\ \mathbf{v}=\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}}} \left(\theta_{\lceil\mathbf{v}\rceil_{d,\mathbf{x}}} - \theta_{\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}}\right)\left(\Phi_\mathbf{v}(\mathbf{x}) + \Phi_{\lceil\mathbf{v}\rceil_{d,\mathbf{x}}}(\mathbf{x})\right) \\ &= \sum_{\mathbf{v}\in\mathcal{N}(\mathbf{x})} \Phi_\mathbf{v}(\mathbf{x})\left(\theta_{\lceil\mathbf{v}\rceil_{d,\mathbf{x}}} - \theta_{\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}}\right).\end{aligned}$$

□

12.3. Monotonic Dominance Results

Here are formal statements and proofs for our results on monotonic dominance.

12.3.1. Monotonic Dominance Criteria for GAMs

Proposition 1. *Let $f(\mathbf{x}) = \sum_{d=1}^D f_d(\mathbf{x}[d])$ be a GAM. Then feature a monotonically dominates feature b in f iff $f'_a \geq f'_b \geq 0$*

Proof. Directly follows from the definition and that $\partial_d f = f'_d$. □

Corollary 1. *Let $f : \mathbb{R}^D \to \mathbb{R}$ be a D-dimensional linear model with parameters $\{\alpha_d\}_{d=0}^D$. Then feature a monotonically dominates feature b iff $\alpha_a \geq \alpha_b \geq 0$.*

Proof. Follows from Proposition 1 because f can be regarded as a GAM with $f_d(x) = \alpha_d x + \alpha_0/D$. □

Corollary 2. *Consider the GAM-PLF $f(\mathbf{x})=\sum_{d=1}^D \mathrm{PLF}(\mathbf{x}[d];\{(\xi_{d,k},\beta_{d,jk})\}_{k=1}^{K_d})$. Then feature a monotonically dominates feature b iff $\gamma_{a,j} \geq \gamma_{b,k}$ for all $j \in [K_a-1], k \in [K_b-1]$.*

12.3.2. Monotonic Dominance Criteria for Lattice Models

Proposition 2. *Let f be a function for a lattice with size $\mathbf{V}$, dimension D, and vertex values $\{\theta_\mathbf{v}\}$. Assume f is monotonically increasing in features a and b. Then feature a monotonically dominates feature b iff for all $\mathbf{v}\in\mathcal{M}_\mathbf{V}$ with $\mathbf{v}[a] \leq \mathbf{V}[a]-2$ and $\mathbf{v}[b] \leq \mathbf{V}[b]-2$,*

$$\theta_{\mathbf{v}+\mathbf{e}_b} \leq \frac{\theta_\mathbf{v} + \theta_{\mathbf{v}+\mathbf{e}_a+\mathbf{e}_b}}{2} \leq \theta_{\mathbf{v}+\mathbf{e}_a}. \quad (6)$$

Multidimensional Shape Constraints

Name	Edgeworth Cotter et al. (2019a)	Trapezoid Cotter et al. (2019a)	Jointly Monotonic
Definition	$\frac{\partial}{\partial \mathbf{x}[b]}\left(\frac{\partial f(x)}{\partial \mathbf{x}[a]}\right) \geq 0$	$\frac{\partial f(\mathbf{x})}{\partial a} \geq 0$ & $\frac{\partial f(L_a(\mathbf{x}))}{\partial \mathbf{x}[b]} \leq 0$ & $\frac{\partial f(U_a(\mathbf{x}))}{\partial \mathbf{x}[b]} \geq 0$	$\frac{\partial f(x)}{\partial \mathbf{x}[a]} + \frac{\partial f(x)}{\partial \mathbf{x}[b]} \geq 0$
Example	bugs & bug spray	star rating & # reviews	hotel guests & hotel beds
Symmetric	yes	no	yes
Transitive	no	no	no
Linear	holds degenerately	doesn't hold (if $\alpha_b \neq 0$)	$\alpha_a + \alpha_b \geq 0$
GAM-PLF	holds degenerately	doesn't hold (if f_b is not constant)	$\gamma_{a,i} + \gamma_{b,j} \geq 0 \, \forall i \in [K_a], j \in [K_b]$
Lattice[a]	$\theta_{\mathbf{v}+\mathbf{e}_a} - \theta_{\mathbf{v}} \geq \theta_{\mathbf{v}+\mathbf{e}_a+\mathbf{e}_b} - \theta_{\mathbf{v}+\mathbf{e}_b}$, $\forall \mathbf{v} \in \mathcal{M}_{\mathbf{V}}$, $\mathbf{v}[a] \leq \mathbf{V}[a] - 2, \mathbf{v}[b] \leq \mathbf{V}[b] - 2$	f is increasing in feature a & $\theta_{L_a(\mathbf{v})} \geq \theta_{L_a(\mathbf{v})+\mathbf{e}_b}$ & $\theta_{U_a(\mathbf{v})} \leq \theta_{U_a(\mathbf{v})+\mathbf{e}_b}$, $\forall \mathbf{v} \in \mathcal{M}_{\mathbf{V}}, \mathbf{v}[b] \leq \mathbf{V}[b] - 2$	$\theta_{\mathbf{v}} \leq \frac{\theta_{\mathbf{v}+\mathbf{e}_a}+\theta_{\mathbf{v}+\mathbf{e}_b}}{2} \leq \theta_{\mathbf{v}+\mathbf{e}_a+\mathbf{e}_b}$, $\forall \mathbf{v} \in \mathcal{M}_{\mathbf{V}}$, $\mathbf{v}[a] \leq \mathbf{V}[a] - 2, \mathbf{v}[b] \leq \mathbf{V}[b] - 2$
Calibrated Lattice	lattice conditions & increasing c_a and c_b	lattice conditions & increasing c_a and c_b, capped c_a	lattice conditions & c_a, c_b are increasing and affine with the same slope
$g(f(A,B))$	f increasing in features a, b & g increasing & convex	g increasing	g increasing
Ensemble $\sum_t f_t(x)$	holds	holds[b]	holds[c]

Table 6. Two-dimensional Complements Shape Constraints, where input A complements input B, and $A, B \in [0, 1]$. Sufficient conditions (but not always necessary) conditions for the shape constraint to hold are given for different function classes, with the marking *degenerate* if the sufficient conditions always lead to a degenerate feasible solution. The last two rows of the table tell whether the constraint holds under composition if one either takes a one-dimensional transform $g : \mathbb{R} \to \mathbb{R}$ of a function f for which the constraint holds, or if one builds an ensemble out of many base models $\{f_t\}$ where the constraint holds for each base model.

[a] Cotter et al. (2019a) only gave conditions for a $2 \times 2 \times 2 \ldots \times 2$ lattice, here we give conditions for multi-cell lattices.
[b] Note that for the Trapezoid constraint this requires that a base model that has B as an input must also have A as an input.
[c] Note for the joint monotonicity constraint, this requires that every base model that has only one of the features as an input is monotonically increasing in that feature.

Proof. By definition, feature a monotonically dominates feature b if $\partial_a f - \partial_b f \geq 0$ for $\mathbf{x} \in \mathcal{M}_{\mathbf{V}}$.

By Lemma 1 we have

$$\partial_a f(\mathbf{x}) - \partial_b f(\mathbf{x}) = \sum_{\mathbf{v} \in \mathcal{N}(\mathbf{x})} \Phi_{\mathbf{v}}(\mathbf{x})(\theta_{\lceil \mathbf{v} \rceil_{a,\mathbf{x}}} - \theta_{\lfloor \mathbf{v} \rfloor_{a,\mathbf{x}}} - \theta_{\lceil \mathbf{v} \rceil_{b,\mathbf{x}}} + \theta_{\lfloor \mathbf{v} \rfloor_{b,\mathbf{x}}}).$$

Thus in each cell $\mathcal{N}(\mathbf{x})$, $\partial_a f(\mathbf{x}) - \partial_b f(\mathbf{x})$ is a multilinear interpolation of particular differences of the lattice parameters. Thus $\partial_a f - \partial_b \geq 0$ holds iff those differences are nonnegative, specifically, for each $\mathbf{v} \in \mathcal{M}_{\mathbf{V}}$ with $\mathbf{v}[a] \leq \mathbf{V}[a] - 2, \mathbf{v}[b] \leq \mathbf{V}[b] - 2$:

$$\begin{aligned}
\theta_{\mathbf{v}+\mathbf{e}_a} - \theta_{\mathbf{v}} - \theta_{\mathbf{v}+\mathbf{e}_b} + \theta_{\mathbf{v}} &\geq 0 \\
\theta_{\mathbf{v}+\mathbf{e}_a} - \theta_{\mathbf{v}} - \theta_{\mathbf{v}+\mathbf{e}_a+\mathbf{e}_b} + \theta_{\mathbf{v}+\mathbf{e}_a} &\geq 0 \\
\theta_{\mathbf{v}+\mathbf{e}_b+\mathbf{e}_a} - \theta_{\mathbf{v}+\mathbf{e}_b} - \theta_{\mathbf{v}+\mathbf{e}_b} + \theta_{\mathbf{v}} &\geq 0 \\
\theta_{\mathbf{v}+\mathbf{e}_b+\mathbf{e}_a} - \theta_{\mathbf{v}+\mathbf{e}_b} - \theta_{\mathbf{v}+\mathbf{e}_a+\mathbf{e}_b} + \theta_{\mathbf{v}+\mathbf{e}_a} &\geq 0
\end{aligned}$$

This set of inequalities is equivalent to the two given in (6). □

12.3.3. Sufficient Conditions for Monotonic Dominance for Calibrated Lattice Models

Proposition 3. *Let $f : \mathbb{R}^D \to \mathbb{R}$ be differentiable, $a, b \in [D]$, and for each $d \in [D]$ let $c_d : \mathbb{R} \to \mathbb{R}$ be differentiable. Denote by $g : \mathbb{R}^D \to \mathbb{R}$ the composed function given by $g(\mathbf{x}) = f(c_1(\mathbf{x}[1]), \ldots, c_D(\mathbf{x}[D]))$. Then if feature a monotonically dominates feature b in f and $c'_a(x_a) \geq c'_b(x_b) \geq 0$ for any choice of $x_a \in \mathbb{R}$ and any choice of $x_b \in \mathbb{R}$, then feature a monotonically dominates feature b in g.*

Proof. Readily follows from the chain rule and our assumptions on f, c'_a, c'_b. □

Corollary 3. *Let f be a calibrated lattice with a lattice function g and piecewise linear calibrators $\{c_d\}$ with each calibrator c_d having K_d knots and $K_d - 1$ slopes $\gamma_{d,k}$. Let $a, b \in [D]$ and assume that for all $d \in a, b$ and $k \in [K_d - 1]$, $\gamma_{d,k} \geq 0$. If feature a monotonically dominates feature b in g and for any choice of $k \in [K_a - 1]$ and $j \in [K_b - 1]$ it holds that $\gamma_{a,k} \geq \gamma_{b,j}$, then feature a monotonically dominates feature b in f.*

Proof. Follows directly from Proposition 3. □

12.3.4. Sufficient Conditions for Monotonic Dominance Under Monotone Transformations

Proposition 4. *Let $f : \mathbb{R}^D \to \mathbb{R}$ be a D dimensional function. Assume that for $a, b \in [D]$ feature a monotonically dominates feature b in f. Let $g : \mathbb{R} \to \mathbb{R}$ be an increasing differentiable function. Then feature a monotonically dominates feature b in $g(f(\mathbf{x}))$.*

Proof. By the chain rule $\partial_d(g(f(\mathbf{x}))) = g'(f(\mathbf{x}))\partial_d f(\mathbf{x})$, and since $g'(f(\mathbf{x})) \geq 0$ by assumption, the result holds. □

12.3.5. Sufficient Conditions for Monotonic Dominance For Ensembles

Proposition 5. *Let $f = \sum_{t=1}^{T} f_t$ be a D-dimensional ensemble of T D-dimensional models. Let $a, b \in [D]$. If feature a monotonically dominates feature b in every base model f_t, for $t = 1, \ldots, T$, then feature a monotonically dominates feature b in f.*

Proof. By our assumption for each $t \in [T]$, $\partial_a f_t \geq \partial_b f_t \geq 0$. The result follows from the fact that for $d \in \{a, b\}$, $\partial_d f = \sum_{t=1}^{T} \partial_d f_t$. □

Note that in practice when building ensembles with random subsets of features that in order to ensure each of the base models actually does satisfy monotonic dominance of a over b, one must ensure that every base model f_t that depends on input b also depends on input a, formally:

Proposition 6. *If a function f depends on b and is monotonically increasing in b, but does not depend on a, then a cannot be monotonically dominant over b in f_t.*

Proof. Monotonic dominance requires $\partial_a f_t \geq \partial_b f_t$, but if f_t does not depend on a, then $\partial_a f_t = 0$, and monotonic dominance cannot hold unless f_t also has no dependence on b. □

12.3.6. Monotonic Dominance is Transitive

Proposition 7. *If f is monotonically increasing in a, b, and c, and if a dominates b, and b dominates c, then a dominates c.*

Proof. By assumption, $\frac{\partial f(x)}{\partial \mathbf{x}[a]} \geq \frac{\partial f(x)}{\partial \mathbf{x}[b]}$ and $\frac{\partial f(x)}{\partial \mathbf{x}[b]} \geq \frac{\partial f(x)}{\partial \mathbf{x}[c]}$, thus it must be that $\frac{\partial f(x)}{\partial \mathbf{x}[a]} \geq \frac{\partial f(x)}{\partial \mathbf{x}[c]}$. □

12.3.7. Monotonic Dominance in Angles

In Fig. 1, we illustrate that monotonic dominance is increasing in the diagonal direction in each $a - b$ slice. Since monotonic dominance requires the function to also be increasing in a and b, it's also increasing in other directions, in fact, as shown below, it's increasing in any direction in the $[-\pi/4, \pi/2]$ cone.

Proposition 8. *Let $f : \mathbb{R}^D \to \mathbb{R}$ be a D-dimensional differentiable function, and let $a, b \in [D]$. Then the following statements are equivalent.*

1. *Feature a monotonically dominates feature b.*
2. *Every slice $g : \mathbb{R}^2 \to \mathbb{R}$ formed from f by fixing all the input features except features a and b, is increasing along directions whose angle with the $\mathbf{x}[a]$-axis is in $[-\pi/4, \pi/2]$. That is, for every $\mathbf{x} \in \mathbb{R}^D$, and $\Delta_a, \Delta_b \in \mathbb{R}$ with $\Delta_a \geq \max\{0, -\Delta_b\}$, it holds that $f(\mathbf{x}) \leq f(\mathbf{x} + \Delta_a \mathbf{e}_a + \Delta_b \mathbf{e}_b)$.*

Proof. **1 ⇒ 2**. Let $h : [0, 1] \to \mathbb{R}$ be the function given by: $h(t) = f(\mathbf{x} + t\Delta_a \mathbf{e}_a + t\Delta_b \mathbf{e}_b)$. By the chain rule we have $dh/dt = \Delta_a \partial f/\partial \mathbf{x}[a] + \Delta_b \partial f/\partial \mathbf{x}[b]$. Since a monotonically dominates b, $\Delta_a \geq 0$ and $\Delta_b \geq -\Delta_a$, it follows that $dh/dt \geq \Delta_a \partial f/\partial \mathbf{x}[b] - \Delta_a \partial f/\partial \mathbf{x}[b] = 0$. Thus h is increasing, and in particular $h(0) \leq h(1)$, which implies statement 2.

$\mathbf{2} \Rightarrow \mathbf{1}$. Fix a positive $\Delta \in \mathbb{R}$. Since we assume statement 2 holds, we have $f(\mathbf{x} + 0 \cdot \mathbf{e}_a + \Delta\mathbf{e}_b) \geq f(\mathbf{x})$. Thus

$$\frac{f(\mathbf{x} + \Delta\mathbf{e}_b) - f(\mathbf{x})}{\Delta} \geq 0. \tag{7}$$

Statement 2 also implies that $f(\mathbf{x} + \Delta\mathbf{e}_b) \leq f(\mathbf{x} + \Delta\mathbf{e}_b + \Delta\mathbf{e}_a - \Delta\mathbf{e}_b)$. Thus

$$\frac{f(\mathbf{x} + \Delta\mathbf{e}_a) - f(\mathbf{x})}{\Delta} \geq \frac{f(\mathbf{x} + \Delta\mathbf{e}_b) - f(\mathbf{x})}{\Delta} \tag{8}$$

Taking the limit of both sides of (7) and (8) as $\Delta \to 0$, it follows that $\partial f/\partial \mathbf{x}[a] \geq \partial f/\partial \mathbf{x}[b] \geq 0$. □

12.3.8. Range Dominance Results

We prove the following properties for range dominance.

12.3.9. Range Dominance Criteria for GAMs

Proposition 9. *Let $f(\mathbf{x}) = \sum_{d=1}^{D} f_d(\mathbf{x}[d])$ be a GAM increasing in features a,b. Then feature a range dominates feature b in f iff $f_a(u_a) - f_a(\ell_a) \geq f_b(u_b) - f_b(\ell_b)$.*

Proof. For $\mathbf{x} \in \mathbb{R}^D$, recall that a GAM satisfies $f(x) = \sum_{d=1}^{D} f_d(x)$, so for $d \in [D]$, $f(U_d(\mathbf{x})) - f(L_d(\mathbf{x})) = f_d(u_d) - f_d(\ell_d)$. The result follows. □

Corollary 4. *Let $f : \mathbb{R}^D \to \mathbb{R}$ be a D-dimensional linear model with parameters $\{\alpha_d\}_{d=0}^{D}$. Assume that f is increasing in features a,b for $a, b \in [D]$. Then feature a range dominates feature b iff $\alpha_a(u_a - \ell_a) \geq \alpha_b(u_b - \ell_b)$.*

Proof. Linear models are a special case of a GAM with $f_d(x) = \alpha_d x + \alpha_0/D$, thus the result follows from Proposition 9. □

Corollary 5. *Consider the GAM-PLF specified by $f(\mathbf{x}) = \sum_{d=1}^{D} \mathrm{PLF}(\mathbf{x}[d]; \{(\xi_{d,k}, \beta_{d,jk})\}_{k=0}^{K_d})$. Assume that f is increasing in features a and b. Then feature a range dominates feature b iff $\beta_{a,K_a} - \beta_{a,0} \geq \beta_{b,K_b} - \beta_{b,0}$.*

Proof. The GAM-PLF is a special case of the GAM, and the result follows directly from Proposition 9. □

12.3.10. Range Dominance Criteria for Lattice Models

Proposition 10. *Let f be a function for a lattice with size $\mathbf{V}$, dimension D, and vertex values $\{\theta_{\mathbf{v}}\}$ Let $f : \mathbb{R}^D \to \mathbb{R}$ be a function for a lattice of size $\mathbf{V}$, dimension d and vertex values $\{\theta_{\mathbf{v}}\}$. Suppose f is monotonically increasing in features a and b. Then feature a range dominates feature b in f iff for all $\mathbf{v} \in \mathcal{M}_{\mathbf{V}}$,*

$$\theta_{U_a(\mathbf{v})} - \theta_{L_a(\mathbf{v})} \geq \theta_{U_b(\mathbf{v})} - \theta_{L_b(\mathbf{v})}. \tag{9}$$

Proof. Clearly, the proposition's condition is necessary because it is the definition of range dominance given in Table 1 applied to each vertex $\mathbf{v}$ in the grid $\mathcal{M}_{\mathbf{V}}$. It remains to show that satisfying range dominance for each vertex in the grid is sufficient for it to be satisfied for every input $\mathbf{x} \in \overline{\mathcal{M}_{\mathbf{V}}}$.

From (1), the range of possible outputs as you vary $\mathbf{x}[d]$ is:

$$f(U_d(\mathbf{x})) - f(L_d(\mathbf{x})) = \sum_{\mathbf{v} \in \mathcal{N}(U_d(\mathbf{x}))} \Phi_{\mathbf{v}}(U_d(\mathbf{x}))\theta_{\mathbf{v}} - \sum_{\mathbf{v} \in \mathcal{N}(L_d(\mathbf{x}))} \Phi_{\mathbf{v}}(L_d(\mathbf{x}))\theta_{\mathbf{v}}.$$

Note that for $\mathbf{v} \in \mathcal{N}(L_d(\mathbf{x}))$, $\mathbf{v}[d]$ is either 0 or 1. Moreover, since $L_d(\mathbf{x})[d] = 0$, it follows from (2) that if $\mathbf{v}[d] = 1$, then $\Phi_{\mathbf{v}}(L_d(\mathbf{x})) = 0$. A similar argument shows that for $\mathbf{v} \in \mathcal{N}(U_d(\mathbf{x}))$, if $\mathbf{v}[d] = \mathbf{V}[d]$ then $\Phi_{\mathbf{v}}(U_d(\mathbf{x})) = 0$. Thus we have

$$f(U_d(\mathbf{x})) - f(L_d(\mathbf{x})) = \sum_{\substack{\mathbf{v} \in \mathcal{N}(U_d(\mathbf{x})) \\ \mathbf{v}[d] = \mathbf{V}[d]-1}} \Phi_{\mathbf{v}}(U_d(\mathbf{x}))\theta_{\mathbf{v}} - \sum_{\substack{\mathbf{v} \in \mathcal{N}(L_d(\mathbf{x})) \\ \mathbf{v}[d] = 0}} \Phi_{\mathbf{v}}(L_d(\mathbf{x}))\theta_{\mathbf{v}}.$$

Also, observe that the mapping $\mathbf{v} \mapsto L_d(\mathbf{v})$ is a bijection from $\{\mathbf{v} \in \mathcal{N}(\mathbf{x}) : \mathbf{v}[d] = \lfloor \mathbf{x}[d] \rfloor\}$ to $\{\mathbf{v} \in \mathcal{N}(L_d(\mathbf{x})) : \mathbf{v}[d] = 0\}$, and similarly the mapping $\mathbf{v} \mapsto U_d(\mathbf{v})$ is a bijection from $\{\mathbf{v} \in \mathcal{N}(\mathbf{x}) : \mathbf{v}[d] = \lfloor \mathbf{x}[d] \rfloor\}$ to $\{\mathbf{v} \in \mathcal{N}(U_d(\mathbf{x})) : \mathbf{v}[d] = \mathbf{V}[d] - 1\}$. Thus, we can change the sets we sum over in the two sums to a common set and obtain

$$f(U_d(\mathbf{x})) - f(L_d(\mathbf{x})) = \sum_{\substack{\mathbf{v} \in \mathcal{N}(\mathbf{x}) \\ \mathbf{v}[d] = \lfloor \mathbf{x}[d] \rfloor}} \Big(\Phi_{U_d(\mathbf{v})}(U_d(\mathbf{x}))\theta_{U_d(\mathbf{v})} - \Phi_{L_d(\mathbf{v})}(L_d(\mathbf{x}))\theta_{L_d(\mathbf{v})} \Big). \tag{10}$$

Now, for $v \in \mathbb{N}$, $x \in \mathbb{R}$, let $\lambda(v, x) = 1 + (x - v) \cdot (-1)^{I_{v=x}}$. Thus, for $\mathbf{v} \in \mathcal{N}(\mathbf{x})$ with $\mathbf{v}[d] = \lfloor \mathbf{x}[d] \rfloor$, it follows from (2) that $\phi_{L_d(\mathbf{v})}(L_d(\mathbf{x})) = \prod_{i=1}^{D} \lambda\big(L_d(\mathbf{v})[i], L_d(\mathbf{x})[i]\big)$. Clearly for $i \neq d$, we have $L_d(\mathbf{x})[i] = \mathbf{x}[i]$, $L_d(\mathbf{v})[i] = \mathbf{v}[i]$. Also, note that $\lambda\big(L_d(\mathbf{v})[d], L_d(\mathbf{x})[d]\big) = 1 = \lambda(\mathbf{v}[d], \mathbf{x}[d]) + \lambda(\mathbf{v}[d]+1, \mathbf{x}[d])$. Thus,

$$\begin{aligned} \phi_{L_d(\mathbf{v})}(L_d(\mathbf{x})) &= \big(\lambda(\mathbf{v}[d], \mathbf{x}[d]) + \lambda(\mathbf{v}[d]+1, \mathbf{x}[d])\big) \cdot \prod_{i \neq d} \lambda(\mathbf{v}[i], \mathbf{x}[i]) \\ &= \Phi_{\mathbf{v}}(\mathbf{x}) + \Phi_{\mathbf{v}+\mathbf{e}_d}(\mathbf{x}). \end{aligned} \tag{11}$$

A similar argument also shows that

$$\Phi_{U_d(\mathbf{v})}(U_d(\mathbf{x})) = \Phi_{\mathbf{v}}(\mathbf{x}) + \Phi_{\mathbf{v}+\mathbf{e}_d}(\mathbf{x}). \tag{12}$$

Plugging (11) and (12) into (10), we get

$$f(U_d(\mathbf{x})) - f(L_d(\mathbf{x})) = \sum_{\substack{\mathbf{v}\in\mathcal{N}(\mathbf{x})\\ \mathbf{v}[d]=\lfloor \mathbf{x}[d]\rfloor}} (\Phi_{\mathbf{v}}(\mathbf{x}) + \Phi_{\mathbf{v}+\mathbf{e}_d}(\mathbf{x}))\cdot (\theta_{U_d(\mathbf{v})} - \theta_{L_d(\mathbf{v})}) = \sum_{\mathbf{v}\in\mathcal{N}(\mathbf{x})} \Phi_{\mathbf{v}}(\mathbf{x})(\theta_{U_d(\mathbf{v})} - \theta_{L_d(\mathbf{v})}), \tag{13}$$

which gives the range for direction d. To prove range dominance we need to show that the range for direction a is at least the range for direction b, or equivalently that the difference $\big(f(U_a(\mathbf{x}))-f(L_a(\mathbf{x}))\big) - \big(f(U_b(\mathbf{x}))-f(L_b(\mathbf{x}))\big)$ is nonnegative. Substituting a and b for d in (13), we have that that difference is

$$\sum_{\mathbf{v}\in\mathcal{N}(\mathbf{x})} \Phi_{\mathbf{v}}(\mathbf{x})\big((\theta_{U_a(\mathbf{v})}-\theta_{L_a(\mathbf{v})}) - (\theta_{U_b(\mathbf{v})}-\theta_{L_b(\mathbf{v})})\big).$$

By (2) and (9) every term in this sum is nonnegative and thus the result follows. □

12.3.11. Sufficient Conditions for Range Dominance for Calibrated Lattice Models

Proposition 11. *Let f be a D-dimensional function given by $f(\mathbf{x}) = g(c_1(\mathbf{x}), \ldots, c_D(\mathbf{x}))$, where $g : \mathbb{R}^D \to \mathbb{R}$ is a D-dimensional function and each $c_i : \mathbb{R} \to \mathbb{R}$ is a 1-dimensional function. Fix $a, b \in [D]$. For $i \in \{a, b\}$, denote by $[\ell_i^{(g)}, u_i^{(g)}]$ and $[\ell^{(c_i)}, u^{(c_i)}]$ the domains of feature i of g and f, respectively and assume that c_i is increasing and satisfies $c_i(\ell^{(c_i)}) = \ell_i^{(g)}$ and $c_i(u^{(c_i)}) = u_i^{(g)}$. If feature a range dominates feature b in g, then feature a range dominates feature b in f.*

Proof. Assume that feature a range dominates feature b in g. Since c_a and c_b are increasing and g is increasing in features a and b then f is also increasing in features a and b. For a vector $\mathbf{x} \in \mathbb{R}^D$, denote by $\mathbf{c}(\mathbf{x})$ the vector $(c_1(x[1]), \ldots, c_D(x[D]))$ and, for $i \in [D]$, by $L_i^{(g)}(\mathbf{x})$, $U_i^{(g)}(\mathbf{x})$, the vectors obtained from $\mathbf{x}$ by replacing its ith entry with $\ell_i^{(g)}$ and $u^{(g)_i}$, respectively. Now, let $\mathbf{x} \in \mathbb{R}$. By our assumptions on c_a and c_b, $f(L_i(\mathbf{x})) = g(\mathbf{c}(L_i(\mathbf{x}))) = g(L_i^{(g)}(\mathbf{c}(\mathbf{x})))$ and $f(U_i(\mathbf{x})) = g(\mathbf{c}(U_i(\mathbf{x}))) = g(U_i^{(g)}(\mathbf{c}(\mathbf{x})))$. Thus $f(U_i(\mathbf{x})) - f(L_i(\mathbf{x})) = g(U_i^{(g)}(\mathbf{c}(\mathbf{x}))) - g(L_i^{(g)}(\mathbf{c}(\mathbf{x})))$. Since feature a range dominates feature b in g it follows that $f(U_a(\mathbf{x})) - f(L_a(\mathbf{x})) \geq f(U_b(\mathbf{x})) - f(L_b(\mathbf{x}))$ □

Corollary 6. *Let f be a calibrated lattice given by $f(\mathbf{x}) = g(c_1(x[1]), \ldots, c_D[\mathbf{x}[D]])$. If feature a range dominates feature b in g for $a, b \in [D]$ and c_a and c_b are increasing and capped, then feature a range dominates feature b in f.*

Proof. Follows directly from the Prop 11. □

12.3.12. Sufficient Conditions for Range Dominance Under Composition

Proposition 12. *Let $f : \mathbb{R}^D \to \mathbb{R}$ be a D-dimensional function, and assume that for $a, b \in [D]$ feature a range dominates feature b in f. Let $g(x) = sx + t$ be a 1-dimensional affine function for $s, t \in \mathbb{R}$ with $s \geq 0$. Then feature a range dominates feature b in $g(f(\cdot))$.*

Proof. Since g is increasing and f is monotonic in features a and b, $g(f(\cdot))$ is also increasing in features a and b. Let $\mathbf{x} \in \mathbb{R}^D$ and $i \in [D]$. Clearly, $g(f(U_i(\mathbf{x}))) - g(f(L_i(\mathbf{x}))) = s(f(U_i(\mathbf{x})) - f(L_i(\mathbf{x})))$. The result follows. □

12.3.13. Sufficient Conditions for Range Dominance for Ensembles

Proposition 13. *Let $\{f_t\}_{t=1}^T$, with $f_i : \mathbb{R}^D \to \mathbb{R}$ for $i \in [T]$ be a set of T D-dimensional functions, and let $a, b \in [D]$. If feature a range dominates feature b in each f_i, then feature a range dominates feature b in the ensemble $f = \sum_{i=1}^T f_i$.*

Proof. Since each f_i is monotonic in features a and b, f is also monotonic in features a and b. Let $\mathbf{x} \in \mathbb{R}^D$ and $i \in [D]$. The result follows since $f(U_i(\mathbf{x})) - f(L_i(\mathbf{x})) = \sum_i (f_i(U_i(\mathbf{x})) - f_i(L_i(\mathbf{x})))$. □

12.4. Edgeworth Results

Proofs for most of the Edgeworth properties given in this paper can be found in Cotter et al. (2019a), with the exception of the following new results.

12.4.1. Edgeworth Is Not Transitive

Proposition 14. *Suppose f satisfies an Edgeworth constraint on features a and b, and also satisfies an Edgeworth constraint on features a and c, then it cannot be concluded that f satisfies an Edgeworth constraint on a and b.*

Proof. Proof by counterexample: consider the second-degree polynomial function $f(x) = x[a]x[b] + x[a]x[c] - +x[b]\sin(x[c])$. This function satisfies Edgeworth on inputs a and b, and satsifies Edgeworth on b and c, but does not satisfy Edgeworth on b and c. □

12.4.2. EDGEWORTH ENSEMBLES

It was already shown in Cotter et al. (2019a) that an ensemble of base models that satisfy Edgeworth also satisfies Edgeworth (that is, the set of Edgeworth-constrained functions is closed under linearity). However, we next prove a further point that differs between Edgeworth and Trapezoid constraints and is important in the practice when building ensembles with random subsets of features: whether each base model must be a function of the a & b features in order for the constraints to hold. Specifically, for the ensemble to satisfy Edgeworth, we do not need to ensure that each base model has access to inputs a or b:

Proposition 15. *If a function f does not depend on at least one of the inputs a, b then it satisfies the Edgeworth constraint on features a and b.*

Proof. Case 1: f does not depend on input a, such that $\frac{\partial f(x)}{\partial x[a]} = 0$. Then since the partial of 0 is 0, the Edgeworth constraint is satisfied. Case 2: f does not depend on input b, such that $\frac{\partial f(x)}{\partial x[b]} = 0$. Then $\frac{\partial}{\partial x[b]} \frac{\partial f(x)}{\partial x[a]} = 0$. Case 3: f does not depend on input a or input b, again the Edgeworth constraint holds with equality. □

12.5. Trapezoid Results

Proofs for most of the Trapezoid properties given in this paper can be found in Cotter et al. (2019a), with the exception of the following new results.

12.5.1. TRAPEZOID IS NOT TRANSITIVE

Proposition 16. *Suppose f satisfies a Trapezoid constraint on features a and b, and also satisfies a Trapezoid constraint on features a and c, then it cannot be concluded that f satisfies a Trapezoid constraint on b and c.*

Proof. Due to the asymmetry of the trapezoid constraint definition and its requirement that the primary feature be monotonic, transitivity is not even well-defined in the general case, as it could be that neither input b and c is monotonic. □

12.5.2. TRAPEZOID ENSEMBLES

It was already shown in Cotter et al. (2019a) that an ensemble of base models that each satisfy Trapezoid also satisfies Trapezoid (that is, the set of Trapezoid-constrained functions is closed under linearity). However, we next prove a further point that differs between Edgeworth and Trapezoid constraints and is important in the practice when building ensembles with random subsets of features: whether each base model must be a function of the a & b features in order for the constraints to hold. Specifically, for the ensemble to satisfy Trapezoid, in order for each of the base models to satisfy Trapezoid, one must ensure that any base model that includes the conditioning feature b *also* includes the primary monotonic feature a:

Proposition 17. *A function $f(x)$ may not satisfy the Trapezoid constraint on features a and b and be monotonic increasing (or decreasing) on b.*

Proof. Trapezoid requires that f is monotonically increasing in b when a is maximal and monotonically decreasing in b when a is minimal. Thus f cannot be either monotonically increasing in b or monotonically decreasing on b for every choice of the other features. □

Proposition 18. *A function $f(x)$ may not satisfy the Trapezoid constraint on features a and b if f depends on b but does not depend on a.*

Proof. Trapezoid requires that f is monotonically increasing in b when a is maximal and monotonically decreasing in b when a is minimal, but if f does not depend on a it cannot be increasing/decreasing conditioned on a. □

12.6. Joint Monotonicity Results

12.6.1. JOINT MONOTONICITY IS NOT TRANSITIVE

Proposition 19. *Suppose f satisfies a joint monotonicity constraint on features a and b, and also satisfies a joint monotonicity constraint on features a and c, then it cannot be concluded that f satisfies a joint monotonicity constraint on b and c.*

Proof. Note that f can be jointly monotonic in a and b even if f is monotonically decreasing in b. Similarly, f could be jointly monotonic in a and c but be monotonically decreasing in c. However, if f is monotonically decreasing in b and c, then $f'_b + f'_c \leq 0$, and thus f is not jointly monotonic in b and c. □

12.6.2. JOINT MONOTONICITY FOR GAMS AND LINEAR MODELS

Proposition 20. *Let $f(\mathbf{x}) = \sum_{i=1}^{D} f_i(\mathbf{x}[i])$ be a GAM. Then f is jointly monotonic in features a and b iff $f'_a + f'_b \geq 0$*

Proof. Readily follows from the fact that $\partial_i f = f'_i$. □

Corollary 7. *Let $f : \mathbb{R}^D \to \mathbb{R}$ be a D-dimensional linear model with parameters $\{\alpha_i\}_{i=0}^{D}$. Then f is jointly monotonic in features a and b iff iff $\alpha_a + \alpha_b \geq 0$.*

Proof. Note that f can be regarded as a GAM with $f_i(x) = \alpha_i x + \alpha_0/D$. The result now follows from Proposition 20. □

Corollary 8. *Let $f(\mathbf{x}) = \sum_{i=1}^{D} \text{PLF}(\mathbf{x}[i]; \{(\xi_{i,k}, \beta_{i,jk})\}_{k=0}^{K_i})$ be a GAM-PLF. Then f is jointly monotonic in features a and b iff $\gamma_{a,j} + \gamma_{b,k} \geq 0$ for all $j \in [K_a]$,$k \in [K_b]$.*

Proof. Result follows as a special case of Proposition 20. □

12.6.3. Joint Monotonicity for Lattice Models

Proposition 21. *Let f be the function of a lattice with size $\mathbf{V}$, dimension D, and vertex values $\{\theta_{\mathbf{v}}\}$. Let $a, b \in [D]$. Then f is jointly monotonic in features a, b iff for all $\mathbf{v} \in \mathcal{M}_{\mathbf{V}}$ with $\mathbf{v}[a] \leq \mathbf{V}[a] - 2, \mathbf{v}[b] \leq \mathbf{V}[b] - 2$,*

$$\theta_{\mathbf{v}} \leq \frac{\theta_{\mathbf{v}+\mathbf{e}_a} + \theta_{\mathbf{v}+\mathbf{e}_b}}{2} \leq \theta_{\mathbf{v}+\mathbf{e}_a+\mathbf{e}_b}. \tag{14}$$

Proof. By Lemma 1 we have

$$\begin{aligned}&\partial_a f(\mathbf{x}) + \partial_b f(\mathbf{x}) \\ &= \sum_{\mathbf{v} \in \mathcal{N}(\mathbf{x})} \Phi_{\mathbf{v}}(\mathbf{x})(\theta_{\lceil \mathbf{v} \rceil_{a,\mathbf{x}}} - \theta_{\lfloor \mathbf{v} \rfloor_{a,\mathbf{x}}} + \theta_{\lceil \mathbf{v} \rceil_{b,\mathbf{x}}} - \theta_{\lfloor \mathbf{v} \rfloor_{b,\mathbf{x}}})\end{aligned}$$

Thus in each cell, $\mathcal{N}(\mathbf{x})$, $\partial_a f + \partial_b f$ is a multilinear interpolation of particular differences of the lattice values. Thus Since the multilinear weights $\Phi_{\mathbf{v}}(\mathbf{x}) \geq 0$, to ensure $\partial_a f(\mathbf{x}) + \partial_b f(\mathbf{x}) \geq 0$ we must ensure that the lattice-value-differences are all non-negative, that is, for each $\mathbf{v} \in \mathcal{M}_{\mathbf{V}}$ with $\mathbf{v}[a] \leq \mathbf{V}[a] - 2, \mathbf{v}[b] \leq \mathbf{V}[b] - 2$, we require that:

$$\begin{aligned}\theta_{\mathbf{v}+\mathbf{e}_a} - \theta_{\mathbf{v}} + \theta_{\mathbf{v}+\mathbf{e}_b} - \theta_{\mathbf{v}} &\geq 0 \\ \theta_{\mathbf{v}+\mathbf{e}_a} - \theta_{\mathbf{v}} + \theta_{\mathbf{v}+\mathbf{e}_a+\mathbf{e}_b} - \theta_{\mathbf{v}+\mathbf{e}_a} &\geq 0 \\ \theta_{\mathbf{v}+\mathbf{e}_b+\mathbf{e}_a} - \theta_{\mathbf{v}+\mathbf{e}_b} + \theta_{\mathbf{v}+\mathbf{e}_b} - \theta_{\mathbf{v}} &\geq 0 \\ \theta_{\mathbf{v}+\mathbf{e}_b+\mathbf{e}_a} - \theta_{\mathbf{v}+\mathbf{e}_b} + \theta_{\mathbf{v}+\mathbf{e}_a+\mathbf{e}_b} - \theta_{\mathbf{v}+\mathbf{e}_a} &\geq 0.\end{aligned}$$

Simplifying these inequalities row-by-row becomes:

$$\begin{aligned}\theta_{\mathbf{v}+\mathbf{e}_a} + \theta_{\mathbf{v}+\mathbf{e}_b} &\geq 0 \\ \theta_{\mathbf{v}+\mathbf{e}_a+\mathbf{e}_b} - \theta_{\mathbf{v}} &\geq 0 \\ \theta_{\mathbf{v}+\mathbf{e}_a+\mathbf{e}_b} - \theta_{\mathbf{v}} &\geq 0 \\ 2\theta_{\mathbf{v}+\mathbf{e}_b+\mathbf{e}_a} - \theta_{\mathbf{v}+\mathbf{e}_b} - \theta_{\mathbf{v}+\mathbf{e}_a} &\geq 0.\end{aligned}$$

which can be further simplified to (14). □

Proposition 22. *Let $f = g(c_1(\mathbf{x}), \ldots, c_D(\mathbf{x}))$ be a D-dimensional function, where $g : \mathbb{R}^D \to \mathbb{R}$ is a D-dimensional function and each $c_i : \mathbb{R} \to \mathbb{R}$ is a 1-dimensional function. If for $a, b \in [D]$, g is jointly monotonic in features a and b, and c_a and c_b are both affine with the same nonnegative slope then f is also jointly monotonic in features a and b.*

Proof. Let $s \geq 0$ be the slope of c_a and c_b. By the chain rule, $\partial_a f + \partial_b f = s(\partial_a g + \partial_b g) \geq 0$. □

Corollary 9. *Let f be a calibrated lattice with lattice function g and calibrators $\{c_i\}_{i \in [D]}$. Assume that for $a, b \in [D]$, the calibrators c_a and c_b are affine with the same nonnegative slope. Then if g is jointly monotonic in features a and b, so is f.*

Proof. Follows as a special case of Proposition 22. □

Proposition 23. *Let $f : \mathbb{R}^D \to \mathbb{R}$ be a D-dimensional function jointly monotonic in features $a, b \in [D]$. Let $g : \mathbb{R} \to \mathbb{R}$ be a 1-dimensional increasing function. Then $g(f(\cdot))$ is jointly monotonic in features a, b.*

Proof. By the chain rule $\partial_a g(f(\mathbf{x})) + \partial_b g(f(\mathbf{x})) = g'(f(\mathbf{x}))(\partial_a f(\mathbf{x}) + \partial_b f(\mathbf{x}))$. By our assumptions on g and f the RHS is nonnegative. □

12.6.4. Joint Monotonicity and Ensembles

We show that joint monotonicity is closed under addition:

Proposition 24. *Let $\{f_t\}_{t=1}^{T}$, with $f_t : \mathbb{R}^D \to \mathbb{R}$ for $t \in [T]$ and let $a, b \in [D]$. If each f_t is jointly monotonic in a and b then so is f.*

Proof. $\partial_a f + \partial_b f = \sum_{t=1}^{T} (\partial_a f_t + \partial_b f_t)$. By our assumption, the RHS is non-negative. □

Note that when constraining an ensemble to be jointly monotonic on inputs a and b, if any base model f_t only takes one of the inputs, then the joint monotonicity constraint on the ensemble will force that base model to be monotonically increasing in the other input (see proposition below). This need not be a problem but should be taken into consideration.

Proposition 25. *For a function $f(x)$ to satisfy the joint monotonicity constraint on two features a and b where f does not depend on a, it must be that $f(x)$ is monotonically increasing in b.*

Proof. If f does not depend on a, its partial is zero by definition, such that $\partial_a f + \partial_b f = \partial_b f$, and thus to satisfy joint monotonicity it must be that $\partial_b f \geq 0$. □

12.7. Unimodality Results

For the sake of building intuition, we first present a simple proof for the one-dimensional case of a unimodal PLF. That proof can also be distilled as a special case from the proof for a multidimensional lattice that follows. We also provide a picture and example of a two-dimensional unimodal lattice model.

12.7.1. Proof for Unimodal PLF

First we show that a two-layer 1D model where the first layer is monotonically increasing and the second layer is unimodal is overall unimodal.

Proposition 26. *Let $f : \mathbb{R} \to \mathbb{R}$ be a 1-dimensional unimodal function. Let $c : \mathbb{R} \to \mathbb{R}$ be a 1-dimensional monontonically increasing function. Then $h(x) = f(c(x))$ is unimodal.*

Proof. To show that h is unimodal, we must show that there exists a minimizer in the domain of h, call it z^*, such that for $z \leq z^*$ the function h is decreasing, and for $z \geq z^*$ the function h is increasing.

Let z^* denote the minimizer of the unimodal function f, such that for $z \geq z^*$, $f(z)$ is increasing by the definition of unimodality. Similarly, for $z \leq z^*$, $f(z)$ is decreasing by the definition of unimodality.

Choose a x^* such that $c(x^*) = z^*$: if c is strictly monotonic, then c is invertible and $x^* = c^{-1}(z^*)$; if c is only non-decreasing, then x^* can be any value such that $c(x^*) = z^*$.

Then the domain $z \leq z^*$ maps to $x \leq x^*$ because c is monotonically increasing. And since f is decreasing for $z \leq z^*$, by the rules of composition $h(x)$ is a decreasing function for $x \leq x^*$. Analogously, for $z \geq z^*$, h is increasing. Thus h is unimodal. □

12.7.2. Ceterus Paribus Unimodality

Below we extend the definition for a multidimensional function f to satisfy the unimodal shape constraint to hold for just a subset of the features. To simplify notation we will assume that this subset is $[s]$ for some $s \in D$. This definition and related properties can be easily generalized to an arbitrary subset of features.

Let $f : \mathbb{R}^D \to \mathbb{R}$ be a D-dimensional function and $s \in [D]$. We say that f satisfies the unimodal shape constraint in features $1, 2, \ldots, s$ if every restriction $g : \mathbb{R}^s \to \mathbb{R}$ of f, obtained by fixing the features $s+1, s+2, \ldots, D$ in the input to f to some constants, satisfies the unimodal shape constraint on all its non-fixed features.

The following proposition characterizes lattices that satisfy the unimodal shape constraint in some subset of features with the minimizer fixed in the lattice "center".

Proposition 27. *Let $f : \mathbb{R}^D \to \mathbb{R}$ be the function of a D-dimensional lattice of size $\mathbf{V}$. Let $s \in [D]$ and assume that $\mathbf{V}[d]$ is odd for $d \leq s$. Define $\mathbf{o}^* \in \mathbb{R}^s$ by $\mathbf{o}^*[d] = (\mathbf{V}[d] - 1)/2$, for $d \in [s]$. The following statements are equivalent.*

1. *f satisfies the unimodal shape constraint in directions $1, \ldots, s$ with minimizer $\mathbf{o}^*$ for any restrictions on the last $D - s$ features.*
2. *For every $\mathbf{v} \in \mathcal{M}_{\mathbf{V}}$, $\delta_1, \ldots, \delta_s \in \{0, 1\}$ and all $d \in [s]$ such that $\mathbf{v} + \delta_d \mathbf{e}_d, \mathbf{v} - (1 - \delta_d)\mathbf{e}_d \in \mathcal{M}_{\mathbf{V}}$,*

$$\sum_{d=1}^{s} (\theta_{\mathbf{v}+\delta_d \mathbf{e}_d} - \theta_{\mathbf{v}-(1-\delta_d)\mathbf{e}_d})(\mathbf{v}[d] - \mathbf{o}^*[d]) \geq 0 \quad (15)$$

Proof. By the definition, f satisfies the unimodal shape constraint in directions 1,2,...,s with minimizer $\mathbf{o}^*$ if and only if every restriction obtained from f by fixing the last $D - s$ features is increasing along rays originating in $\mathbf{o}^*$. This condition can be equivalently restated as: for each $\mathbf{x} \in \overline{\mathcal{M}_{\mathbf{V}}}$, the function $h_{\mathbf{x}} : [0, 1] \to \mathbb{R}$, given by $h_{\mathbf{x}}(t) = f(\mathbf{r}_{\mathbf{x}}(t))$, with $\mathbf{r}_{\mathbf{x}}(t) = (\mathbf{o}^*[1] + t(\mathbf{x}[1] - \mathbf{o}^*[1]), \ldots, \mathbf{o}^*[s] + t(\mathbf{x}[s] - \mathbf{o}^*[s]), \mathbf{x}[s+1], \mathbf{x}[s+2], \ldots, \mathbf{x}[D])$, is increasing. Since each such $h_{\mathbf{x}}$ is continuous and piecewise-differentiable with finitely many pieces, the last condition is equivalent to requiring that $h'_{\mathbf{x}}(t) \geq 0$ for all $t \in [0, 1]$ where the derivative is defined. Observe that it's sufficient to require that for all $\mathbf{x} \in \overline{\mathcal{M}_{\mathbf{V}}}$, $h'_{\mathbf{x}}(1) \geq 0$, when its defined, since $h_{\mathbf{x}}(t) = h_{\mathbf{r}_{\mathbf{x}}(t)}(1)$. Therefore, statement 1 holds if and only if

$$\forall \mathbf{x} \in \overline{\mathcal{M}_{\mathbf{V}}}, h'_{\mathbf{x}}(1) \geq 0 \quad (16)$$

By the chain rule, $h'_{\mathbf{x}}(1) = \sum_{d=1}^{s} \partial_d f(\mathbf{x}) \cdot (\mathbf{x}[d] - \mathbf{o}^*[d])$ and hence using Lemma 1 we have

$$\begin{aligned} h'_{\mathbf{x}}(1) &= \sum_{d \in [s], \mathbf{v} \in \mathcal{N}(\mathbf{x})} \Phi_{\mathbf{v}}(\mathbf{x})(\theta_{\lceil \mathbf{v} \rceil_{d,\mathbf{x}}} - \theta_{\lfloor \mathbf{v} \rfloor_{d,\mathbf{x}}})(\mathbf{x}[d] - \mathbf{o}^*[d]) \\ &= \sum_{d,\mathbf{v}} \Phi_{\mathbf{v}}(\mathbf{x})(\theta_{\lceil \mathbf{v} \rceil_{d,\mathbf{x}}} - \theta_{\lfloor \mathbf{v} \rfloor_{d,\mathbf{x}}})(\mathbf{x}[d] - \lfloor \mathbf{x}[d] \rfloor) \\ &\quad + \sum_{d,\mathbf{v}} \Phi_{\mathbf{v}}(\mathbf{x})(\theta_{\lceil \mathbf{v} \rceil_{d,\mathbf{x}}} - \theta_{\lfloor \mathbf{v} \rfloor_{d,\mathbf{x}}})(\lfloor \mathbf{x}[d] \rfloor - \mathbf{o}^*[d]), \end{aligned}$$

where to get the last equality we added to and subtracted from each summand the quantity $\Phi_{\mathbf{v}}(\mathbf{x})(\theta_{\lceil \mathbf{v} \rceil_{d,\mathbf{x}}} - \theta_{\lfloor \mathbf{v} \rfloor_{d,\mathbf{x}}}) \lfloor \mathbf{x}[d] \rfloor$.

Next, for a fixed $d \in [D]$, partitioning the set $\mathcal{N}(\mathbf{x})$ of size 2^D into the 2^{D-1} pairs $\{(\mathbf{v}, \lfloor \mathbf{v} \rfloor_{d,\mathbf{x}}) : \mathbf{v} \in \mathcal{N}(\mathbf{x}), \mathbf{v} = \lceil \mathbf{v} \rceil_{d,\mathbf{x}}\}$, we regroup the terms in the summation and get

$$\begin{aligned} h'_{\mathbf{x}}(1) = \sum_{d, \mathbf{v}: \mathbf{v} = \lceil \mathbf{v} \rceil_{d,\mathbf{x}}} \Big(&(\Phi_{\mathbf{v}}(\mathbf{x}) + \Phi_{\lfloor \mathbf{v} \rfloor_{d,\mathbf{x}}}(\mathbf{x}))(\mathbf{x}[d] - \lfloor \mathbf{x}[d] \rfloor) \\ &\cdot (\theta_{\lceil \mathbf{v} \rceil_{d,\mathbf{x}}} - \theta_{\lfloor \mathbf{v} \rfloor_{d,\mathbf{x}}}) \Big) \\ + \sum_{d,\mathbf{v}} \Phi_{\mathbf{v}}(\mathbf{x})(\theta_{\lceil \mathbf{v} \rceil_{d,\mathbf{x}}} - \theta_{\lfloor \mathbf{v} \rfloor_{d,\mathbf{x}}})(\lfloor \mathbf{x}[d] \rfloor - \mathbf{o}^*[d]). \end{aligned} \quad (17)$$

Now, using (2) and defining $\lambda(v, x) = 1 + (x - v)(-1)^{I_{v=\lfloor x \rfloor}}$ for $x \in \mathbb{R}$ and $v \in \mathbb{N}$, we have for each $\mathbf{v} \in \mathcal{N}(\mathbf{x})$, with

$\mathbf{v}{=}\lceil\mathbf{v}\rceil_{d,\mathbf{x}}$

$$\begin{aligned}
&\Big(\Phi_{\mathbf{v}}(\mathbf{x}) + \Phi_{\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}}(\mathbf{x})\Big)(\mathbf{x}[d] - \lfloor\mathbf{x}[d]\rfloor) \\
&\quad = (\mathbf{x}[d] - \lfloor\mathbf{x}[d]\rfloor)\sum_{\mathbf{u}\in\{\mathbf{v},\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}\}}\prod_{i=1}^{D}\lambda(\mathbf{u}[i],\mathbf{x}[i]) \\
&\quad = (\mathbf{x}[d] - \lfloor\mathbf{x}[d]\rfloor)\Big(\prod_{i\neq d}\lambda(\mathbf{v}[i],\mathbf{x}[i])\Big)\cdot \\
&\qquad \Big(\lambda(\lfloor\mathbf{x}[d]\rfloor{+}1,\mathbf{x}[d]) + \lambda(\lfloor\mathbf{x}[d]\rfloor,\mathbf{x}[d])\Big),
\end{aligned}$$

where to get the last equality, observe that for $i \neq d$, the ith entry of $\mathbf{v}$ and $\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}$ is the same. Noting that $\lambda(\lfloor\mathbf{x}[d]\rfloor{+}1,\mathbf{x}[d]) + \lambda(\lfloor\mathbf{x}[d]\rfloor,\mathbf{x}[d]) = 1$ and that $\mathbf{x}[d] - \lfloor\mathbf{x}[d]\rfloor = \lambda(\mathbf{x}[d],\mathbf{v}[d])$, we get

$$\begin{aligned}
\Big(\Phi_{\mathbf{v}}(\mathbf{x}) + \Phi_{\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}}(\mathbf{x})\Big)(\mathbf{x}[d] - \lfloor\mathbf{x}[d]\rfloor) &= \prod_{i=1}^{D}\lambda(\mathbf{v}[i],\mathbf{x}[i]) \\
&= \Phi_{\mathbf{v}}(\mathbf{x}) \qquad (18)
\end{aligned}$$

Plugging (18) into (17), we get

$$\begin{aligned}
h'_{\mathbf{x}}(1) &= \sum_{d,\mathbf{v}:\mathbf{v}=\lceil\mathbf{v}\rceil_{d,\mathbf{x}}}\Phi_{\mathbf{v}}(\mathbf{x})(\theta_{\lceil\mathbf{v}\rceil_{d,\mathbf{x}}}{-}\theta_{\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}}) \\
&\quad + \sum_{d,\mathbf{v}}\Phi_{\mathbf{v}}(\mathbf{x})(\theta_{\lceil\mathbf{v}\rceil_{d,\mathbf{x}}}{-}\theta_{\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}})(\lfloor\mathbf{x}[d]\rfloor - \mathbf{o}^*[d]) \\
&= \sum_{d,\mathbf{v}}\Big(\Phi_{\mathbf{v}}(\mathbf{x})(\theta_{\lceil\mathbf{v}\rceil_{d,\mathbf{x}}}{-}\theta_{\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}}) \\
&\qquad \cdot (I_{\mathbf{v}=\lceil\mathbf{v}\rceil_{d,\mathbf{x}}}{+}\lfloor\mathbf{x}[d]\rfloor{-}\mathbf{o}^*[d])\Big) \\
&= \sum_{\mathbf{v}\in\mathcal{N}(\mathbf{x})}\Phi_{\mathbf{v}}(\mathbf{x})\sum_{d=1}^{s}(\theta_{\lceil\mathbf{v}\rceil_{d,\mathbf{x}}}{-}\theta_{\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}})(\mathbf{v}[d]{-}\mathbf{o}^*[d]),
\end{aligned}$$

where the last equality holds since, for each $\mathbf{v}{\in}\mathcal{N}(\mathbf{x})$, $I_{\mathbf{v}=\lceil\mathbf{v}\rceil_{d,\mathbf{x}}}{+}\lfloor\mathbf{x}[d]\rfloor{=}\mathbf{v}[d]$.

Hence $h'_{\mathbf{x}}(1)$ is a multilinear interpolation of the values $\{\sum_{d=1}^{s}(\theta_{\lceil\mathbf{v}\rceil_{d,\mathbf{x}}}{-}\theta_{\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}})(\mathbf{v}[d] - \mathbf{o}^*[d])\}_{\mathbf{v}}$ on $\mathcal{N}(\mathbf{x})$. Thus requiring that it would be nonnegative for all $\mathbf{x} \in \overline{\mathcal{M}_{\mathbf{V}}}$ is equivalent to requiring that

$$\sum_{d=1}^{s}(\theta_{\lceil\mathbf{v}\rceil_{d,\mathbf{x}}}{-}\theta_{\lfloor\mathbf{v}\rfloor_{d,\mathbf{x}}})(\mathbf{v}[d]{-}\mathbf{o}^*[d]) \geq 0,\ \forall\mathcal{N}(\mathbf{x}), \mathbf{v}{\in}\mathcal{N}(\mathbf{x}).$$

It's easy to verify that these are precisely the inequalities in Statement 2. □

Given a lattice function that satisfies the unimodal shape constraint, what are the conditions needed to be satisfied by calibrators so that the resulting calibrated lattice function also satisfies that constraint? One might think that it's sufficient that the calibrators for the features involved in the unimodality constraint be increasing. However, that is not true as the following example shows.

Let f be the function of the 2-dimensional lattice with size $(3, 3)$ and vertex values: $\theta_{1,1} = 0, \theta_{0,1} = \theta_{2,1} = 1, \theta_{1,0} = \theta_{1,2} = 3, \theta_{0,0} = \theta_{2,0} = \theta_{0,2} = \theta_{2,2} = 2$. It's easy to verify that Statement 2 of Proposition 27 holds. Thus f satisfies the unimodality shape constraint with minimizer $(1, 1)$. Now, let $c_1 : [0, 3] \to [0, 3]$ and $c_2 : [0, 3] \to [0, 3]$ be the calibrators given by:

$$c_1(x) = \begin{cases} x & \text{if } 0 \leq x < 1 \\ (x+1)/2 & \text{if } 1 \leq x \leq 3 \end{cases},$$

and

$$c_2(x) = \begin{cases} x & \text{if } 0 \leq x < 2 \\ 2 & \text{if } 2 \leq x \leq 3 \end{cases}.$$

Let $g(x, y) = f(c_1(x), c_2(y))$. Then it can be easily verified that the global minimum of g is at $(1, 1)$ and it is unique. Thus for g to satisfy the unimodal shape constraint, it must do so with minimizer $(1, 1)$. However g is not increasing along the ray $r(t) = (1, 1) + t(1, 1)$, since $g(r(1)) = g(2, 2) = f(3/2, 2) = (\theta_{2,2} + \theta_{1,2})/2 = 5/2$ and $g(r(2)) = g(3, 3) = f(2, 2) = \theta_{2,2} = 2$.

12.7.3. Strong Ceterus Paribus Unimodality

It turns out, however, that there is a stronger constraint that a function can satisfy (i.e., it implies the unimodal shape constaint), that is invariant under composition with increasing calibrators. We call this shape constraint *strong unimodality*. Essentially, it means that f is increasing in any direction that's parallel to an axis and is pointing away from the minimizer.

Strong unimodality: Let $f : \mathbb{R}^D \to \mathbb{R}$ be a D-dimensional differentiable[1] function and let $\mathbf{x}^* \in \mathbb{R}$ be the minimizer of f. We say that f satisfies the strong unimodality shape constraint with minimizer $\mathbf{x}^*$ if for all $\mathbf{x} \in \mathbb{R}^D$, $d \in [D]$, $\partial_d f(\mathbf{x}) \cdot (\mathbf{x}[d] - \mathbf{x}^*[d]) \geq 0$

The following proposition shows that this is indeed a stronger constraint than regular unimodality.

Proposition 28. *Let $f : \mathbb{R}^D \to \mathbb{R}$ be a D-dimensional differentiable function and $\mathbf{x}^* \in \mathbb{R}^D$ be its minimizer. If f satisfies the strong unimodality shape constraint with minimizer $\mathbf{x}^*$ then f satisfies the unimodality shape constraint with the same minimizer.*

Proof. We need to show that f is increasing along rays originating from $\mathbf{x}^*$. Since f is differentiable, that is equivalent to showing that for all $\mathbf{x} \in \mathbb{R}^D$ the directional derivative of

[1] One can generalize this definition to non differentiable functions using differences, but though the properties below still hold, their proofs become more complicated. To simplify the exposition we restrict to differentiable functions.

f in direction $\mathbf{x} - \mathbf{x}^*$ at $\mathbf{x}$ is nonnegative. That is we need to show $\sum_{d=0}^{D} \partial_d f(\mathbf{x})(\mathbf{x}[d] - \mathbf{x}^*[d]) \geq 0$, which holds by our assumption that each term in the above sum is nonnegative. □

Similarly to conventional unimodality, we extend this definition to allow specifying the constraint on only a subset of the features $1, 2, \ldots, s$ for some $s \in [D]$ by requiring that every restriction of f to the features $1, 2, \ldots, s$ satisfies the strong unimodal shape constraint with some minimizer. The following proposition characterizes lattices that satisfy the strong unimodal shape constraint in some subset of the features with the minimizer fixed in the lattice center.

Proposition 29. *Let $f : \mathbb{R}^D \to \mathbb{R}$ be the function of a D-dimensional lattice of size $\mathbf{V}$. Let $s \in [D]$ and assume that $\mathbf{V}[d]$ is odd for $d \leq s$. Define $\mathbf{o}^* \in \mathbb{R}^s$ by $\mathbf{o}^*[d] = (\mathbf{V}[d] - 1)/2$, for $d \in [s]$. Then the following statements are equivalent.*

1. *f satisfies the strong unimodal shape constraint in directions $1, \ldots, s$ with minimizer $\mathbf{o}^*$ for any restrictions on the last $D - s$ features.*
2. *For every $d \in [s]$ and $\mathbf{v} \in \mathcal{M}_{\mathbf{v}}$ such that $\mathbf{v} + \mathbf{e}_d \in \mathcal{M}_{\mathbf{v}}$*
   (a) If $\mathbf{v}[d] \geq \mathbf{o}^*[d]$ then $\theta_{\mathbf{v}} \leq \theta_{\mathbf{v}+\mathbf{e}_d}$.
   (b) If $\mathbf{v}[d] < \mathbf{o}^*[d]$ then $\theta_{\mathbf{v}} \geq \theta_{\mathbf{v}+\mathbf{e}_d}$.

Proof. **1 ⇒ 2**. Let $\mathbf{v} \in \mathcal{M}_{\mathbf{V}}$,$d \in [s]$ satisfy the conditions of Statement 2, and let g be the restriction of f obtained by fixing the last $D - s$ features to $\mathbf{v}[s+1], \ldots, \mathbf{v}[D]$. Let $\epsilon \in (0, 1)$ and denote by $\mathbf{v}_\epsilon = \mathbf{v} + \epsilon \mathbf{e}_d$. Since g satisfies the strong unimodality shape constraint with minimizer $\mathbf{o}^*$, we have $\partial_d f(\mathbf{v}_\epsilon)(\mathbf{v}_\epsilon[d] - \mathbf{o}^*[d]) \geq 0$. It follows from (1) that $\partial_d f(\mathbf{v}_\epsilon) = \theta_{\mathbf{v}+\mathbf{e}_d} - \theta_{\mathbf{v}}$. Thus,

$$(\theta_{\mathbf{v}+\mathbf{e}_d} - \theta_{\mathbf{v}})(\mathbf{v}_\epsilon[d] - \mathbf{o}^*[d]) \geq 0 \tag{19}$$

Now, if $\mathbf{v}[d] \geq \mathbf{o}^*[d]$ then $\mathbf{v}_\epsilon[d] - \mathbf{o}^*[d] > 0$ and thus it follows from (19) that $\theta_{\mathbf{v}+\mathbf{e}_d} \geq \theta_{\mathbf{v}}$. On the other hand, if $\mathbf{v}[d] < \mathbf{o}^*[d]$, then since $\epsilon < 1$, we get $\mathbf{v}_\epsilon[d] < \mathbf{o}^*[d]$; thus, it follows from (19) that $\theta_{\mathbf{v}+\mathbf{e}_d} \leq \theta_{\mathbf{v}}$. Therefore Statement 2 holds in every case.

2 ⇒ 1. Let $\mathbf{x} \in \overline{\mathcal{M}_{\mathbf{V}}}$, $d \in [s]$. We need to show that $\partial_d f(\mathbf{x})(\mathbf{x}[d] - \mathbf{o}^*[d]) \geq 0$. By Lemma 1 that inequality can be rewritten as

$$\sum_{\mathbf{v} \in \mathcal{N}(\mathbf{x})} \Phi_{\mathbf{v}}(\mathbf{x})(\theta_{\lceil \mathbf{v} \rceil_{d,\mathbf{x}}} - \theta_{\lfloor \mathbf{v} \rfloor_{d,\mathbf{x}}})(\mathbf{x}[d] - \mathbf{o}^*[d]) \geq 0 \tag{20}$$

If $\mathbf{x}[d] \geq \mathbf{o}^*[d]$, then since $\mathbf{o}^*$ has integer entries, we get $\lfloor \mathbf{x}[d] \rfloor \geq \mathbf{o}^*[d]$ and, by the condition in Statement 2, $\theta_{\lceil \mathbf{v} \rceil_{d,\mathbf{x}}} - \theta_{\lfloor \mathbf{v} \rfloor_{d,\mathbf{x}}} \geq 0$. Thus each term in the sum in (20) is nonnegative and hence (20) holds.

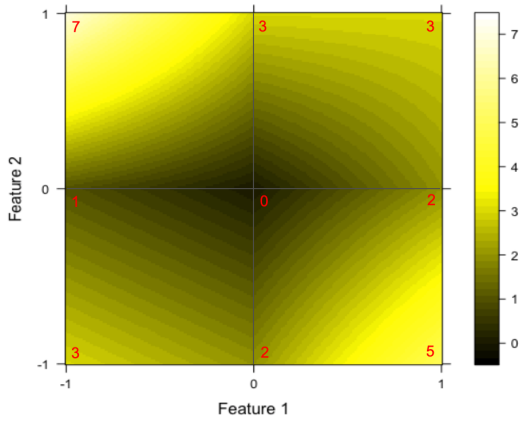

Figure 5. An illustrative 3^2 min-unimodal lattice l with the color representing the heatmap of model output $l(x)$. The nine numbers on the figure represents the model output at nine grid points, $\{-1, 0, 1\}^2$, which are the nine model parameters of the lattice model, i.e, $\theta_{[-1,-1]} \equiv l([-1,-1]) = 3$, $\theta_{[0,-1]} \equiv l([0,-1]) = 2$, $\theta_{[1,-1]} \equiv l([1,-1]) = 5$, $\theta_{[-1,0]} \equiv l([-1,0]) = 1$, $\theta_{[0,0]} \equiv l([0,0]) = 0$, $\theta_{[1,-0]} \equiv l([1,0]) = 2$, $\theta_{[-1,1]} \equiv l([-1,1]) = 7$, $\theta_{[0,1]} \equiv l([0,1]) = 3$ and $\theta_{[1,1]} \equiv l([1,1]) = 3$. Model output at $x \notin \{-1, 0, 1\}^2$ is interpolated based on the model parameter at adjacent grid points.

Similarly, if $\mathbf{x}[d] < \mathbf{o}^*[d]$ then $\lfloor \mathbf{x}[d] \rfloor < \mathbf{o}^*[d]$ and by the condition in Statement 2 we get $\theta_{\lceil \mathbf{v} \rceil_{d,\mathbf{x}}} - \theta_{\lfloor \mathbf{v} \rfloor_{d,\mathbf{x}}} \leq 0$. So again each term in the sum in (20) is nonnegative and thus (20) holds. □

Fig. 5 illustrates a 3×3 strong unimodal lattice model. The model has its minimizer at the center, and is increasing in any ray originating from the center.

For a function $f : \mathbb{R}^D \to \mathbb{R}$ and $s \in [D]$, we denote the restriction obtained from f by fixing the last $D - s$ inputs to $r_1, \ldots, r_{D-s} \in \mathbb{R}$ by $f|_{r_1,\ldots,r_{D-s}}$. That is $f|_{r_1,\ldots,r_{D-s}} : \mathbb{R}^s \to \mathbb{R}$ is given by $f|_{r_1,\ldots,r_{D-s}}(\mathbf{y}) = f(\mathbf{y}[1], \ldots, \mathbf{y}[s], r_1, \ldots, r_{D-s})$.

Proposition 30. *Let $f : \mathbb{R}^D \to \mathbb{R}$ be a D-dimensional function that satisfies the strong unimodality shape constraint in directions $1, 2, \ldots, s$. Let $\{c_d\}_{d=1}^D$ be 1-dimensional functions $c_d : \mathbb{R} \to \mathbb{R}$ with c_d increasing for $d = 1, \ldots, s$. Let $h : \mathbb{R}^D \to \mathbb{R}$ denote the composed function given by $h(\mathbf{x}) = f(\mathbf{c}(\mathbf{x}))$, where $\mathbf{c}(\mathbf{x}) = (c_1(\mathbf{x}[1]), \ldots, c_D(\mathbf{x}[D]))$. Assume that for every restriction $h|_{r_{s+1},\ldots,r_D}$ there exists a point $\mathbf{x}^* \in \mathbb{R}^s$ such that $(c_1(\mathbf{x}^*[1]), \ldots, c_s(\mathbf{x}^*[s]))$ is a minimizer for the restriction $f|_{c_{s+1}(r_{s+1}),\ldots,c_D(r_D)}$ of f. Then $h : \mathbb{R}^D \to \mathbb{R}$ also satisfies the strong unimodality constraint in directions $1, 2, \ldots, s$.*

Proof. Let $\mathbf{x} \in \mathbb{R}^D$. We'll show that $h|_{\mathbf{x}[s+1],\ldots,\mathbf{x}[D]}$ satisfies the strong unimodality constraint. Let $\mathbf{x}^* \in \mathbb{R}^s$

be such that $\mathbf{y}^* = (c_1(\mathbf{x}^*[1]), \ldots, c_s(\mathbf{x}^*[s]))$ is a minimizer of $f|_{c_{s+1}(\mathbf{x}[s+1]),\ldots,c_D(\mathbf{x}[D])}$. We need to show that $\partial h \cdot (\mathbf{x}[d] - \mathbf{x}^*[d]) \geq 0$ for all $d \in [s]$. By the chain rule $\partial_d h = \partial_d f(\mathbf{c}(\mathbf{x}))c'_d(\mathbf{x}[d])$. Thus we need to show that

$$\partial_d f(\mathbf{c}(\mathbf{x}))c'_d(\mathbf{x}[d])(\mathbf{x}[d] - \mathbf{x}^*[d]) \geq 0. \tag{21}$$

Since $f|_{c_{s+1}(\mathbf{x}[s+1]),\ldots,c_D(\mathbf{x}[D])}$ satisfies the strong unimodality constraint with minimizer $\mathbf{y}^*$ and $\mathbf{y}^*[d] = c_d(\mathbf{x}^*[d])$, we have

$$\partial_d f(\mathbf{c}(\mathbf{x}))(c_d(\mathbf{x}[d]) - c_d(\mathbf{x}^*[d])) \geq 0. \tag{22}$$

Now, assume first that $c_d(\mathbf{x}[d]) - c_d(\mathbf{x}^*[d]) \neq 0$. Since c_d is increasing, $c'_d(\mathbf{x}[d])(\mathbf{x}[d] - \mathbf{x}^*[d])/(c_d(\mathbf{x}[d]) - c_d(\mathbf{x}^*[d]))$ is nonnegative. Multiplying both sides of inequality (22) by the last quantity, we get (21) as desired.

Otherwise, assume that $c_d(\mathbf{x}[d]) - c_d(\mathbf{x}^*[d]) = 0$. Clearly (21) holds if $\mathbf{x}[d] = \mathbf{x}^*[d]$. So assume $\mathbf{x}[d] \neq \mathbf{x}^*[d]$. Since c_d is increasing it must follow that $c_d(t) = c_d(\mathbf{x}^*[d])$ for all t "in between" $\mathbf{x}^*[d]$ and $\mathbf{x}[d]$, that is for all $t \in [\min\{\mathbf{x}[d], \mathbf{x}^*[d]\}, \max\{\mathbf{x}[d], \mathbf{x}^*[d]\}]$. Thus $c'_d(\mathbf{x}[d]) = 0$ and (21) follows here as well. □

13. Appendix: Additional Experiments

13.1. Result Matching Prediction (Regression)

This section compares train times of models with and without 2D shape constraints. We used a real-world dataset from a large internet services company. The problem is to predict how well a candidate result matches a given query, expressed as a numeric score. The dataset has 1,104,439 train examples. Of the $D = 14$ features, based on domain expertise, 11 features were constrained to be monotonic, 12 feature pairs were chosen for Edgeworth, and 4 feature pairs were chosen for range dominance. We trained an ensemble of 50 calibrated lattices with each base model seeing 6-10 of the $D = 14$ possible features and experimented with batch sizes of 128 and 1024. All runs were trained for 1000 epochs on a workstation with 6 Intel Xeon W-2135 CPUs. Table 7 reports the train time of each run in hours. Train time with 2D shape constraints was 10-20% longer than train time without 2D shape constraints.

Table 7. Train time comparison of models with and without 2D shape constraints in hours.

Batch Size	Monotonic	Monotonic + 2D constraints
128	30.73	37.45
1024	11.52	12.48

13.2. Law School Admission Dataset (Classification)

In this section, we describe experiments conducted on the Law School Admissions dataset (Wightman, 1998). The problem is to predict whether a student would pass the bar exam based on their LSAT score and undergraduate GPA. The dataset has 27,234 students in total and is randomly split 80-20 into train and test sets. Note that the intention of this experiment is simply to provide nice illustrations of how a calibrated lattice model looks with different choices of constraints, as shown in Fig. 2. Table 8 shows both train and test accuracy numbers. The numbers are no worse when the model is further constrained with any of the 2D shape constraints, but learns a more controlled function providing extra interpretability of the feature interactions and guarantees on the model behavior as illustrated by the contour plots in Fig. 2.

Table 8. Performance of a calibrated lattice model on the Law School Admission dataset with different 2D shape constraints.

Model	Train Acc.	Test Acc.
Monotonic	94.98%	94.79%
Mono. + Monotonic Dom.	94.95%	94.81%
Mono. + Range Dom.	94.95%	94.77%
Mono. + Edgeworth	94.95%	94.79%
Mono. + Trapezoid	94.89%	94.83%

13.3. User Intent Prediction (Classification) and Result Matching Prediction (Regression)

This section presents full results of training several calibrated lattice ensemble models on the User Intent Prediction dataset discussed in section 9.4 and the Result Matching Prediction dataset discussed in section 13.1. Models were trained on US data and evaluated on non-US data. We experimented with different levels of regularization: monotonic, dominance, Edgeworth, trapezoid; different numbers of base lattice models: 10, 50; and different numbers of training epochs: 30, 100, 200, 300.

Table 9 shows the Train and Test MSE as well as the Test AUC for the User Intent Prediction dataset for all runs. As expected, the Train MSE is a little better for the 50 lattice ensemble, since it has more model flexibility than the 10 lattice ensemble. In all cases except for the runs with trapezoid, the 100 train epoch runs performed better than the 200 and 300 train epoch runs. We also see that models with dominance and Edgeworth constraints had comparable or lower Test MSE than models without 2D shape constraints across model specifications.

Table 10 shows the Train and Test MSE for random IID splits of the User Intent Prediction dataset averaged over 5 runs. We see that Test MSEs of the models with domi-

nance and Edgeworth constraints are all within the margin of error of the monotonic only model with an additional benefit of these models being further explained by 2D shape constraints.

Table 11 shows the Train and Test MSE for Result Matching Prediction for all runs. Similarly as for the User Intent Prediction dataset, the Train MSE is a little better for the 50 lattice ensemble. There is less of a consistent pattern for the effect of the number of training epochs. Models with Edgeworth constraints perform slightly better for 10 lattice ensembles, while other 2D constraints increase Test MSE.

Table 9. Full results of different calibrated lattice ensemble models on non-IID split of the User Intent Prediction dataset.

Model	Lattices	Epochs	Train MSE	Test MSE	Test AUC
Monotonic	10	100	0.6710±0.0003	0.7544±0.0004	0.8378±0.0004
		200	0.6697±0.0004	0.7547±0.0005	0.8369±0.0005
		300	0.6694±0.0003	0.7563±0.0004	0.8360±0.0003
	50	100	0.6660±0.0001	0.7516±0.0002	0.8390±0.0001
		200	0.6622±0.0010	0.7529±0.0008	0.8376±0.0007
		300	0.6615±0.0001	0.7533±0.0004	0.8375±0.0003
Mono. + Dominance	10	100	0.6729±0.0001	0.7534±0.0001	0.8377±0.0002
		200	0.6719±0.0002	0.7545±0.0003	0.8368±0.0002
		300	0.6717±0.0004	0.7555±0.0006	0.8361±0.0003
	50	100	0.6685±0.0001	0.7509±0.0002	0.8387±0.0002
		200	0.6657±0.0002	0.7516±0.0002	0.8379±0.0001
		300	0.6643±0.0003	0.7529±0.0003	0.8369±0.0002
Mono. + Edgeworth	10	100	0.6750±0.0001	0.7522±0.0003	0.8390±0.0002
		200	0.6742±0.0003	0.7529±0.0001	0.8384±0.0001
		300	0.6739±0.0003	0.7547±0.0004	0.8371±0.0004
	50	100	0.6707±0.0001	0.7506±0.0001	0.8394±0.0001
		200	0.6684±0.0002	0.7511±0.0004	0.8385±0.0002
		300	0.6674±0.0002	0.7518±0.0005	0.8381±0.0002
Mono. + Edge. + Dom.	10	100	0.6766±0.0001	0.7518±0.0001	0.8389±0.0002
		200	0.6758±0.0003	0.7525±0.0002	0.8378±0.0003
		300	0.6753±0.0002	0.7543±0.0003	0.8369±0.0001
	50	100	0.6726±0.0001	0.7506±0.0002	0.8389±0.0002
		200	0.6706±0.0001	0.7512±0.0003	0.8379±0.0002
		300	0.6697±0.0003	0.7521±0.0002	0.8371±0.0002
Mono. + Trapezoid	10	100	0.6883±0.0002	0.7650±0.0004	0.8354±0.0002
		200	0.6878±0.0001	0.7642±0.0005	0.8357±0.0004
		300	0.6887±0.0003	0.7656±0.0006	0.8356±0.0005
	50	100	0.6856±0.0001	0.7637±0.0002	0.8362±0.0001
		200	0.6839±0.0001	0.7632±0.0004	0.8363±0.0002
		300	0.6836±0.0003	0.7634±0.0007	0.8359±0.0004
Mono. + Trap. + Dom.	10	100	0.6968±0.0002	0.7675±0.0010	0.8315±0.0003
		200	0.6973±0.0007	0.7679±0.0007	0.8307±0.0006
		300	0.6969±0.0007	0.7664±0.0007	0.8317±0.0008
	50	100	0.6954±0.0002	0.7663±0.0001	0.8318±0.0002
		200	0.6946±0.0003	0.7658±0.0004	0.8321±0.0001
		300	0.6938±0.0002	0.7653±0.0004	0.8326±0.0004

Table 10. Full results of different calibrated lattice ensemble models on IID splits of the User Intent Prediction dataset.

Model	Lattices	Epochs	Train MSE	Test MSE
Monotonic	10	100	0.7002±0.0010	0.7092±0.0028
		200	0.6998±0.0011	0.7094±0.0026
		300	0.6998±0.0010	0.7097±0.0031
	50	100	0.6958±0.0008	0.7059±0.0029
		200	0.6937±0.0009	0.7064±0.0027
		300	0.6925±0.0011	0.7070±0.0028
Mono. + Dominance	10	100	0.7023±0.0007	0.7098±0.0031
		200	0.7022±0.0009	0.7103±0.0033
		300	0.7022±0.0010	0.7105±0.0029
	50	100	0.6971±0.0016	0.7073±0.0033
		200	0.6957±0.0011	0.7077±0.0031
		300	0.6954±0.0007	0.7081±0.0026
Mono. + Edgeworth	10	100	0.7033±0.0007	0.7101±0.0031
		200	0.7030±0.0012	0.7107±0.0027
		300	0.7031±0.0011	0.7109±0.0029
	50	100	0.6981±0.0014	0.7071±0.0028
		200	0.6976±0.0006	0.7071±0.0031
		300	0.6972±0.0008	0.7084±0.0038
Mono. + Edge. + Dom.	10	100	0.7046±0.0008	0.7108±0.0028
		200	0.7048±0.0009	0.7116±0.0029
		300	0.7051±0.0011	0.7121±0.0030
	50	100	0.7002±0.0013	0.7084±0.0030
		200	0.6999±0.0012	0.7094±0.0022
		300	0.6997±0.0007	0.7097±0.0038
Mono. + Trapezoid	10	100	0.7191±0.0009	0.7229±0.0025
		200	0.7191±0.0006	0.7231±0.0032
		300	0.7195±0.0010	0.7234±0.0026
	50	100	0.7152±0.0008	0.7198±0.0026
		200	0.7147±0.0005	0.7194±0.0028
		300	0.7143±0.0009	0.7191±0.0025
Mono. + Trap. + Dom.	10	100	0.7265±0.0013	0.7306±0.0026
		200	0.7270±0.0011	0.7310±0.0028
		300	0.7264±0.0014	0.7304±0.0025
	50	100	0.7251±0.0009	0.7293±0.0029
		200	0.7255±0.0011	0.7300±0.0027
		300	0.7244±0.0013	0.7291±0.0027

Table 11. Full results of different calibrated lattice ensemble models on the Result Matching Prediction dataset.

Model	Lattices	Epochs	Train MSE	Test MSE
Monotonic	10	30	0.5673	0.7945
		100	0.5646	0.7953
		300	0.5681	0.7964
	50	30	0.5657	0.7925
		100	0.5659	0.7921
		300	0.566	0.7926
Mono. + Dominance	10	30	0.5708	0.7961
		100	0.5753	0.7996
		300	0.5768	0.7979
	50	30	0.5733	0.7965
		100	0.5764	0.7963
		300	0.5738	0.7941
Mono. + Edgeworth	10	30	0.5689	0.7927
		100	0.5736	0.7949
		300	0.5739	0.7923
	50	30	0.5691	0.7949
		100	0.5685	0.7961
		300	0.565	0.7925
Mono. + Edge. + Dom.	10	30	0.5776	0.796
		100	0.5789	0.7987
		300	0.5766	0.7968
	50	30	0.5694	0.7954
		100	0.5723	0.7952
		300	0.5717	0.7958
Mono. + Trapezoid	10	30	0.575	0.8003
		100	0.5764	0.7997
		300	0.5748	0.7974
	50	30	0.5702	0.798
		100	0.571	0.7968
		300	0.5667	0.7958
Mono. + Trap. + Dom.	10	30	0.5785	0.8025
		100	0.5801	0.8002
		300	0.5768	0.7976
	50	30	0.5745	0.7984
		100	0.5708	0.7985
		300	0.5733	0.7979