

An Extended Validity Domain for Constraint Learning

Yilin Zhu* Samuel Burer†

May 27, 2025

Abstract

We consider embedding a predictive machine-learning model within a prescriptive optimization problem. In this setting, called *constraint learning*, we study the concept of a *validity domain*, i.e., a constraint added to the feasible set, which keeps the optimization close to the training data, thus helping to ensure that the computed optimal solution exhibits less prediction error. In particular, we propose a new validity domain which uses a standard convex-hull idea but in an extended space. We investigate its properties and compare it empirically with existing validity domains on a set of test problems for which the ground truth is known. Results show that our extended convex hull routinely outperforms existing validity domains, especially in terms of the function value error, that is, it exhibits closer agreement between the true function value and the predicted function value at the computed optimal solution. We also consider our approach within two stylized optimization models, which show that our method reduces feasibility error, as well as a real-world pricing case study.

1 Introduction

The fields of optimization and machine learning (ML) are closely intertwined, a prominent example being the use of optimization as a subroutine for training ML models. ML assists optimization in a number of interesting and beneficial ways, too. [Kotary et al. \(2021\)](#) classify such approaches into two groups: *ML-augmented optimization*, which uses ML to enhance existing optimization algorithms, e.g., when ML models emulate expensive branching rules within a branch-and-bound algorithm; and *end-to-end optimization*, which combines ML and optimization techniques to obtain the optimal solutions of optimization problems. [De Filippo et al. \(2018\)](#) and [Bengio et al. \(2021\)](#) also survey the use of ML techniques for modeling various components of combinatorial-optimization algorithms, thus improving the algorithms’ accuracy and efficiency.

[Sadana et al. \(2025\)](#) have recently proposed another taxonomy to describe the interactions between ML and optimization. They focus on *contextual optimization*, i.e., when an

*Department of Mathematics, University of Iowa, Iowa City, IA, 52242–1994, USA. Email: yilin-zhu@uiowa.edu.

†Department of Business Analytics, University of Iowa, Iowa City, IA, 52242–1994, USA. Email: samuel-burer@uiowa.edu.

optimization problem depends on uncertain parameters that are themselves correlated with available side information, covariates, and features. In particular, the authors identify three subcategories of contextual optimization: (i) *decision-rule optimization* uses an ML model to predict an optimal solution directly; (ii) *sequential learning and optimization* first uses an ML model to predict uncertain parameters in an optimization problem, which is subsequently solved to obtain an optimal solution; and (iii) *integrated learning and optimization* combines both training and optimization with the goal of improving the quality of the final optimal solution, not specifically the quality of the ML prediction.

In this paper, we consider a specific case of sequential learning and optimization called *constraint learning (CL)*; see Maragno et al. (2025) and the survey of Fajemisin et al. (2023). In this setting, a predictive model is included as a function—which we denote by \hat{h} —within an optimization problem. This function is learned from empirical data, for example, because it lacks an explicit formula and hence cannot be employed within an optimization model in a traditional manner. After \hat{h} is learned and inserted into a constraint or objective, the optimization model is solved to obtain an optimal solution.

Because many popular ML models (e.g., linear regression, neural networks, and decision trees) are mixed-integer-programming (MIP) representable, the final optimization model can often be solved by off-the-shelf software such as Gurobi. To simplify the implementation and solution of CL models as MIPs, several software packages have recently been developed. These include OptiCL (Maragno et al. 2025), JANOS (Bergman et al. 2022), OMLT (Ceccon et al. 2022), PyEPO (Tang and Khalil 2024), and Gurobi versions 10 and later (Gurobi 2023). While the precise details of how these packages embed an ML model into an optimization problem are critical for the overall efficiency of the solver’s performance, we do not consider such details in this paper.

The CL paradigm has attracted significant attention recently. Tjeng et al. (2019) reformulated neural networks as MIPs to evaluate their adversarial accuracy. Maragno et al. (2025) learned a so-called *palatability constraint* for the optimization of food baskets provided by the World Food Programme. The same authors also optimized chemotherapy regimens for gastric-cancer patients. This involved learning and embedding a clinically relevant total-toxicity function within a constraint. In a hypothetical context related to university admissions, Bergman et al. (2022) maximized the expected incoming class size using an ML model that predicts the probability of individual students accepting an admission offer. Each student’s probability was a function of his or her high-school background as well as a university decision variable for the amount of scholarship offered to that student. The university also faced a fixed overall scholarship budget. Mistry et al. (2021) used a gradient-boosted tree to model the relationship between the proportions and properties of ingredients within a concrete mixture, the goal being to optimize the strength of such a mixture.

One challenge for CL is that the error inherent in the embedded model \hat{h} , as described above, may manifest as error in the final optimization model. Indeed, researchers have realized that the CL approach can sometimes lead to an unreasonable computed optimal solution \hat{x} . One particular downside occurs when \hat{x} , although feasible and optimal for the final optimization model based on \hat{h} , is nevertheless too far from the original data on which \hat{h} has been trained. Because the accuracy of \hat{h} can deteriorate far from the data (i.e., poor extrapolation), \hat{x} can be optimal with respect to \hat{h} , while in reality being severely suboptimal—or even infeasible.

Researchers have proposed a remedy for this downside, called a *validity domain*. (Another common term in the literature is *trust region*, but since this term has already been used extensively in the nonlinear-programming literature, we prefer *validity domain* in this paper.) To guard against poor extrapolation and its downstream effect on the optimization, a validity domain further constrains the feasible region of the optimization to be closer to the data, i.e., to a subset where the predictions of \hat{h} are likely to be more reliable.

Many different types of validity domains have been proposed in the literature. Courrieu (1994) defined several validity domains for the specific case when the learned function \hat{h} is a neural network. In particular, we will explore one of these in this paper: the *convex hull* validity domain, i.e., when the variable x of the optimization is constrained to be within the convex hull of the training data. Schweidtmann et al. (2022) used persistent homology to study the topological structure of the data and then constructed a validity domain using the convex hull combined with one-class support vector machines. Maragno et al. (2025) used an enlarged convex hull of the data set to define a validity domain, thus allowing the final optimal solution to be slightly outside the data. Shi et al. (2024) compared six different validity domain techniques, including one of their own design based on an *isolation forest*. An isolation forest is a type of one-class classification model that is, in particular, MIP-representable. They showed that their isolation-forest validity domain was generally the most accurate among the six, while still being computationally efficient.

In this paper, we introduce a new validity domain, which is based on applying the convex-hull idea in an extended space, which concatenates the optimization variable x with the output of the learned function \hat{h} . This approach arises from an intuition that the convex-hull idea is helpful not only for describing the original data set but also for “learning” the constraints, objective, and optimal solution of the optimization problem. We explain this intuition and provide theoretical support in Section 3. In Sections 4–6, we test our approach and observe, generally speaking, significant improvement in the error at the optimal solution \hat{x} of the final model. We also show that our approach does not require significantly more time than the regular convex-hull idea, which acts only in the x space, while being faster than the method of Shi et al. (2024).

Note that we focus on regression models instead of classification models; that is, the output of \hat{h} is continuous, not discrete. Also, it is important to note that our approach is agnostic to the type of predictive model \hat{h} used, i.e., our approach can be applied no matter the functional form of \hat{h} .

This paper is organized as follows. In Section 2, we recount required background on constraint learning, and then in Section 3, we introduce our new validity domain in the extended space. Sections 4–6 contain numerical results and examples illustrating our method. We conclude the paper in Section 7.

2 Background on Constraint Learning

In this section, we provide the relevant background on constraint learning.

2.1 Fundamentals

Formally, we study the following standard-form model introduced in [Fajemisin et al. \(2023\)](#):

$$\hat{v} := \min_{(x,y) \in \hat{F}} f(x) \quad \text{where} \quad \hat{F} := \{(x,y) : x \in X, g(x) \leq 0, \theta(y) \leq 0, y = \hat{h}(x)\}. \quad (1)$$

Here, $x \in \mathbb{R}^{n_1}$ is the vector of decision variables, $f : \mathbb{R}^{n_1} \rightarrow \mathbb{R}$ is a function capturing the objective function in x , X is a simple ground set such as a box or a sphere, and $g : \mathbb{R}^{n_1} \rightarrow \mathbb{R}^{m_1}$ is a function capturing the constraints on x . In addition, $y \in \mathbb{R}^{n_2}$ is a vector of auxiliary variables, and $\theta : \mathbb{R}^{n_2} \rightarrow \mathbb{R}^{m_2}$ is a function expressing constraints on y . Finally, $\hat{h} : \mathbb{R}^{n_1} \rightarrow \mathbb{R}^{n_2}$ is a function, corresponding to a predictive ML model, which maps the decision variables x to the auxiliary variables y . In words, y are the predictions of x via the learned function \hat{h} . In total, there are $n := n_1 + n_2$ variables, $m := m_1 + m_2$ inequality constraints, and n_2 equality constraints. We assume that $\hat{F} \neq \emptyset$ and that (1) attains its optimal value \hat{v} , and we use the notation $\widehat{\text{Opt}}$ to denote the optimal solution set.

The function \hat{h} plays an important modeling role by acting as a surrogate, or approximation, for a true function $h : \mathbb{R}^{n_1} \rightarrow \mathbb{R}^{n_2}$, which is not known explicitly but can be learned via a suitable ML technique applied to empirical data. Accordingly, we consider the *true optimization model*

$$v^* := \min_{(x,y) \in F} f(x) \quad \text{where} \quad F := \{(x,y) : x \in X, g(x) \leq 0, \theta(y) \leq 0, y = h(x)\}, \quad (2)$$

which differs from (1) only by the equation $y = h(x)$ in the constraints. The value v^* denotes the *true optimal value*, and we use Opt^* to denote the *true optimal solution set*, which we assume to be nonempty. Roughly speaking, if \hat{h} is a good approximation of h , then one expects \hat{v} and $\widehat{\text{Opt}}$ to be good approximations of v^* and Opt^* .

We assume that the approximation \hat{h} of h is learned from a data set in X of size N ,

$$D_N := \{x^{(i)} \in X : i = 1, \dots, N\}, \quad (3)$$

and from the corresponding N observed function values $h(D_N) := \{h(x^{(i)}) : i = 1, \dots, N\}$. In practice, the observed function values contain noise, and ML techniques that learn \hat{h} from D_N and $h(D_N)$ should account for this noise, e.g., to avoid overfitting. In this paper, we will not explicitly model the noise. Rather, just as ML techniques learn \hat{h} using the noisy evaluations of h , our method described in Section 3 will make direct use of the noisy evaluations as well when applied in practice; see Sections 4-6.

Other variations of (1) are possible. For example, the objective and constraint functions f and g could involve both x and y , not just x . In addition, there could be auxiliary features, say w , incorporating contextual information about the optimization problem. Another possibility is additional variables, say z , that are (nonlinear) combinations of the main decision variables x . Together, w and z could then be used to build a better approximation $\hat{h}(x, w, z)$ of the predicted outputs. For the sake of simplicity, we will not incorporate w and z in this paper, but we refer the reader to [Fajemisin et al. \(2023\)](#) for additional references, which do take into account w and z .

From time to time in this paper, we will also refer to an even simpler form of (1), which optimizes a predictive model over a simplified feasible set:

$$\min_x \{ \hat{f}(x) : x \in X \}. \quad (4)$$

This problem is an instance of (1) but with y appearing in the objective: $\min_{x,y} \{y : x \in X, y = \hat{f}(x)\}$ and g and θ nonexistent.

2.2 Errors

When comparing (1) with (2), hopefully \hat{v} and $\widehat{\text{Opt}}$ are respectively close to v^* and Opt^* . To measure this precisely, we formally define the errors

$$\begin{aligned} \text{optimal value error} &:= |\hat{v} - v^*|, \\ \text{optimal solution error} &:= \|\hat{x} - x^*\|, \end{aligned}$$

where \hat{x} and x^* are given optimal solutions for (1) and (2), respectively. Note that the optimal solution error depends on a *particular* pair of optimal solutions and hence could vary for different pairs if there are indeed multiple optimal solutions, but we choose this definition since it is both natural and practical.

Another type of error refers directly to the quality of the approximation \hat{h} of h for a particular feasible solution (\hat{x}, \hat{y}) of (1). We define the

$$\text{function value error at } (\hat{x}, \hat{y}) := \|\hat{y} - h(\hat{x})\| = \|(\hat{h} - h)(\hat{x})\|.$$

That is, the function value error at \hat{x} measures the Euclidean difference in \mathbb{R}^{n_2} between the true value of h at \hat{x} and its predicted value under \hat{h} . Although the function value error is defined for any (\hat{x}, \hat{y}) , we will typically measure it at an optimal solution of (1).

A final type of error measures how close to true feasibility a given feasible $(\hat{x}, \hat{y}) \in \widehat{F}$ is:

$$\text{feasibility error at } (\hat{x}, \hat{y}) := \|\max\{0, \theta(h(\hat{x}))\}\|,$$

where the maximum between 0 and the vector $\theta(h(\hat{x}))$ is taken component-wise. Recall that \widehat{F} and the true feasible set F differ only in that \widehat{F} uses the constraint $y = \hat{h}(x)$, while F uses $y = h(x)$. In particular, given $(\hat{x}, \hat{y}) \in \widehat{F}$, the naturally corresponding element of F is $(\hat{x}, h(\hat{x}))$. Because \hat{x} already satisfies $\hat{x} \in X$ and $g(\hat{x}) \leq 0$, true feasibility of $(\hat{x}, h(\hat{x}))$ then holds if and only if $\theta(h(\hat{x})) \leq 0$. The feasibility error is simply a measure of how much this constraint is violated.

We remark that both the function value and feasibility errors are exactly measurable only when h can be evaluated without noise.

2.3 Validity domains

Solving (1) may lead to large errors compared to the true optimization problem (2). The *validity domain* concept has thus been introduced to mitigate these errors, particularly the

function value error. Given a validity domain $V \subseteq \mathbb{R}^{n_1}$, we introduce the following optimization problem, which is closely related to (1):

$$\hat{v}(V) := \min_{x,y} \left\{ f(x) : (x,y) \in \hat{F}, x \in V \right\} \quad (5)$$

In words, (5) is simply (1) with the added constraint $x \in V$. We denote the optimal value as $\hat{v}(V)$ to reflect the presence of V , and $\hat{x}(V)$ denotes an optimal solution of (5).

We next describe several validity domains, which have been proposed in the literature.

2.3.1 Simple bounds

A simple validity domain is the smallest coordinate aligned hypercube, which surrounds the empirical data D_N defined in (3):

$$\text{Box} := \left\{ x \in \mathbb{R}^{n_1} : \min_{i=1}^N x_j^{(i)} \leq x_j \leq \max_{i=1}^N x_j^{(i)} \quad \forall j = 1, \dots, n_1 \right\}. \quad (6)$$

2.3.2 The convex hull

Recall that D_N in (3) denotes the N points in X over which \hat{h} has been learned using the (possibly noisy) function values $h(D_N)$. Let $\text{conv}(D_N)$ be the smallest convex subset containing D_N , i.e., its convex hull. Courrieu (1994) first used the convex hull as a validity domain in the case that \hat{h} modeled feed-forward neural networks, and Maragno et al. (2025) also advocated the use of $\text{conv}(D_N)$ as a validity domain. Throughout the rest of this paper, we will use the notation

$$\text{CH} := \text{conv}(D_N) \quad (7)$$

to denote this specific validity domain. Including the constraint $x \in \text{CH}$ in (1) amounts to adding extra variables and linear constraints, which can be expensive when N is large. Maragno et al. (2025) discuss strategies to ease the computational burden, e.g., column generation. They also proposed several variants of CH, which can be implemented in modern optimization solvers.

2.3.3 The enlarged convex hull

Balestriero et al. (2021) argue that, for higher-dimensional datasets, say, with $n \geq 100$, the phenomenon of interpolation (i.e., when a new sample belongs to the convex hull of the existing samples) almost surely never occurs. Hence, generalization of a fitted model must occur via extrapolation, not interpolation. As a result, in the context of CL, the convex hull method just mentioned may be too conservative. To alleviate this conservatism, Maragno et al. (2025) proposed an *enlarged convex hull* formally defined as the Minkowski sum

$$\text{CH}_\varepsilon := \text{CH} + B_p(0, \varepsilon), \quad (8)$$

where $B(0, \varepsilon)$ is the p -norm ball of a user-supplied radius $\varepsilon > 0$ centered at 0. When $p = 1$ or $p = \infty$, this can be modeled using linear programming, or when $p = 2$, using second-order cone programming. The authors also demonstrated the effectiveness of CH_ε in their setting.

2.3.4 Support vector machines

Schweidtmann et al. (2022) proposed the use of a one-class support vector machine to define a validity domain. In particular, the authors expressed x as being in-distribution using a validity domain defined by the constraint $\sum_i \alpha_i K(x_i, x) + b \geq 0$, where i indexes the support vectors, x_i are the support vectors, $K(x_i, x)$ is the RBF kernel function $\exp(-\gamma \cdot \|x - x_i\|^2)$ with a hyperparameter γ , and α_i, b are the learned parameters. Although this constraint is neither convex nor MIP-representable, the authors demonstrated numerical success using a custom global optimization algorithm.

2.3.5 Isolation forests

Similar to the preceding subsection, another model for outlier detection is the *isolation forest* (Liu et al. 2008), which is an ensemble-tree method to isolate outliers based on their quantitative characteristics, e.g., their atypical feature values. Each tree in the forest is trained to isolate instances by their features, and by design, outliers are closer to the root of the tree because they are more susceptible to isolation. The isolation forest then defines an instance to be an outlier if it has a short average path length to the root node, where the average is taken over all isolation trees in the forest.

Using a given isolation forest trained on the observed input data set, Shi et al. (2024) defined a validity domain to be the set of points x , where the path length of x in every tree of the forest is larger than a pre-determined threshold d . Note that this is equivalent to disallowing any x that is an outlier in some tree, thus slightly deviating from the use of the average measure described in the preceding paragraph. They also developed a MIP representation of the isolation forest enabling it to be embedded into an optimization model and showed the effectiveness of this validity domain compared to others in the literature.

3 An Extended Validity Domain

We introduce a new validity domain, which is based on a simple observation—though perhaps not well known—that (2) is equivalent to a convex optimization problem in an extended space. Note that, as discussed in Section 2, we assume throughout this section that noiseless observations of the true function $h(x)$ are available, even as the precise functional form of $h(x)$ is unknown.

3.1 Intuition

Recall that F denotes the true feasible region of (2) in the space $(x, y) \in \mathbb{R}^n$. Now define

$$F^+ := \{(x, y, f(x)) : (x, y) \in F\}$$

to be the *extended feasible set*, which is also sometimes called the *graph of f over F* . In words, F^+ is the set of all triples (x, y, ϕ) in \mathbb{R}^{n+1} , where (x, y) is feasible and the scalar $\phi = f(x)$ encodes the objective value. The following proposition states that (2) is equivalent to the minimization of ϕ over $(x, y, \phi) \in \text{conv}(F^+)$, i.e., a convex minimization:

Proposition 1. $v^* = \min\{\phi : (x, y, \phi) \in \text{conv}(F^+)\}$.

Proof. Note that ϕ is a scalar variable, which represents the function value $f(x)$ in the extended space. Let ν denote the optimal value of $\min\{\phi : (x, y, \phi) \in \text{conv}(F^+)\}$, and let (x^*, y^*) denote an optimal solution of (2) such that $f(x^*) = v^*$. Then (x^*, y^*, v^*) is a member of $\text{conv}(F^+)$, and hence $v^* \geq \nu$. To show $v^* \leq \nu$, consider an optimal $(x, y, \nu) \in \text{conv}(F^+)$. There exists a finite index set $K \ni k$, feasible points (x^k, y^k) for (2), and multipliers $\lambda_k \geq 0$ such that $\sum_{k \in K} \lambda_k = 1$ and

$$(x, y, \nu) = \sum_{k \in K} \lambda_k (x^k, y^k, f(x^k)) = \left(\sum_{k \in K} \lambda_k x^k, \sum_{k \in K} \lambda_k y^k, \sum_{k \in K} \lambda_k f(x^k) \right).$$

In particular, ν is greater than or equal to the minimum $f(x^k)$, which is itself greater than or equal to v^* . \square

Proposition 1 establishes that solving (2) is equivalent to the minimization of a linear function over the convex hull of F^+ . Hence, in a certain sense, understanding $\text{conv}(F^+)$ provides a key to solving (2). Indeed, in what follows, we take the point of view that $\text{conv}(F^+)$ can help us “learn” the optimization problem (2) from a new perspective—one which is complementary to the use of \hat{h} described in Section 2.

Proposition 2 below provides further intuition that $\text{conv}(F^+)$ learns and describes the true optimization model (2)—but this time from a data-driven perspective. Let D_N be the data set given in (3) over which \hat{h} has been trained. We define two related data sets:

$$D_N^+ := \{(x^{(i)}, y^{(i)}, \phi^{(i)}) : y^{(i)} = h(x^{(i)}), \phi^{(i)} = f(x^{(i)}), i = 1, \dots, N\} \quad (9)$$

$$F_N^+ := \{(x, y, \phi) \in D_N^+ : (x, y) \in F\} \quad (10)$$

Here, $D_N^+ \subseteq \mathbb{R}^{n+1}$ is an extension of D_N , which appends the true values of h and f to each data point $x^{(i)}$. Further, F_N^+ is the restriction of D_N^+ to just those data points that are feasible for (2). In particular, D_N^+ may contain both feasible and infeasible $(x^{(i)}, y^{(i)})$. For these definitions, recall our assumption that noiseless evaluations of h are available, which in particular allows one to check which data points in D_N^+ are feasible.

Using the extended data sets D_N^+ and F_N^+ and viewing the sample size N as a parameter, we now show that, if one takes larger and larger samples such that F_N^+ becomes dense in the extended feasible set F^+ , then solving the surrogate optimization over $\text{conv}(F_N^+)$ eventually solves (2).

Proposition 2. Let $\{D_N\}_{N=1}^\infty \subseteq X$ be a sequence of data sets with

$$\text{conv}(F^+) \subseteq \lim_{N \rightarrow \infty} \text{conv}(F_N^+),$$

and define $v_N^* := \min\{\phi : (x, y, \phi) \in \text{conv}(F_N^+)\}$. Then $v^* = \lim_{N \rightarrow \infty} v_N^*$.

Proof. Because $\text{conv}(F_N^+) \subseteq \text{conv}(F^+)$ for all N , it holds that $v_N^* \geq v^*$ for all N . Now consider $(x^*, y^*, \phi^*) \in F^+$, where (x^*, y^*) is an optimal solution of (2) and hence $\phi^* = v^*$. By assumption, there exists a sequence $\{(x_N, y_N, \phi_N) \in \text{conv}(F_N^+)\}$ converging to (x^*, y^*, v^*) . In particular, $\{\phi_N\} \rightarrow v^*$. Note also that $\phi_N \geq v_N^*$ since (x_N, y_N, ϕ_N) is a member of $\text{conv}(F_N^+)$. In total, $\phi_N \geq v_N^* \geq v^*$ for all N with $\{\phi_N\} \rightarrow v^*$. Hence, $\{v_N^*\} \rightarrow v^*$. \square

3.2 Our new validity domain and its variants

Based on the prior subsection, we propose the following validity domain, where F_N^+ is the extended data set of feasible points defined in (10) based on D_N and D_N^+ from (3) and (9), respectively:

$$\text{CH}^+ := \text{conv}(F_N^+). \quad (11)$$

While $\text{CH} \subseteq \mathbb{R}^{n_1}$ from (7) is defined in the space of x , this new validity domain $\text{CH}^+ \subseteq \mathbb{R}^{n+1}$ is defined in the space of (x, y, ϕ) . However, it is easy to see that the projection of CH^+ onto the variable x is contained in CH , i.e., $\text{proj}_x(\text{CH}^+) \subseteq \text{CH}$.

In order to use CH^+ as a validity domain, we simply add the constraint $(x, y, f(x)) \in \text{CH}^+$ to (1):

$$\hat{v}(\text{CH}^+) := \min_{x,y} \left\{ f(x) : (x, y) \in \hat{F}, (x, y, f(x)) \in \text{CH}^+ \right\}.$$

Note that \hat{h} appears in both constraints through the equation $y = \hat{h}(x)$. Based on the fact that $\text{proj}_x(\text{CH}^+) \subseteq \text{CH}$, we have the following immediate relationship between $\hat{v}(\text{CH}^+)$ and $\hat{v}(\text{CH})$.

Proposition 3. $\hat{v}(\text{CH}^+) \geq \hat{v}(\text{CH})$.

In fact, we believe the property $\text{proj}_x(\text{CH}^+) \subseteq \text{CH}$ is a defining feature of our approach. The validity domain CH acts as a natural geometric restriction on x to keep the optimization close to the training data—with a goal to ameliorate the function value error. By further constraining CH in the space x , our intuition is that the new validity domain CH^+ further aids the optimization as argued in Propositions 1–2. Practically speaking, in Sections 4–6, we will show via example that CH^+ does indeed further reduce the function value error relative to CH .

Other variations of CH^+ , which maintain the property that the projection onto x is contained in CH , are possible, and such variations can be useful in practice (see Section 6). For example, consider a case in which the historical data F_N^+ employed in the definition of CH^+ in (11) is incomplete in the sense that the present-day optimization problem includes additional decision variables for which past data has not been collected. Said differently, one may have historical data on only a subset of the decision variables. Another example occurs when, even if the past data is complete, a new constraint is imposed in the present, rendering much of the historical data infeasible and hence theoretically unfit for the construction of CH^+ . In such a case, it may prove beneficial to make use of all past data—both feasible and infeasible points—so as to not discard too much data.

More formally, let us consider the case that we do indeed have historical data D_N on all variables x , and define $F_N := D_N \cap F$ to be those data points that are feasible for (2). Given arbitrary subsets $J \subseteq \{1, \dots, n_2\}$ and $D \subseteq D_N$ such that $F_N \subseteq D$, consider the following variant of CH^+ :

$$\text{CH}' := \text{conv}(\{(x, y_J, f(x)) : x \in D, y_J = h_J(x)\}) \subseteq \mathbb{R}^{n_1+|J|+1}.$$

Note that CH' equals CH^+ when $J = \{1, \dots, n_2\}$ and $D = F_N$; see the definitions (9) and (10). It is then clear that $\text{proj}_x(\text{CH}^+) \subseteq \text{proj}_x(\text{CH}') \subseteq \text{CH}$. In this sense, CH' is sandwiched between CH^+ and CH . Another variant of CH^+ could drop the function value $f(x)$ in

the definition of CH' . In Sections 5–6, we will show practical cases in which we choose a validity domain CH' between CH and CH^+ . Even in those cases, we will call these the “ CH^+ approach” for simplicity.

Following Section 2.3.3, we also consider the enlarged version of CH^+ , which is based on a user-supplied $\varepsilon > 0$:

$$\text{CH}_\varepsilon^+ := \text{CH}^+ + B_p(0, \varepsilon), \quad (12)$$

where $B_p(0, \varepsilon)$ is the p -norm ball of radius ε in \mathbb{R}^{n+1} . In our experiments in Section 4, we will test $p = 2$, which can be modeled using second-order programming.

3.3 Illustration

For illustration, we refer to the simplified form (4) described in Section 2.1, which optimizes a true but unknown objective function $f(x)$ over a simple domain $x \in X$ by substituting a learned approximation \hat{f} of f . Our problem in this subsection is thus $\min\{\hat{f}(x) : x \in X\}$.

Figure 1 depicts a one-dimensional example ($n_1 = 1$) in which the true function to optimize is $f(x) = (x - 1.75)^2$ over all $x \in X := [0, 4]$. The true optimal value is $v^* = 0$, and the true unique optimal solution is $x^* = 1.75$. We depict $f(x)$ as a dotted curve in both panels of Figure 1, keeping in mind that this curve would be unknown in practice. We take $N = 4$ and $D_N = \{1.00, 1.75, 2.25, 3.00\}$, and then $\hat{f}(x)$ is taken to be the least-squares regression line based on the true function f evaluated at D_N without noise; this is depicted as the solid, sloped line in both panels.

In the left panel, we also depict the validity domain CH as the horizontal line segment overlaid on the x -axis from 1 to 3, which is the convex hull of D_N . Then $\min\{\hat{f}(x) : x \in X \cap \text{CH}\}$ is optimized at $\hat{x}(\text{CH}) = 1.00$. On the graph of $\hat{f}(x)$, this is depicted as the star in the left panel. In contrast, the right panel depicts the set CH^+ , which is the filled quadrilateral in (x, y) , spanning the four sampled points in D_N^+ , which lie on the graph of $f(x)$. Then $\min\{\hat{f}(x) : x \in X, (x, \hat{f}(x)) \in \text{CH}^+\}$ optimizes the function value over the intersection of the regression line and the filled quadrilateral. In this case, the optimal solution occurs at $\hat{x}(\text{CH}^+) \approx 1.37$ and is depicted by the star in the right panel. The plots show that the approach based on CH^+ exhibits better function-value and optimal-solution errors but worse optimal-value error.

Figure 1 also demonstrates that the CH^+ method does not simply return the sampled point with minimum f value. Indeed, the interaction of the sampled values $(x^{(i)}, f(x^{(i)}))$ with the function approximation \hat{f} is critical to the behavior of CH^+ .

4 Numerical Results

To evaluate the extended validity domain CH^+ introduced in Section 3, we adopt a procedure introduced by Shi et al. (2024). To this end, in the following subsections, we define the concept of a basic experiment, then describe our method for generating multiple experiments, and finally detail the optimization results. All numerical experiments were coded using Python 3.10.8 and Gurobi 11.0 and conducted on a single Xeon E5-2680v4 core running at 2.4 GHz with 16 GB of memory under the CentOS Linux operating system. The code and results are publicly shared at <https://github.com/yillzhu/extvdom>.




Figure 1: Illustration of CH on the left and CH^+ on the right. CH^+ exhibits better function-value and optimal-solution errors.

4.1 Definition of an experiment

In our testing, we define an *experiment* to be the full specification of six design options:

- *Ground truth*, i.e., a true optimization problem of the form $v^* = \min_x \{f(x) : x \in X\}$ for which the function $f(x)$ and the true optimal value v^* are known. A true optimal solution $x^* \in \text{Opt}^*$ is known as well. This corresponds to the simplified optimization (4) discussed in Section 2.1 in which the learned function lies in the objective, not the constraints. We will interchangeably identify a ground truth with its function f .
- *Sampling rule*, i.e., a rule \mathcal{R} describing how to sample points in X .
- *Sample size*, i.e., the number of points N to sample in X .
- *Noise level*, i.e., a scale factor σ corresponding to the amount of random noise added to the function evaluations of f during sampling.
- *Seed*, i.e., the seed s used to initiate the random number generator before sampling.
- *ML technique*, i.e., the machine-learning technique \mathcal{M} used to learn \hat{f} from noisy evaluations of f .

Once the ground truth f , sampling rule \mathcal{R} , sample size N , noise factor σ , seed s , and ML model \mathcal{M} are specified, an experiment proceeds as follows. Setting the seed s , we randomly sample $D_N \subseteq X$ following the sampling rule \mathcal{R} , and then we evaluate f at all points in D_N adding random noise scaled by the noise factor σ . Then we use the ML technique \mathcal{M} to learn the approximation \hat{f} using the empirical data D_N and the noisy function values $f(D_N)$. The experiment continues by solving (4) based on \hat{f} with an additional validity-domain constraint for several different domains.

In short, an experiment specifies the six design options, builds the optimization model based on \hat{f} , and solves the model multiple times, each time testing a different validity domain.

Because the ground truth is known, the errors defined in Section 2.2 are easily computed for each validity domain. This allows us to determine, for a given experiment, which validity domain yields smaller errors. Finally, by aggregating these errors over multiple experiments, we can identify trends in the performance of different validity domains.

4.2 Generating multiple experiments

Following Shi et al. (2024), we test seven different ground truths corresponding to seven challenging nonlinear benchmark functions from the literature (see, for example, Surjanovic and Bingham (2023)), namely $f \in \{Beale, Griewank, Peaks, Powell, Qing, Quintic, Rastrigin\}$. The input dimension n_1 of these functions varies from 2 to 10, and in each case, X is an n_1 -dimensional box. Further, we consider two sampling rules $\mathcal{R} \in \{Uniform, Normal\}$, where *Uniform* indicates a uniform sample over the box domain X and *Normal* indicates a jointly independent normal sample around the global minimum $x^* \in X$ of f with covariance matrix ρI , where I is identity matrix of size n_1 and ρ is problem-specific. In particular, we choose ρ to be 1/6 of the distance of x^* to the boundary of X . Due to the nature of the normal distribution in each dimension, this ensures samples D_N that tend to be close to x^* and highly unlikely to be outside X .

The sample size is taken as $N \in \{500, 1000, 1500\}$, and the noise is a zero-mean univariate normal distribution with standard deviation equal to σ times the standard deviation of the noisy $f(D_N)$ function values, where the scale factor $\sigma \in \{0.0, 0.1, 0.2\}$. In particular, $\sigma = 0.0$ corresponds to the no-noise case. Furthermore, we take 100 seeds, specifically $s \in \{2023, 2024, \dots, 2122\}$.

In addition, we test three different machine learning models \mathcal{M} using the scikit-learn package of Python: $\mathcal{M} \in \{RandomForestRegressor, GradientBoostingRegressor, MLPRegressor\}$. In particular, we implement *RandomForestRegressor* with 100 trees and maximum depth of 5; *GradientBoostingRegressor* with 100 boosting stages and a maximum depth of 5 for the individual regression estimators; and *MLPRegressor* with 2 hidden layers and 30 neurons in each layer. (For each experiment independently, we also tried using grid search on the parameters to find the best model for that experiment. Ultimately, we found that the overall conclusions about the various validity domains in Section 4.3 were very similar. So we have fixed the parameters of \mathcal{M} to simplify our experiments and reduce testing time.) Before constructing the ML models, the input features are standardized in $[0, 1]$ using min-max scaling, and the noisy function values are normalized to mean 0 and standard deviation 1.

We generate experiments by looping over all combinations of $(f, \mathcal{R}, N, \sigma, s, \mathcal{M})$ for a total of $7 \times 2 \times 3 \times 3 \times 100 \times 3 = 37,800$ experiments. To provide a snapshot of the quality of the ML models, as well as the time required to compute them, Table 1 shows the average R^2 scores broken down by ground truth and ML technique, and similarly Table 2 shows the median training times (in seconds). Specifically, for the calculation of each R^2 value, we train and test the model using an 80-20 split of the data, and Table 1 shows the performance over the test set. However, the models used for optimization as described in the next subsection are trained on the complete data set.

	Beale	Griewank	Peaks	Powell	Qing	Quintic	Rastrigin
<i>RandomForestRegressor</i>	0.88	0.34	0.87	0.55	0.66	0.82	0.17
<i>GradientBoostingRegressor</i>	0.90	0.39	0.97	0.83	0.89	0.92	0.42
<i>MLPRegressor</i>	0.93	0.62	0.90	0.91	0.89	0.95	0.04

Table 1: Mean test R^2 scores over all 37,800 experiments, grouped by function and ML technique.

	Beale	Griewank	Peaks	Powell	Qing	Quintic	Rastrigin
<i>RandomForestRegressor</i>	0.35	0.46	0.34	0.47	0.65	0.50	0.75
<i>GradientBoostingRegressor</i>	0.24	0.41	0.24	0.41	0.72	0.49	0.89
<i>MLPRegressor</i>	2.42	1.75	1.64	1.76	1.55	1.23	2.19

Table 2: Median training times (in seconds) over all 37,800 experiments, grouped by function and ML technique.

4.3 Optimization results for four validity domains

For each experiment, we solve the corresponding optimization model four times by varying the validity domain $V \in \{\text{BOX}, \text{CH}, \text{ISOFOR}, \text{CH}^+\}$, where BOX, CH, and CH^+ are defined in Sections 2–3, and ISOFOR refers to the isolation-forest validity domain of Shi et al. (2024), which has also been described in Section 2. For ISOFOR, we follow Shi et al. (2024) by setting hyperparameters to their default values and by taking the maximum depth of a tree to be 5 for the *Beale* and *Peaks* functions and 6 otherwise. In the online supplemental material, we show the median optimization setup and solve times for all combinations of ML technique and validity domains.

We now examine the median errors associated with each validity domain. Table 3 presents results for all experiments grouped by function, type of error, and sampling rule. Each group then has four errors corresponding to the four validity domains. Furthermore, within each group of four errors, we scale so that the median error corresponding to BOX equals 1.00, and we show the smallest error in a bold font, thus facilitating comparison of the different validity domains. Note that we choose to report medians rather than means to avoid any out-sized influence by the tails of the associated distributions. However, in the online supplemental material, we report the means of the same experiments for comparison; we find that the qualitative conclusions do not change.

For both sampling rules, we see small function value errors for CH^+ . In particular, for six of the seven functions, CH^+ achieves the best median function value error for *Uniform*. For the seventh function (*Peaks*), CH^+ achieves nearly the best—1.02 versus 1.00 for both BOX and CH. For the *Normal* sampling rule, CH^+ also performs well, where it achieves the best median function value error for all seven functions.

For the optimal value and optimal solution errors, CH^+ performs competitively. With respect to *Uniform*, for five of the seven functions, CH^+ has either the best median optimal value error or the best median optimal solution error. For *Normal*, we see this performance for all seven functions.

In Figure 2, we examine more closely the behavior of the function value error of CH^+ compared to that of CH over all experiments. We plot the empirical distribution of the

Function	Validity Domain	Median Function Value Error		Median Optimal Value Error		Median Optimal Solution Error	
		<i>Uniform</i>	<i>Normal</i>	<i>Uniform</i>	<i>Normal</i>	<i>Uniform</i>	<i>Normal</i>
Beale	BOX	1.00	1.00	1.00	1.00	1.00	1.00
	CH	0.97	0.87	1.00	0.97	1.01	1.18
	IsoFOR	0.63	0.76	0.73	0.91	0.81	0.52
	CH ⁺	0.09	0.35	0.16	0.72	0.86	0.79
Griewank	BOX	1.00	1.00	1.00	1.00	1.00	1.00
	CH	0.73	0.95	1.09	1.00	0.92	0.97
	IsoFOR	0.86	1.00	1.56	1.00	0.21	0.90
	CH ⁺	0.49	0.53	1.18	0.90	0.89	1.08
Peaks	BOX	1.00	1.00	1.00	1.00	1.00	1.00
	CH	1.00	1.00	1.00	1.00	1.01	1.01
	IsoFOR	1.47	1.25	1.16	1.01	1.04	0.88
	CH ⁺	1.02	0.68	1.00	0.94	0.93	0.92
Powell	BOX	1.00	1.00	1.00	1.00	1.00	1.00
	CH	0.99	1.03	0.98	0.91	0.95	0.95
	IsoFOR	0.90	1.06	0.89	0.79	0.66	0.59
	CH ⁺	0.09	0.17	0.15	0.23	0.78	0.63
Qing	BOX	1.00	1.00	1.00	1.00	1.00	1.00
	CH	0.61	0.54	1.00	1.00	0.83	0.83
	IsoFOR	1.20	0.73	1.00	1.00	0.75	0.72
	CH ⁺	0.41	0.45	0.79	0.95	0.70	0.54
Quintic	BOX	1.00	1.00	1.00	1.00	1.00	1.00
	CH	0.63	0.42	0.90	0.36	0.97	0.97
	IsoFOR	0.25	0.16	0.94	0.26	0.77	0.75
	CH ⁺	0.13	0.06	1.00	0.25	0.92	0.74
Rastrigin	BOX	1.00	1.00	1.00	1.00	1.00	1.00
	CH	0.87	0.66	1.02	1.49	0.88	0.65
	IsoFOR	0.92	0.84	1.04	1.39	0.82	0.63
	CH ⁺	0.68	0.49	1.13	1.54	0.70	0.63

Table 3: Errors for all 37,800 experiments, grouped by function, type of error, and sampling rule. In each group of four errors corresponding to four validity domains, the errors are scaled so that BOX has value 1.00 and the smallest error is shown in a bold font in order to facilitate comparison among the different validity domains.

ratio of the function value error for CH^+ divided by the function value error for CH . When the ratio is less than 1.0 for a given experiment, CH^+ has a better function value error; when the ratio is greater than 1.0, CH^+ has a worse error on that experiment. In Figure 2, the distribution of ratios is plotted on a logarithmic scale, and two additional pieces of information are shown. First, a vertical dotted line is plotted to mark the ratio 1.0. Second, the percentage of experiments to the left of 1.0, i.e., when CH^+ is performing better, is annotated. We see in particular that CH^+ achieves a better function value error than CH on over 60% of experiments. In addition, the left tail of the distributions show that CH^+ can reduce the error by up to a factor of 1,000, whereas the right tail indicates that the error from CH^+ is usually no more than 100 times the error from CH . Finally, we note that in both panels of Figure 2, the mode is 1.0, indicating that the most common situation is for CH^+ and CH to yield the same function value error. Although we do not show them here, analogous plots comparing the function value error of CH^+ compared to that of Box and respectively ISOFor showed similar results.




Figure 2: Empirical distributions of the ratio of the function value error of CH^+ divided by the function value error of CH , grouped by sampling rule.

In the online supplemental material, we further provide sensitivity analysis of our method with respect to certain design choices, such as the noise level σ , the sample size N , and the sampling rules, e.g., *Normal* versus *Uniform*.

4.4 Results with the enlarged convex hull

Next we recreate the results of the previous subsection except that here we test the validity domains $V \in \{\text{Box}, \text{CH}^+, \text{CH}_{0.05}^+, \text{CH}_{0.10}^+\}$, where $\text{CH}_{0.05}^+$ and $\text{CH}_{0.10}^+$ refer to enlarged versions of CH^+ corresponding to ε values 0.05 and 0.10, respectively. The purpose of these experiments are to highlight the performance of the enlarged convex hulls relative to CH^+ and Box , and the results are shown in Table 4. Note that the results for Box are scaled to 1.00 as before, and the rows for CH^+ in Table 4 match the rows for CH^+ in Table 3, i.e., Table 4 contains some repeated, relevant information from Table 3.

We also mention that, in actuality, the ε values are 0.05 and 0.10 times a scaling factor that is proportional to the diameter of the ground set X , which is an n_1 -dimensional box in each case. Specifically, we use the scale factor $\text{diam}(X)$, the diameter of X , to account for the natural size of the domain.

Examining the function value errors in Table 4, we see that in nearly all groups of four errors (each group corresponding to the four validity domains), as ε increases from 0 to 0.05 to 0.10, we see a corresponding increase in the function value error. In words, increased flexibility of the enlarged convex hull typically comes with a sacrifice in the function value error. Examining the optimal value and optimal solution errors, we do not see a consistent trend, perhaps because CH^+ and its enlarged version CH_ε^+ are designed to manage the function value error and hence there is naturally a weaker link with the optimal value and optimal solution errors.

4.5 Results for larger dimensions

The computational scalability of any CL approach faces several obstacles, which may grow more severe with the base dimension n_1 , the sample size N , the complexity of the learned model \mathcal{M} , and the resultant size and structure of the mixed-integer programming model. Moreover, each validity domain will scale differently than other domains.

Even though CL will require more time and memory for larger, more complex problems, one would hope the quality and effectiveness of the approach does not simultaneously degrade. Accordingly, in this subsection, we test CH^+ on larger problems than in Tables 3 and 4. In particular, we test an additional function for dimension $n_1 = 20$:

$$f(x) := \sum_{i=1}^n (x_1 + \cdots + x_i)^2,$$

where each $x_i \in [-200, 200]$. This function is called the *rotated hyper ellipsoid* function in Long et al. (2019).

In terms of computational scalability, for these experiments, we found that whenever $\mathcal{R} = \text{Uniform}$,

$$\mathcal{M} \in \{\text{RandomForestRegressor}, \text{GradientBoostingRegressor}\},$$

or $V = \text{ISOFOR}$, the run time of each experiment was prohibitive for our testing, e.g., in such cases, a single experiment could take over 30 minutes. So in this subsection, we tested all other combinations, i.e.,

$$\begin{aligned} \mathcal{R} &\in \{\text{Normal}\}, \\ N &\in \{500, 1000, 1500\}, \\ s &\in \{2023, 2024, \dots, 2122\} \\ \sigma &\in \{0.0, 0.1, 0.2\} \\ \mathcal{M} &\in \{\text{MLPRegressor}\} \\ V &\in \{\text{Box}, \text{CH}, \text{CH}^+, \text{CH}_{0.05}^+, \text{CH}_{0.10}^+\}. \end{aligned}$$

Function	Validity Domain	Median Function Value Error		Median Optimal Value Error		Median Optimal Solution Error	
		<i>Uniform</i>	<i>Normal</i>	<i>Uniform</i>	<i>Normal</i>	<i>Uniform</i>	<i>Normal</i>
Beale	Box	1.00	1.00	1.00	1.00	1.00	1.00
	CH ⁺	0.09	0.35	0.16	0.72	0.86	0.79
	CH _{0.05} ⁺	0.83	0.85	0.87	0.85	0.68	0.65
	CH _{0.10} ⁺	0.99	1.03	1.00	0.96	0.75	0.81
Griewank	Box	1.00	1.00	1.00	1.00	1.00	1.00
	CH ⁺	0.49	0.53	1.18	0.90	0.89	1.08
	CH _{0.05} ⁺	0.44	0.69	1.04	0.82	0.95	0.56
	CH _{0.10} ⁺	0.51	0.83	1.02	0.82	0.96	0.58
Peaks	Box	1.00	1.00	1.00	1.00	1.00	1.00
	CH ⁺	1.02	0.68	1.00	0.94	0.93	0.92
	CH _{0.05} ⁺	1.89	1.18	1.00	0.90	0.77	0.54
	CH _{0.10} ⁺	1.88	1.24	1.00	1.00	0.78	0.62
Powell	Box	1.00	1.00	1.00	1.00	1.00	1.00
	CH ⁺	0.09	0.17	0.15	0.23	0.78	0.63
	CH _{0.05} ⁺	0.44	0.50	0.42	0.36	0.87	0.59
	CH _{0.10} ⁺	0.86	0.99	0.82	0.67	0.93	0.65
Qing	Box	1.00	1.00	1.00	1.00	1.00	1.00
	CH ⁺	0.41	0.45	0.79	0.95	0.70	0.54
	CH _{0.05} ⁺	0.31	0.41	0.69	0.92	0.51	0.36
	CH _{0.10} ⁺	0.48	0.57	0.67	0.92	0.50	0.39
Quintic	Box	1.00	1.00	1.00	1.00	1.00	1.00
	CH ⁺	0.13	0.06	1.00	0.25	0.92	0.74
	CH _{0.05} ⁺	0.30	0.16	0.70	0.18	0.86	0.75
	CH _{0.10} ⁺	0.60	0.27	0.78	0.21	0.93	0.84
Rastrigin	Box	1.00	1.00	1.00	1.00	1.00	1.00
	CH ⁺	0.68	0.49	1.13	1.54	0.70	0.63
	CH _{0.05} ⁺	0.77	0.71	1.07	1.31	0.66	0.77
	CH _{0.10} ⁺	0.82	0.82	1.04	1.13	0.70	0.82

Table 4: Errors for all 37,800 experiments, grouped by function, type of error, and sampling rule. In each group of four errors corresponding to four validity domains, the errors are scaled so that BOX has value 1.00 and the smallest error is shown in a bold font in order to facilitate comparison among the different validity domains.

This amounted to 900 experiments with five validity domains tested in each experiment. We also mention that, for *MLPRegressor*, we chose a hidden layer size of 50 for these tests, compared to 30 in the results cataloged in Tables 3 and 4. All other details of our *MLPRegressor* implementation remained the same.

The results are shown in Table 5, which displays similar patterns as seen in the previous Tables 3 and 4. In other words, we see consistent performance of CH^+ and CH_ϵ^+ in larger dimensions for this function.

Function	Validity Domain	Median Function Value Error	Median Optimal Value Error	Median Optimal Solution Error
Rotated Hyper Ellipsoid	Box	1.00	1.00	1.00
	CH	0.09	0.12	0.26
	CH^+	0.01	0.03	0.21
	$\text{CH}_{0.05}^+$	0.05	0.06	0.25
	$\text{CH}_{0.10}^+$	0.10	0.15	0.32

Table 5: Errors for all 900 experiments, grouped by function, type of error, and sampling rule. In each group of five errors corresponding to five validity domains, the errors are scaled so that BOX has value 1.00 and the smallest error is shown in a bold font in order to facilitate comparison among the different validity domains.

5 Two Stylized Optimization Models

In this section, we examine the performance of our extended validity domain CH^+ in the context of two stylized optimization problems. In both cases, we see that CH^+ and its enlarged variants more effectively manage the function value and feasibility errors.

5.1 A simple nonlinear optimization

Consider the true optimization problem $v^* := \min_{x \in \mathbb{R}^{n_1}} \{c^T x : \|x\| \leq 1, x \in [-1, 1]^{n_1}\}$ where $c \in \mathbb{R}^{n_1}$ is an arbitrary vector satisfying $\|c\| = 1$. It is easy to see that $v^* = -1$ and the unique optimal solution is $x^* = -c$. This is an instance of (2) with $f(x) := c^T x$, $X := [-1, 1]^{n_1}$, g nonexistent, $\theta(y) := y - 1$, and $h(x) := \|x\|$. In this subsection, we will consider small values of n_1 , specifically $n_1 \in \{5, 10\}$.

We conduct a single experiment by randomly generating c uniformly on the surface of the Euclidean unit ball, sampling $N = 1,000$ points x such that the Euclidean norm of x is uniform in $[0.5, 1.5]$, and evaluating $h(x) + 0.05\epsilon$ on all samples, where ϵ is a standard normal random variable. We train the function \hat{h} using the Python function *MLPRegressor* just as in Section 4 using the same hyperparameter choices.

For a single experiment, we then test the validity domains BOX, CH, CH^+ , $\text{CH}_{0.05}^+$, and $\text{CH}_{0.10}^+$, where, as in Section 4, the ϵ values are in fact scaled by the diameter of the ground set X . We also attempted to test ISOFOR, but its overall run times were on the order of hours per experiment compared with seconds. So we decided not to test ISOFOR in order to streamline our computational testing.

n_1	Validity Domain	Median Function Value Error	Median Optimal Value Error	Median Optimal Solution Error	Median Feasibility Error
5	Box	1.00	1.00	1.00	1.00
	CH	0.90	1.00	0.94	0.90
	CH ⁺	0.18	2.00	1.00	0.00
	CH _{0.05} ⁺	0.20	0.86	0.86	0.00
	CH _{0.10} ⁺	0.30	0.38	0.71	0.24
10	Box	1.00	1.00	1.00	1.00
	CH	0.19	0.14	0.51	0.19
	CH ⁺	0.05	1.22	0.71	0.00
	CH _{0.05} ⁺	0.07	0.77	0.64	0.00
	CH _{0.10} ⁺	0.10	0.40	0.55	0.06

Table 6: Median errors of the Box, CH, CH⁺, CH_{0.05}⁺, and CH_{0.10}⁺ validity domains in the stylized simple nonlinear optimization model, grouped by dimension n_1 . The median is taken over 100 experiments, and each sub-grouping of five errors is scaled so that BOX has error 1.00 with a bold font used to highlight the minimum error in that sub-group.

CH⁺ is constructed as the convex hull of F_N^+ according to (11), where F_N^+ is the set of feasible samples in the full extended space defined by (10). Despite the fact that the evaluations of h are noisy, we include a point x in F_N^+ if $\|x\| + 0.05\epsilon \leq 1$, where ϵ is the added noise (not to be confused with ε in the definition of CH _{ε} ⁺). In particular, the noise affects the construction of the set F_N^+ . Because x has been sampled with radius uniform in $[0.5, 1.5]$, the expected cardinality of F_N^+ is approximately $N/2$.

For $n_1 = 5$, we run 100 experiments and show the median errors in Table 6, and then we repeat the same experiments for $n_1 = 10$. Similar to Tables 3-5 in Section 4, we collect the results in Table 6, grouped by n_1 . For each sub-grouping of five errors, we scale such that BOX has error 1.00. First and foremost, we see that CH⁺ achieves much better function and feasibility errors than BOX and CH. Secondly, there appears to be a tradeoff between achieving good function and feasibility errors versus achieving good optimal value and optimal solution errors. For example, CH⁺ does best in terms of function and feasibility errors but worst in terms of the optimal value and optimal solution errors. In contrast, CH_{0.10}⁺ sacrifices function and feasibility errors while improving the other two errors. In this particular case, CH_{0.05}⁺ seems to strike a nice balance between all error types.

Since CH⁺ has been constructed as the convex hull of F_N^+ , which includes only feasible samples—and since the true feasible set is convex—it is perhaps not surprising that CH⁺ achieves the smallest feasibility error. So we repeated the same tests but for the convex hull of D_N^+ defined in (9), i.e., both feasible and infeasible sample points are included. This is in fact an example of using a convex hull CH', which is sandwiched between CH and the full CH⁺, as discussed in Section 3.2. The results showed the same trends, in particular that CH' and its enlarged variants manage feasibility well in this case, too.

5.2 A price optimization

We also test CH^+ on a stylized price optimization problem, which serves as a preview of the case study in Section 6. Imagine a company with two substitute products, labeled as products 1 and 2. Product 1 is the low-price product, and product 2 is the high-price product. Demands d_1 and d_2 for the respective products are functions of their prices, p_1 and p_2 . The company would like to set prices so as to maximize revenue under various constraints. The true model is:

$$\max_{p_1, p_2} \quad d_1 p_1 + d_2 p_2 \quad (13a)$$

$$\text{s.t.} \quad p_1 \in [6.5, 9.5] \quad p_2 \in [7.5, 10.5] \quad p_1 + 0.5 \leq p_2 \leq p_1 + 1.5 \quad (13b)$$

$$d_1 = 10^7 p_1^{-3.2} p_2^{0.5} \quad d_2 = 10^7 p_1^{1.5} p_2^{-2.2} \quad d_1 + d_2 \leq 1.8 \times 10^6. \quad (13c)$$

Here, the objective function (13a) is the total revenue of the company; constraint (13b) describes the allowable prices; and constraint (13c) describes the demand functions for products 1 and 2 as well as a cap on demand corresponding what can actually be sold. Gurobi solves this problem efficiently, reporting a true optimal solution of $(p_1^*, p_2^*) = (9.50, 10.31)$. At optimality, the demand constraint is active.

To perform numerical tests, we replace the demand functions with learned models and enforce different validity-domain constraints. To learn the demand functions, we uniformly sampled 1,000 points in the box domain of (p_1, p_2) , evaluated the corresponding (d_1, d_2) without noise, and learned the approximation (\hat{d}_1, \hat{d}_2) using a simple quadratic regression. (Note that if the data had been learned with a log-log regression, then the learned model would be exact.) We tested BOX, CH, IsoFOR, CH^+ , $\text{CH}_{0.05}^+$, and $\text{CH}_{0.10}^+$. The convex hull used for CH is taken over the 2-dimensional samples (p_1, p_2) , while our extended convex hull CH^+ is taken over the 4-dimensional samples (p_1, p_2, d_1, d_2) . In particular, we make use of both feasible and infeasible samples, and we do not append the objective value. As such, we actually implement a variant CH' of CH^+ , which is sandwiched between CH and CH^+ , as discussed in Section 3.2. However, in our discussion here, we keep the name CH^+ for simplicity.

The median errors of the validity domains are summarized in Table 7, where as before we scale the BOX errors to 1.00. We see clearly that CH^+ achieves much lower function value and feasibility errors than either BOX or CH, although the optimal value and optimal solution errors are notably higher. We believe that overall this constitutes an advantage of CH^+ and its variants because, generally speaking, one cannot have true optimality without true feasibility. Regarding the enlarged versions of CH^+ , while we have chosen $\varepsilon = 0.05$ and $\varepsilon = 0.10$ to be consistent with the experiments in prior sections of the paper, it would appear that both enlarged validity domains achieve the same errors as BOX, indicating that these values of ε are too loose in this particular case.

6 A Case Study

In this section, we investigate an [avocado-price optimization model](#) recently described by Gurobi Optimization, the makers of [Gurobi \(2023\)](#). The goal is to set the prices and supply

	Median Function Value Error	Median Optimal Value Error	Median Optimal Solution Error	Median Feasibility Error
BOX	1.00	1.00	1.00	1.00
CH	0.99	1.06	1.07	0.99
ISOFOR	0.67	3.81	4.97	0.66
CH ⁺	0.16	10.81	16.92	0.11
CH _{0.05} ⁺	1.00	1.00	1.00	0.99
CH _{0.10} ⁺	1.00	1.00	1.00	1.00

Table 7: Median errors of all validity domains in the price optimization model. The median is taken over 10,000 experiments, and each sub-grouping of three errors is scaled so that BOX has error 1.00 with a bold font used to highlight the minimum error in that sub-group.

quantities for avocados across eight regions of the United States while incorporating transportation and other costs and maximizing the total national profit from avocado sales. A critical component of the model is the relationship between avocado prices and the demand for avocados in each region. Gurobi Optimization proposed to learn the demand function using ML techniques based on observed sales data. Then the demand function could be embedded in the optimization model, thus creating a case study for constraint learning (CL). We revisit this study in light of our extended validity domain CH⁺ proposed in Section 3.

Note that, in this case study, there is no ground truth, and so we are unable to measure the errors defined in Section 2.2. Instead, we seek experimental insights from this case.

6.1 Avocado price model

The eight regions are indexed by $r \in \{1, \dots, 8\}$, and the total units of avocado imported into a single port in the United States is denoted by B . The per-unit cost of waste is α , which is independent of the region, and the transportation cost from the port to region r is β_r . The learned demand function for avocados in region r at sales price p is denoted $\hat{d}(p, r)$. The optimization variables are the unit price p_r and units supplied x_r for region $r \in \{1, \dots, 8\}$. Auxiliary variables s_r and w_r represent the units sold and units wasted per region, respectively. The model formulation is

$$\max \sum_{r=1}^8 (p_r s_r - \alpha w_r - \beta_r x_r) \quad (14a)$$

$$\text{s.t.} \quad \sum_{r=1}^8 x_r = B \quad (14b)$$

$$s_r \leq \min\{x_r, \hat{d}(p_r, r)\} \quad w_r = x_r - s_r \quad p_r \geq 0, x_r \geq 0 \quad \forall r \in \{1, \dots, 8\}. \quad (14c)$$

The nonconvex, bilinear objective (14a) calculates profit by subtracting the cost of shipping and waste from the revenue over all regions. Constraint (14b) ensures that the total units supplied equals the import quantity, and constraint (14c): defines the number of units sold in a region to be no larger than the minimum of supply and predicted demand in the region; sets the number of units wasted to be the difference between number of units supplied and

the number sold; and finally enforces nonnegativity of p_r and x_r . In practice, there may also be upper bounds on p_r and x_r .

6.2 Dataset and predictive model

The dataset to learn $\hat{d}(p, r)$ has been prepared by Gurobi Optimization from two sources, and the combined dataset is hosted at [Gurobi Optimization’s GitHub account](#). One source is the [Hass Avocado Board](#) (HAB), from which Gurobi Optimization has retrieved data for the years 2019 to 2022, and the second source is [Kaggle](#), which hosts data originating from HAB for the years 2015 to 2018. The total time span of the data is thus from 2015 to 2022. Note that, since true demand is not observed directly, sales data is used instead as the best available proxy of demand, and we do not consider other issues such as non-stationarity of demand over time.

The data has been cleaned and processed such that each observation includes a date (indicating the start of a calendar week), a seasonality indicator for that specific week (corresponding to 1 for peak and 0 for off-peak), the region, the number of units of avocados sold (in millions of units), and the average price per unit. For this data set, the average units sold over all observations is 3.9 million, the average price is \$1.14, and the average revenue is \$4.23 million; these values describe a typical week in a typical region. Aggregating up to the entire United States, in a typical week, the average units sold is 31.0 million at an average price of \$1.14 for an average revenue of \$33.85 million. Restricting to just off-peak data, in a typical off-peak week, the average units sold across the U.S. is 28.6 million at an average price of \$1.14 for an average revenue of \$31.10 million.

Following Gurobi Optimization’s example, we use a gradient boosting regressor implemented using the Python’s *scikit-learn* package to learn the demand function $\hat{d}(p, r)$ based on four features of the data: year (not the specific week), the seasonality indicator, region, and price. For convenience, we label these features *Year*, *Peak*, *Region*, and *Price*, respectively. The label for the response variable in the data is *Units Sold*.

Figure 3 depicts eight panels, one for each region in the data. Here, the regions are labeled with their descriptive names, e.g., *Northeast* and *Plains*, as opposed to their index numbers $r \in \{1, \dots, 8\}$. Each panel depicts a scatter plot of light-color, circular dots, which show the observed data, where the horizontal axis is *Price* and the vertical axis is *Units Sold* (in millions). We note that each panel is plotted over the region $0.6 \leq p_r \leq 2.0$, representing a unit price between \$0.60 and \$2.00 for all regions. In addition, in each panel the learned function $\hat{d}(p_r, r)$ with fixed values *Year* = 2023 and *Peak* = 0 is plotted as a curve through the data; the vertical axis is labeled as *Predicted Demand*. (Each panel also contains additional information in the form of four additional markers, which we explain in the next subsection.) We observe the expected inverse relationship between price and demand. Further, the fitted demand function does a reasonably good job capturing the relationship between price and demand in each of the eight regions.

6.3 Effect of validity domains

We now solve (14) for various validity domains. All experiments are executed using Gurobi 12.0 as the optimization solver. In these optimizations, we fix *Year* = 2023, i.e., predicting

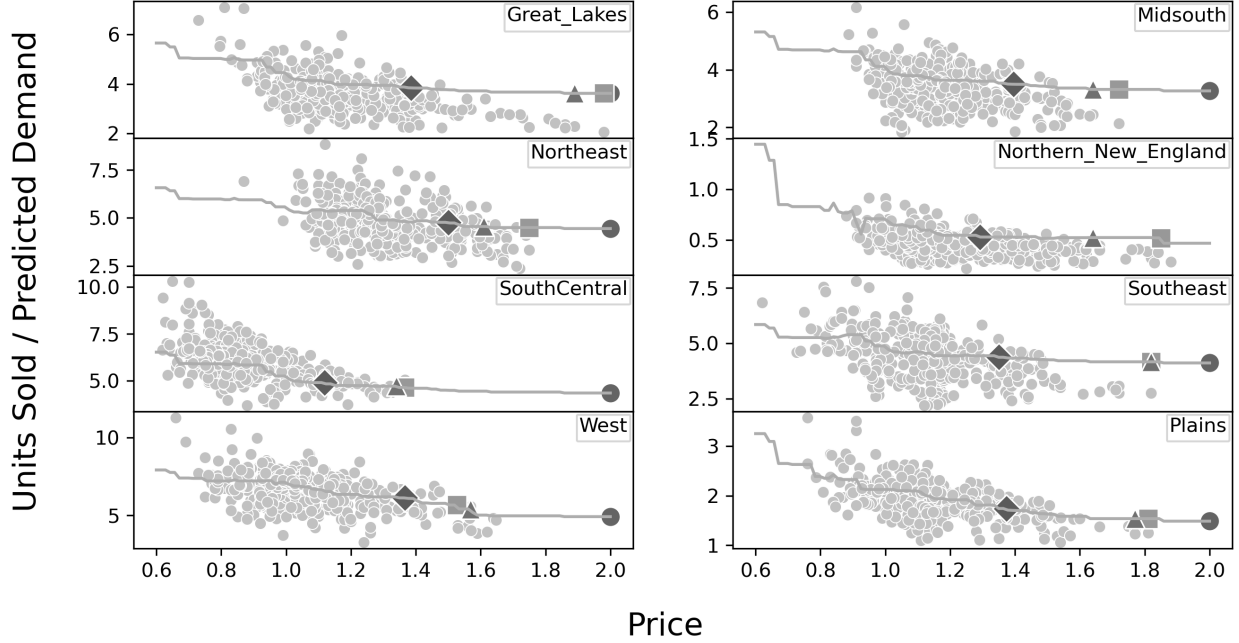


Figure 3: Avocado data (light circle dots), predicted demand function (curve), and four optimal solutions. The default (i.e., no validity domain) is circle, BOX is square, CH is triangle, and CH⁺ is diamond.

into the next year beyond the data, and we also fix $Peak = 0$. Hence, we are optimizing prices for a single off-peak week in 2023. We set the remaining model parameters as Gurobi Optimization has done with $B = 30$, $\alpha = 0.1$, and the following transportation coefficients β_r : Great Lakes = 0.3, Midsouth = 0.1, Northeast = 0.4, Northern New England = 0.5, South Central = 0.3, Southeast = 0.2, West = 0.2, and Plains = 0.2. Gurobi was able to solve all instances to its default optimality tolerance in just a few minutes on a modern laptop.

When (14) is optimized with a uniform upper bound of $p_r \leq 2.0$ for all regions, the optimal value is \$45.84 million. A corresponding best solution is depicted with circle markers in the panels of Figure 3. In particular, for all but the Northern New England region, the best solution sets $p_r = 2.0$, i.e., the price is at the upper bound of \$2.00. Particularly striking is that the solution is visually quite far from the observed data. One may ask if this solution is trustworthy given that the demand predictions are likely to be less reliable away from the observed data.

We next test the three validity domains BOX, CH, and CH⁺. The results for ISOFOR and CH _{ϵ} ⁺ are included in the online supplemental material. Note that all validity domains are considered with respect to the prices p but not the shipped quantities x because the historical data on x are not available. Our hope is that these validity domains can help alleviate the uncertainty inherent in the default solution just mentioned. We would also like to compare and contrast these three validity domains in this setting. Our observations are summarized as follows:

- BOX: For each r independently, we constrain p_r to be within its observed minimum

and maximum values; see (6). The best reported value is \$39.51 million, and the corresponding prices per region are shown in Figure 3 as square markers. Compared to the circle markers of the default solution, the square markers are “pulled back” much closer to the data. It is clear that BOX does not allow as much extrapolation in the optimal solution, and hence one can expect the predicted demand to be more accurate. Hence, one can expect the final profit number to be more reliable.

- CH: This validity domain is defined to be the convex hull in \mathbb{R}^8 of the 8-tuples $(p_1^{(i)}, \dots, p_8^{(i)})$, where $i = 1, \dots, N$ indexes over all observations; see (7). The value is \$38.02 million, and Figure 3 shows the prices as triangle markers. As with BOX, the prices are visually closer to observed price data, and hence one can expect the overall optimization result to be more reliable.
- CH⁺: Finally, we consider our extended validity domain, which enforces that the concatenated variables and demand predictions $(p_1, \dots, p_8, \hat{d}(p_1, 1), \dots, \hat{d}(p_8, 8))$ lie in the convex hull of the observed data $(p_1^{(i)}, \dots, p_8^{(i)}, d_1^{(i)}, \dots, d_8^{(i)})$ where $i = 1, \dots, N$; see (11). This convex hull lies in \mathbb{R}^{16} . In particular, because the historical shipped quantities x are not available in the historical dataset, our CH⁺ here is actually a variant of the form CH' discussed in Section 3.2. The optimized prices are shown as diamond markers in Figure 3, and the value is \$32.54 million. Although this value is significantly less than the preceding objective values (compare to the value of \$38.02 for CH, for example), the positions of the diamonds is considerably closer to the original data set. In fact, based on the two-dimensional panels in Figure 3, it appears that the optimal prices are actually embedded inside the data, although this may be a visual artifact of the projection of a 16-dimensional image down to eight individual 2-dimensional scatter plots. In any case, one can expect that the predictions $\hat{d}(p, r)$ are the most reliable for this validity domain, making the overall optimization more reliable.

As a final comment, we recall that the empirical data shows an average revenue of \$31.10 million in a typical off-peak week; see the discussion in Section 6.2. The final optimal value of \$32.54 million for CH⁺ is certainly in line with this empirical average and constitutes an increase of 4.6%.

7 Conclusions

We have studied the use of validity domains to reduce the errors associated with the constraint-learning (CL) framework. Based on the intuition of using the convex hull to learn both the data set and the optimization problem itself, we have proposed a new extended validity domain, called CH⁺. Our numerical studies have shown that CH⁺, compared to other common methods in the literature, is competitive in terms of computational effort and tends especially to reduce the function value and feasibility errors. Beyond stylized numerical results, the avocado case study has shown the applicability and adaptability of our approach to real-world situations.

While we would certainly advocate the use of CH^+ as a valuable option in many situations (e.g., to reduce the function value and feasibility errors, as just mentioned), it is of course natural to ask which validity domain is best suited for any specific situation. For example, if one is primarily interested in estimating the true optimal value v^* , then choosing the validity domain that minimizes optimal value error would be the best choice. At this time, we unfortunately lack a solid set of practical guidelines for choosing the best validity domain in common situations. Nevertheless, we hope that this research contributes to the growing literature on how to make the best use of constraint learning.

We mention some opportunities for future research. First, the model (1) based on \hat{h} may be infeasible in general, and when validity domains such as CH or CH^+ are also enforced, infeasibility will be, in a sense, even more likely. It will be interesting to investigate the underlying properties that make (1) feasible, to understand when CH^+ maintains this feasibility, and if not, then to develop alternate validity domains that do maintain it. Second, we have used the convex hull in an extended space in part because CH in the original space is well-studied and possesses good properties. Of course, there exist other techniques for creating validity domains in the original space as described in Section 2. One future idea to explore is whether these other techniques, similar to CH^+ , can also be effective in the extended space. Third, while current validity domains in the literature are model-agnostic, including CH^+ , future research could investigate validity domains specific to the optimization model as well as the ML model \hat{h} used to learn the constraints. Such validity domains might be termed “model aware.” Lastly, this paper has only dealt with regression models \hat{h} . Exactly how CH^+ might work in the case of classification is an extremely interesting question, which, if successful, could increase the versatility of the ideas in this paper.

Acknowledgments

The authors wish to thank Paul Grigas for discussions on the relationship between constraint learning and contextual optimization. The authors are also in debt to the anonymous associate editor and three referees for invaluable suggestions that have greatly improved the paper.

References

- Randall Balestrieri, Jerome Pesenti, and Yann LeCun. Learning in high dimension always amounts to extrapolation, 2021. URL <https://arxiv.org/abs/2110.09485>.
- Yoshua Bengio, Andrea Lodi, and Antoine Prouvost. Machine learning for combinatorial optimization: a methodological tour d’horizon. *European Journal of Operational Research*, 290(2): 405–421, 2021.
- David Bergman, Teng Huang, Philip Brooks, Andrea Lodi, and Arvind U. Raghunathan. JANOS: an integrated predictive and prescriptive modeling framework. *INFORMS J. Comput.*, 34(2):807–816, 2022. ISSN 1091-9856,1526-5528. doi: 10.1287/ijoc.2020.1023. URL <https://doi.org/10.1287/ijoc.2020.1023>.
- Francesco Ceccon, Jordan Jalving, Joshua Haddad, Alexander Thebelt, Calvin Tsay, Carl D. Laird,

- and Ruth Misener. OMLT: optimization & machine learning toolkit. *J. Mach. Learn. Res.*, 23:Paper No. [349], 8, 2022. ISSN 1532-4435,1533-7928.
- Pierre Courrieu. Three algorithms for estimating the domain of validity of feedforward neural networks. *Neural Networks*, 7(1):169–174, 1994.
- Allegra De Filippo, Michele Lombardi, and Michela Milano. Methods for off-line/on-line optimization under uncertainty. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, IJCAI-2018. International Joint Conferences on Artificial Intelligence Organization, July 2018. doi: 10.24963/ijcai.2018/177. URL <http://dx.doi.org/10.24963/ijcai.2018/177>.
- Adejuyigbe O Fajemisin, Donato Maragno, and Dick den Hertog. Optimization with constraint learning: a framework and survey. *European Journal of Operational Research*, 2023.
- Gurobi. Gurobi Optimizer Reference Manual, 2023. URL <https://www.gurobi.com>.
- James Kotary, Ferdinando Fioretto, Pascal Van Hentenryck, and Bryan Wilder. End-to-end constrained optimization learning: A survey. *International Joint Conference on Artificial Intelligence*, 2021. doi: 10.24963/ijcai.2021/610. URL <https://par.nsf.gov/biblio/10337537>.
- Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *2008 eighth ieee international conference on data mining*, pages 413–422. IEEE, 2008.
- Wen Long, Tiebin Wu, Ximing Liang, and Songjin Xu. Solving high-dimensional global optimization problems using an improved sine cosine algorithm. *Expert Systems with Applications*, 123:108–126, 2019. ISSN 0957-4174. doi: <https://doi.org/10.1016/j.eswa.2018.11.032>. URL <https://www.sciencedirect.com/science/article/pii/S0957417418307528>.
- Donato Maragno, Holly Wiberg, Dimitris Bertsimas, Ş İlker Birbil, Dick den Hertog, and Adejuyigbe O Fajemisin. Mixed-integer optimization with constraint learning. *Operations Research*, 73(2):1011–1028, 2025.
- Miten Mistry, Dimitrios Letsios, Gerhard Krennrich, Robert M Lee, and Ruth Misener. Mixed-integer convex nonlinear optimization with gradient-boosted trees embedded. *INFORMS Journal on Computing*, 33(3):1103–1119, 2021.
- Utsav Sadana, Abhilash Chenreddy, Erick Delage, Alexandre Forel, Emma Frejinger, and Thibaut Vidal. A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research*, 320(2):271–289, 2025. ISSN 0377-2217. doi: <https://doi.org/10.1016/j.ejor.2024.03.020>. URL <https://www.sciencedirect.com/science/article/pii/S0377221724002200>.
- Artur M Schweidtmann, Jana M Weber, Christian Wende, Linus Netze, and Alexander Mitsos. Obey validity limits of data-driven models through topological data analysis and one-class classification. *Optimization and engineering*, 23(2):855–876, 2022.
- Chenbo Shi, Mohsen Emadikhiay, Leonardo Lozano, and David Bergman. Constraint learning to define trust regions in optimization over pre-trained predictive models. *INFORMS Journal on Computing*, 36(6):1382–1399, 2024. doi: 10.1287/ijoc.2022.0312. URL <https://doi.org/10.1287/ijoc.2022.0312>.
- Sonja Surjanovic and Derek Bingham. Virtual library of simulation experiments: Test functions and datasets. <https://www.sfu.ca/~ssurjano/optimization.html>, 2023. Accessed: 2024-06-11.
- Bo Tang and Elias B Khalil. Pyepo: A pytorch-based end-to-end predict-then-optimize library for linear and integer programming. *Mathematical Programming Computation*, 16(3):297–335, 2024.

Vincent Tjeng, Kai Y. Xiao, and Russ Tedrake. Evaluating robustness of neural networks with mixed integer programming. In *International Conference on Learning Representations*, 2019.
URL <https://openreview.net/forum?id=HyGIIdRqtm>.

Supplementary Material

S1 Optimization Run Times for Sections 4.3 and 4.4

For the experiments described in Sections 4.3 and 4.4, we show in Tables S1 and S2 the median setup and solve times for all combinations of machine-learning techniques and validity domains. By *setup time*, we mean the time required to build and pass the optimization to Gurobi, including the critical constraint which sets y equal to the output of the learned function $\hat{h}(x)$, where \hat{h} has already been trained and stored in memory. Although computation times are not the main focus in this paper, we see that ISOFOR requires more time in general, while CH^+ can take between 1-4 times as long as BOX and CH. Furthermore, for solve times specifically, the enlarged versions CH_ϵ^+ take times that are more similar to BOX and CH than to CH^+ itself. While we are not exactly sure why this occurs, this could represent a potential advantage of CH_ϵ^+ in practice compared to the default CH^+ .

ML Technique	BOX	CH	ISOFOR	CH^+	$\text{CH}_{0.05}^+$	$\text{CH}_{0.10}^+$
<i>RandomForestRegressor</i>	1.23	1.23	41.25	1.23	0.99	0.97
<i>GradientBoostingRegressor</i>	1.28	1.27	41.47	1.27	1.01	1.01
<i>MLPRegressor</i>	0.01	0.02	40.47	0.02	0.02	0.02

Table S1: Median optimization setup times (in seconds) over all 226,800 experiments, grouped by ML technique and validity domain.

ML Technique	BOX	CH	ISOFOR	CH^+	$\text{CH}_{0.05}^+$	$\text{CH}_{0.10}^+$
<i>RandomForestRegressor</i>	0.86	1.31	3.56	1.48	0.81	0.81
<i>GradientBoostingRegressor</i>	5.40	6.70	14.35	14.76	5.94	4.90
<i>MLPRegressor</i>	0.09	0.23	24.24	0.39	0.45	0.46

Table S2: Median optimization solve times (in seconds) over all 226,800 experiments, grouped by ML technique and validity domain.

S2 Mean Errors for Section 4.3

Recall that Table 3 in Section 4.3 reports *median* error values. To complement Table 3, we now present the exact same analysis except that we calculate *mean* values in Table S3. We find that the mean values reveal a similar pattern as the median values. In particular, CH^+ produces very competitive results in terms of the function value error. In fact, in all but one case, CH^+ has the smallest mean function value error. For the mean optimal value errors and mean optimal solution errors, CH^+ is also competitive with other validity domain techniques perform slightly better in certain cases.

Function	Validity Domain	Mean Function Value Error		Mean Optimal Value Error		Mean Optimal Solution Error	
		<i>Uniform</i>	<i>Normal</i>	<i>Uniform</i>	<i>Normal</i>	<i>Uniform</i>	<i>Normal</i>
Beale	BOX	1.00	1.00	1.00	1.00	1.00	1.00
	CH	0.99	0.80	1.01	1.09	1.00	0.86
	IsoFOR	0.73	0.50	0.78	0.49	0.88	0.64
	CH ⁺	0.21	0.31	0.81	0.77	0.42	0.52
Griewank	BOX	1.00	1.00	1.00	1.00	1.00	1.00
	CH	0.57	0.80	0.82	0.92	1.07	1.00
	IsoFOR	0.63	0.76	0.30	0.54	1.68	1.00
	CH ⁺	0.38	0.31	0.71	0.71	1.25	0.97
Peaks	BOX	1.00	1.00	1.00	1.00	1.00	1.00
	CH	1.00	1.00	0.97	1.00	1.00	1.00
	IsoFOR	1.21	2.13	0.92	0.50	1.10	1.00
	CH ⁺	1.01	0.59	0.93	0.85	1.01	0.99
Powell	BOX	1.00	1.00	1.00	1.00	1.00	1.00
	CH	0.99	0.86	0.95	0.88	0.99	0.91
	IsoFOR	0.90	0.63	0.69	0.55	0.99	0.88
	CH ⁺	0.50	0.29	0.79	0.58	0.63	0.66
Qing	BOX	1.00	1.00	1.00	1.00	1.00	1.00
	CH	0.87	0.44	0.80	0.75	1.00	1.00
	IsoFOR	1.23	0.49	0.66	0.58	1.00	1.00
	CH ⁺	0.64	0.28	0.68	0.56	0.91	0.96
Quintic	BOX	1.00	1.00	1.00	1.00	1.00	1.00
	CH	0.86	0.54	0.89	0.65	0.95	0.66
	IsoFOR	0.33	0.23	0.69	0.48	0.95	0.59
	CH ⁺	0.20	0.09	0.79	0.49	1.00	0.60
Rastrigin	BOX	1.00	1.00	1.00	1.00	1.00	1.00
	CH	0.79	0.37	0.84	0.53	1.02	1.10
	IsoFOR	0.84	0.45	0.82	0.52	1.03	1.08
	CH ⁺	0.69	0.33	0.71	0.50	1.09	1.18

Table S3: Mean errors for all 37,800 experiments, grouped by function, type of error, and sampling rule. In each group of four errors corresponding to the four validity domains, the errors are scaled so that BOX has value 1.00 in order to facilitate comparison among the different validity domains.

S3 Sensitivity Analysis for the Function Value Error in Sections 4.3 and 4.4

In this section, we investigate the sensitivity of the various validity domains studied in this paper to changes in numerical design choices, e.g., the choice of sampling rule. We are specifically concerned with the sensitivity of the function value error in the experiments of Sections 4.3–4.4.

S3.1 Sensitivity with respect to sampling rule

To investigate the sensitivity with respect to the sampling rule, i.e., *Uniform* versus *Normal*, for each 6-tuple instance of ground truth, sample size, noise level, seed, ML technique, and validity domain, we divide the function value error of *Uniform* by that of *Normal*. This ratio is greater than 1 if *Normal* exhibits better error on this instance and less than 1 if *Uniform* does. Then grouping these results by validity domain, Figure S1 shows the distributions of these ratios for each validity domain. Perhaps unsurprisingly, since *Uniform* samples uniformly in a larger box domain, while *Normal* samples more closely around the global minimum in a smaller ball, we find that the ratios are skewed heavily above 1 for all validity domains, indicating that *Normal* typically exhibits less error.

S3.2 Sensitivity with respect to noise level

Analogous to the previous subsection, we next investigate the sensitivity of the function value error with respect to the noise level σ , broken down by validity domain. Recall that we tested three different values of σ , namely 0.0, 0.1, and 0.2, and so for each instance, we calculate two ratios: (i) the function value error of $\sigma = 0.1$ divided by that of $\sigma = 0.0$; and (ii) the function value error of $\sigma = 0.2$ divided by that of $\sigma = 0.0$. In this sense, the noise-free case is our baseline. A ratio greater than 1 for an instance indicates that the increased noise yields a higher function value error compared to the no-noise case.

Figure S2 shows the corresponding densities for each validity domain. We can see that, in general, each validity domain shows little to no sensitivity to the noise level. This is due to the fact that each density is tightly centered around 0 with small skewnesses and large kurtosis. Table S4 shows the corresponding summary statistics. There is perhaps one exception, CH^+ , which has a larger positive skew, indicating that CH^+ is more susceptible to noise, i.e., a higher noise level may slightly increase the function value errors.

S3.3 Sensitivity with respect to sample size

We next recreate the same analysis from the previous subsection but this time for the three samples sizes $N = 500, 1000, 1500$. In particular, we calculate ratios of the function value errors for $N = 1000$ relative to $N = 500$ and for $N = 1500$ relative to $N = 500$. Figure S3 and Table S5 show the same patterns for N as just discussed for σ . In particular, we see that changes in the sample sizes make little difference in the function value errors. Moreover, in this case, the skewness of CH^+ more closely matches the skewness of the other validity domains.

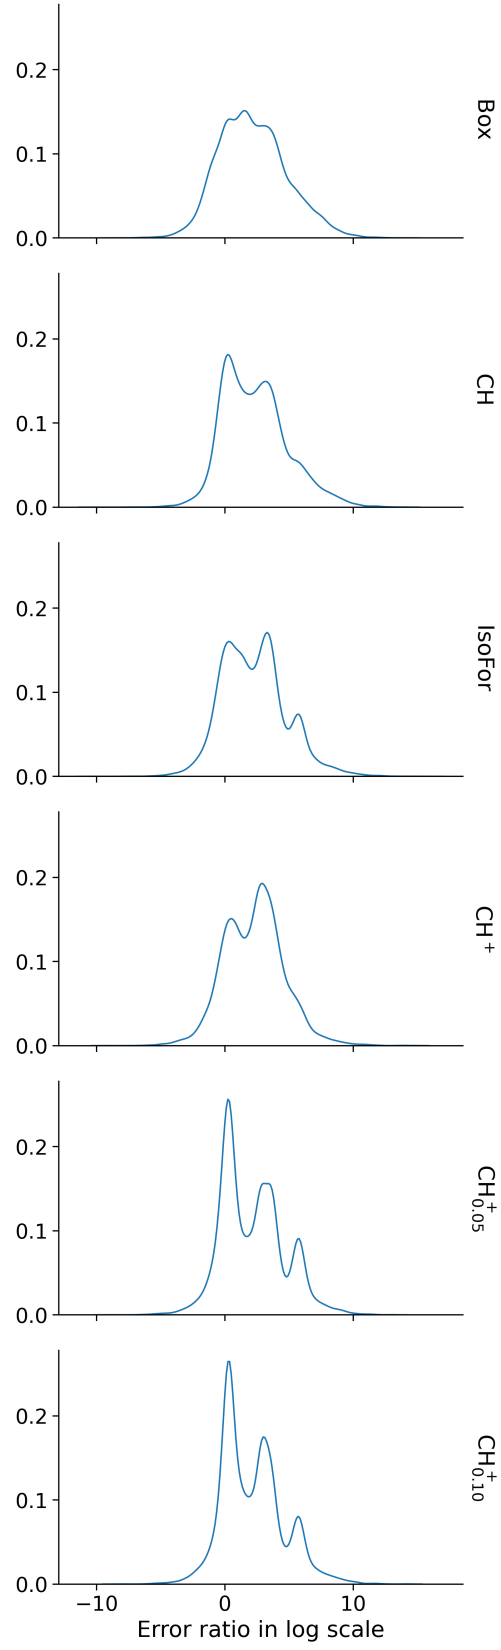


Figure S1: The distribution of the ratio of the function value error obtained by Uniform divided by Normal, grouped by validity domains.

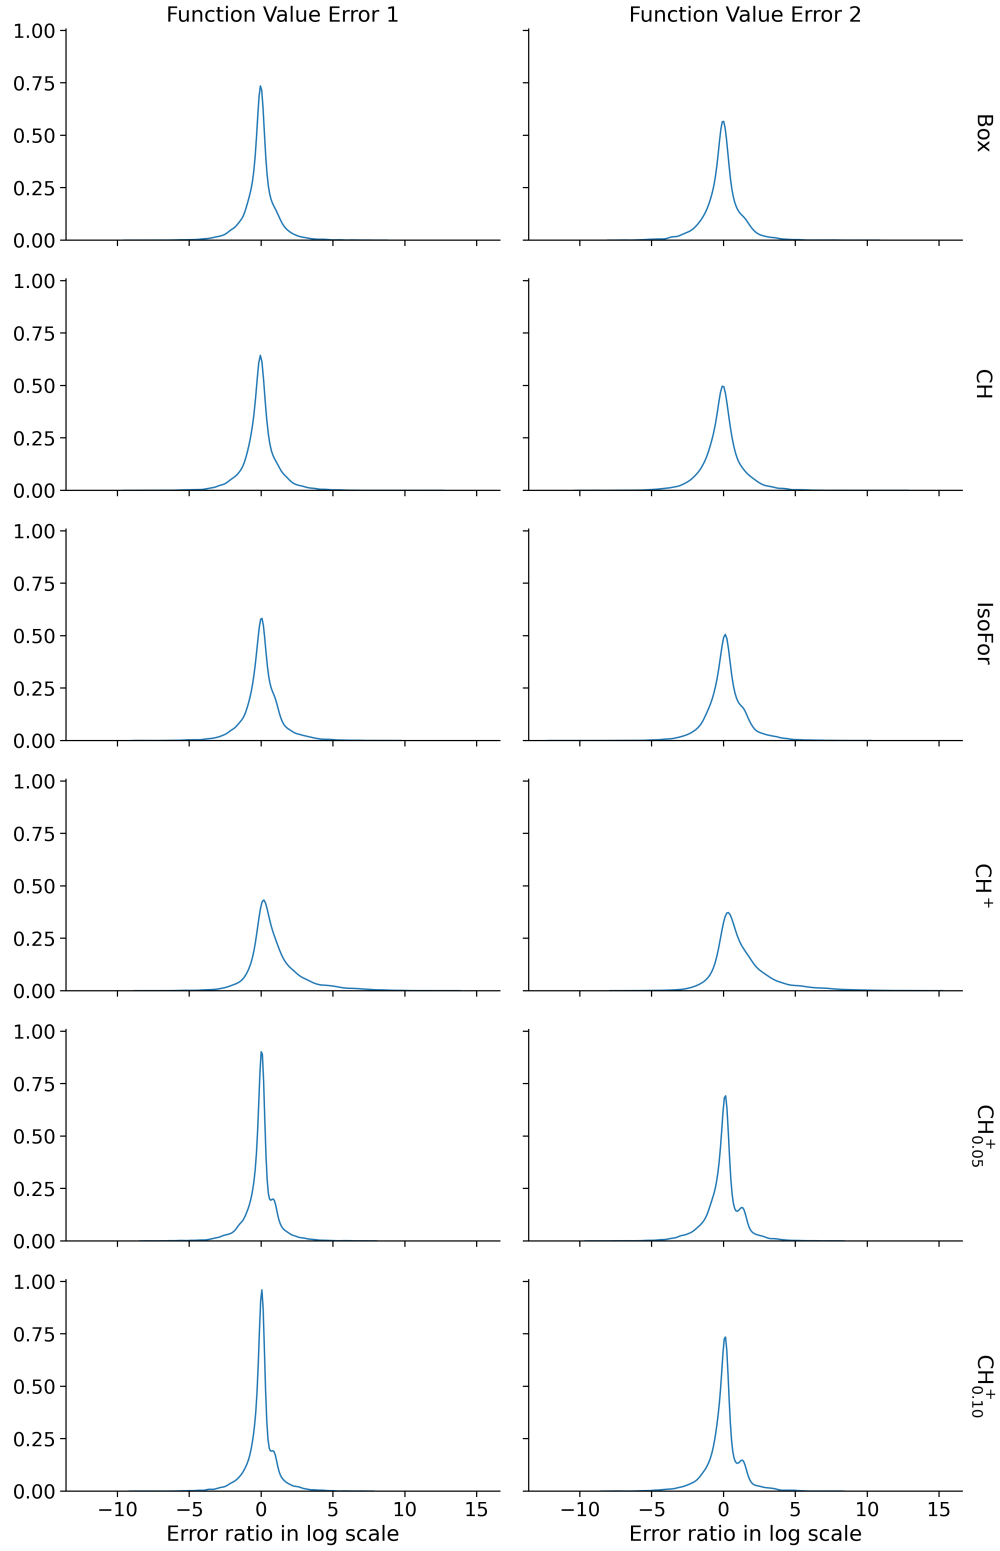


Figure S2: The distribution of the ratio of the function value errors obtained by different noise levels, grouped by validity domains. Error 1 is the ratio between noise level= 0.1 and noise level = 0, Error 2 is the ratio between noise level = 0.2, and noise level = 0.

	Func Val Err 1		Func Val Err 2	
	Skew	Kurt	Skew	Kurt
Box	-0.02	5.36	0.01	4.07
CH	0.09	5.12	0.11	3.96
IsoFor	0.27	4.63	0.31	3.97
CH ⁺	1.36	4.34	1.58	4.62
CH _{0.05} ⁺	-0.26	6.08	-0.29	4.56
CH _{0.10} ⁺	-0.16	6.58	-0.21	4.76

Table S4: The skewness and kurtosis of distribution of the ratio of the errors obtained by different noise levels, grouped by validity domains. Error 1 is the ratio between noise level= 0.1 and noise level = 0, Error 2 is the ratio between noise level = 0.2 and noise level = 0.

	Func Val Err 1		Func Val Err 2	
	Skew	Kurt	Skew	Kurt
Box	-0.10	3.60	-0.11	3.24
CH	0.03	4.49	0.11	3.11
IsoFor	0.04	4.25	-0.10	4.23
CH ⁺	-0.08	7.17	0.18	5.89
CH _{0.05} ⁺	0.09	6.34	0.00	6.58
CH _{0.10} ⁺	0.05	6.57	0.02	6.04

Table S5: The skewness and kurtosis of distribution of the ratio of the errors obtained by different sample sizes, grouped by validity domains. Error 1 is the ratio between sample size= 1000 and sample sizes = 500, Error 2 is the ratio between sample sizes = 1500, and sample size = 500.

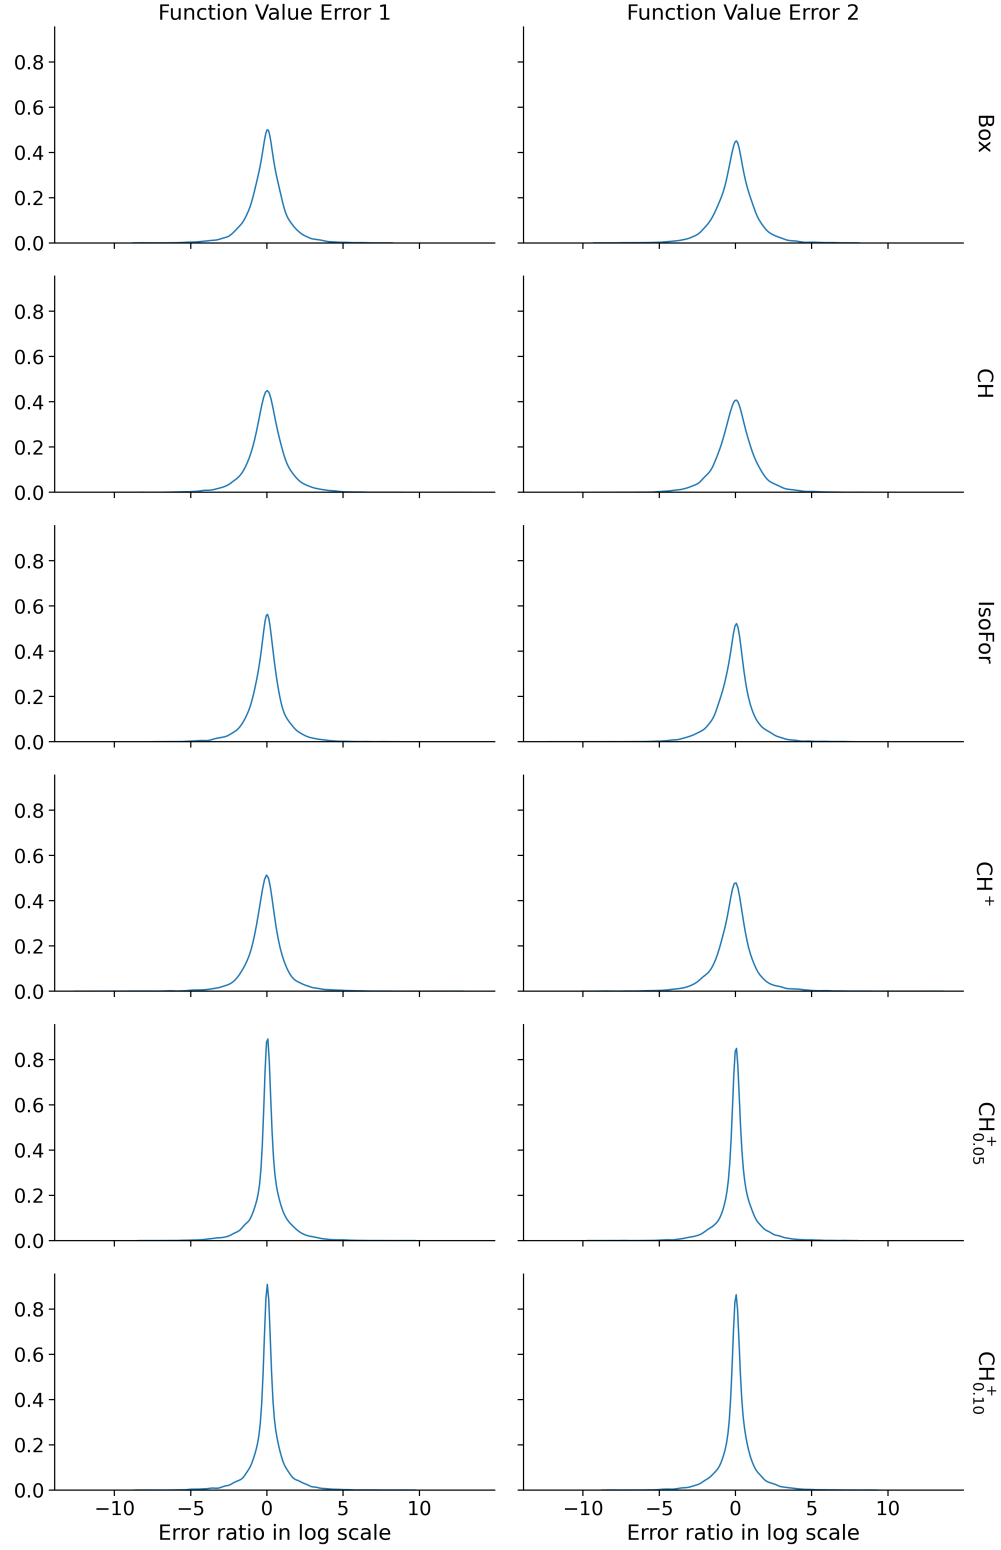


Figure S3: The distribution of the ratio of the function value errors obtained by different sample sizes, grouped by validity domains. Error 1 is the ratio between sample size= 1000 and sample sizes = 500, Error 2 is the ratio between sample sizes = 1000, and sample size = 500.

S3.4 Sensitivity with respect to function type

Our final sensitivity analysis is with respect to whether the ground-truth function is a polynomial (*Beale*, *Qing*, *Quintic*, *Powell*) or not (*Peaks*, *Griewank*, *Rastrigin*). We note in particular that the three non-polynomial functions are based on trigonometric and exponential functions. Figure S4 shows the densities of all function value errors in the experiments of Sections 4.3–4.4 grouped by polynomial versus non-polynomial. We see that the distribution of the errors in the polynomial case are slightly shifted to the left compared to the non-polynomial case, which indicates that, in our experiments, polynomial functions tended to exhibit smaller function value errors.




Figure S4: The distribution of function value errors grouped by whether the function is a polynomial. The blue one is for polynomial functions (*Beale*, *Qing*, *Quintic*, *Powell*), and the orange one is for non-polynomial functions (*Peaks*, *Griewank*, *Rastrigin*).

S4 Additional Experiments for the Avocado Case Study

Figure S5 shows the learned optimal solutions in avocado case study when the validity domains are CH^+ , $CH_{0.05}^+$ and $CH_{0.10}^+$, where in fact the 0.05 and 0.10 values are scaled by the diameter of the box domain for this case. The learned optimal solution from CH^+ is dark blue, from $CH_{0.05}^+$ is orange, and from $CH_{0.10}^+$ is green. We can see the trend that the smaller ε is, the closer the learned optimal solution is to the center of the data points cluster. Note that in some of the graphs, the orange and green point overlap. This result show that, ε is a parameter that controls the conservatism of the validity domain. We also mention that the respective optimal profits were \$32.54 million for CH , \$36.54 million for $CH_{0.05}^+$, and \$38.26 million for $CH_{0.10}^+$.

Finally, we also ran the same experiment with the ISOFOR validity domain. It gave an optimal value of \$34.37 million, which is between CH^+ and $CH_{0.05}^+$.

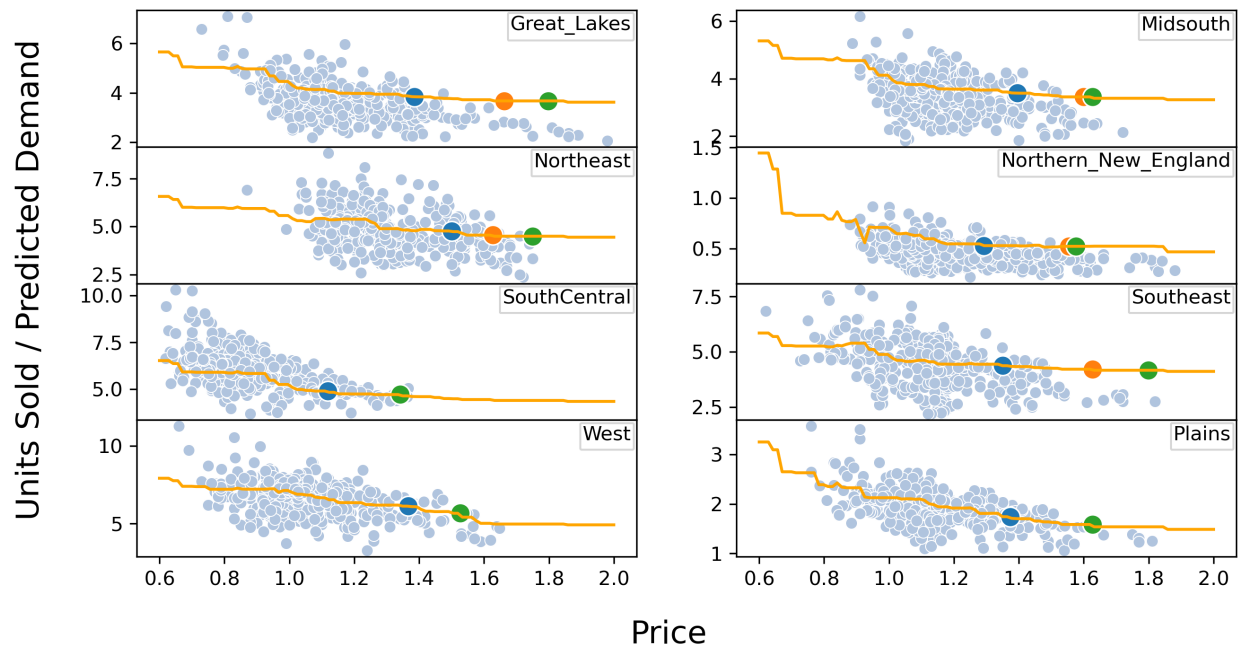


Figure S5: Avocado data (light blue dots), predicted demand function (orange curve), and three optimal solutions obtained from convex hull related validity domains. The CH^+ is dark blue, $CH_{0.05}^+$ is orange, and $CH_{0.10}^+$ is green.