# Learning to optimize with hidden constraints

Aaron Babier, Timothy C. Y. Chan

Mechanical & Industrial Engineering, University of Toronto, Toronto, Canada, {ababier, tcychan}@mie.utoronto.ca

Adam Diamant

Schulich School of Business, York University, Toronto, Canada, adiamant@schulich.yorku.ca

Rafid Mahmood

Mechanical & Industrial Engineering, University of Toronto, Toronto, Canada, rafid.mahmood@mail.utoronto.ca

We consider a data-driven framework for learning to generate decisions to instances of continuous optimization problems where the feasible set varies with an instance-specific auxiliary input. We use a data set of inputs and feasible solutions, as well as an oracle of feasibility, to iteratively train two machine learning models. The first model is a binary classifier for feasibility, which then serves as a barrier function to train the second model via an interior point method. We develop theory and optimality guarantees for interior point methods when given a barrier that relaxes the feasible set, and extend these results to obtain probabilistic out-of-sample guarantees for our learning framework. Finally, we implement our method on a radiation therapy treatment planning problem to predict personalized treatments for head-and-neck cancer patients.

# 1. Introduction

Consider a decision-maker who regularly solves different instances of a continuous optimization problem. There is a fixed objective and a common set of constraints over all instances. However, each instance also includes some exogenous input that may change the feasible set in a way that cannot be easily characterized. This input is not a parameter, but rather, auxiliary data that maps to latent constraints that may not have a convenient or formal mathematical representation.

While conventional optimization methods can produce optimal solutions for the common problem (i.e., the fixed objective and common constraints), these approaches make limited use of the auxiliary data because the latent constraints cannot be described conventionally. For example, a decision-maker may only be able to determine feasibility for any candidate decision via an oracle. The common optimization problem is a relaxation of the instance-specific formulation and if the latent constraints are not considered, the solution to the common problem may not be feasible for the problem instance. In practice, this approach still provides a starting point for the decisionmaker: she can manually modify the output and obtain an instance-specific solution. However, these post-hoc refinements suggest that the final solutions may not be provably optimal decisions.

Systematic approaches may attempt to predict the effect of the auxiliary data on the optimization model. Since these techniques are data-driven, they rely on the collection of past inputs and decisions as well as a prediction model to characterize the latent behavior. In operations research, this has been studied in the context of the auxiliary data affecting the objective. For example, local regression and neighbor or tree-based models can be placed in a conditional stochastic objective so as to penalize decisions predicted to be of poor quality (Bertsimas and Kallus 2019). An alternative is to recast the common problem as a parameterized one and use machine learning to predict the instance-specific parameters from the auxiliary input (Angalakudati et al. 2014, Ferreira et al. 2015, Babier et al. 2018a, Liu et al. 2018). This can be implemented by first predicting the quantities of interest and then optimizing over them or by embedding the machine learning model directly into the optimization problem and solving both simultaneously (Elmachtoub and Grigas 2017, Ban and Rudin 2018). Regardless of the approach, these techniques assume that the feasible region can be described mathematically and thus, the problem reduces to learning an objective function.

Since the common problem is well-specified and both past inputs and implemented decisions exist, it should be possible to learn a representation of the latent constraints in order to directly predict the instance-specific optimal decisions (Larsen et al. 2018). While the notion of using deep learning to predict optimal solutions to an optimization problem has been studied before (Hopfield and Tank 1985, Bengio et al. 2018), there are no results that guarantee the optimality of the predicted solutions for a general input. Thus, in this paper, we consider a framework for learning to generate optimal decisions to a continuous constrained optimization problem with a fixed objective and a feasible region characterized by the latent constraints of the decision-maker. By combining techniques from deep learning and operations research, we capture the best of both worlds: learning unstructured mappings between inputs and decisions while also proving mathematical guarantees on solution quality and generalization to out-of-sample problem instances.

# 1.1. Motivating application

Our work is motivated by the problem of automatically generating personalized radiation therapy (RT) treatment plans for patients diagnosed with cancer. RT is one of the primary methods for cancer treatment and is recommended for over 50% of all diagnoses (Delaney et al. 2005). In RT, a linear accelerator delivers beamlets of radiation from different angles to a tumor. To construct a treatment plan, a dosimetrist solves an optimization problem to sufficiently dose the tumor while minimizing damage to the healthy tissue. Before delivery, the plan must be approved by an oncologist based on its performance across several institutionally mandated criteria. Since it is impossible to simultaneously satisfy all criteria (e.g., tumor dose may be sacrificed to reduce dose to nearby critical structures or vice versa), oncologists make subjective trade-offs y choosing a subset of relevant criteria for a given patient based on prior expertise. These oncologist-driven trade-offs are effectively latent constraints that are parameterized by the patient's information and can be learned from examining past decisions (i.e., treatment plans) approved by the oncologist.

The current clinical practice is an iterative and time-consuming process where the dosimetrist and oncologist work together to generate acceptable treatment plans. The dosimetrist tunes the parameters of a surrogate multi-objective optimization problem and solves it to generate a candidate treatment. The oncologist is a membership oracle who determines whether a solution is satisfactory, and if not, suggests areas to improve. The dosimetrist then re-parametrizes the optimization problem and generates a new solution. Multiple iterations are typically required and it can take several days to generate a single treatment plan, especially for complex cases. With rising demand for RT treatment, particularly in developing countries, these individual delays can lead to a significant increase in wait times, which motivates automated approaches (Atun et al. 2015).

By analyzing past inputs and deliverable treatments, the oncologist's experience in determining satisfactory plans can be automated using predictive modeling. Several recent papers have demonstrated that automated treatment plan generation using machine learning techniques in conjunction with optimization is a viable alternative to the manual process (e.g., Wu et al. 2017, McIntosh and Purdie 2017, Babier et al. 2018a). This area is broadly known as knowledge-based planning (KBP). First, a machine learning model trained on previously delivered plans predicts an acceptable dose distribution (Shiraishi et al. 2015, McIntosh et al. 2017, Mahmood et al. 2018, Babier et al. 2019). The prediction is used as input into an optimization model that generates a treatment plan (i.e., the set of beamlets) that yields a similar dose distribution to the prediction.

Conventional prediction models are trained to generate dose distributions by minimizing some error with respect to a set of clinical plans rather than incorporating characteristics of the optimization problem. However, these approaches have two main drawbacks. First, predicted doses are not guaranteed to be feasible with respect to the latent oncologist constraints or even satisfy any optimality certificates. This may lead to delays in treatment delivery as the dosimetrist and oncologist may need additional iterations to manually correct the plans. Second, the protocols for radiation therapy treatment often vary between institutions (e.g., Geretschläger et al. 2015 versus Babier et al. 2018b). This makes it difficult to deploy the same automated planning pipeline at multiple institutions because off-the-shelf prediction models trained using data from one clinic may not satisfy protocols (e.g., hidden constraints) at other institutions (Wu et al. 2017). A clinic attempting to implement an automated planning pipeline would first need to train a custom prediction model using institution-specific data, which is especially difficult for smaller clinics or those ramping up in developing countries. Thus, it is essential to design dose prediction and automated treatment planning models that can deliver plans that are certifiably safe, as measured by feasibility and optimality guarantees, as well as adaptable to new settings.

# 1.2. Contribution

Our general approach involves transforming the "true" optimization problem (i.e., the model with the latent feasible set) into a prediction problem solved via two machine learning models. First, a binary classifier learns the feasible set for different instances. Then, a generative model navigates the instance-specific feasible region characterized by the support of the classifier to return an optimal solution. The two models are trained sequentially over several iterations; after each iteration, an oracle of feasibility labels the current predictions of the generative model as feasible or infeasible. The newly labelled data then helps train the classifier in the following iteration.

In our training algorithm, the generative model navigates the support of the classifier via an interior point method (IPM). However, IPMs are effective primarily when the feasible set is fully known. As we simultaneously learn feasibility along with optimality, our classifier does not enjoy the conventional properties of a canonical barrier function employed in IPMs. Therefore, we introduce the notion of a weak barrier function—one that may not perfectly discriminate feasibility. We derive a new  $\epsilon$ -optimality guarantee for optimization when given only a weak barrier function. We then show that our generative model, which now predicts solutions rather than optimizing for them, enjoys similar guarantees both for in-sample testing and out-of-sample problem instances. Our technical contributions are as follows:

- We introduce the concept of a δ-barrier, which is a barrier function for a relaxation of a feasible set. We characterize the (δ, ε)-optimality guarantee for solutions that are generated using a δ-barrier, generalizing key results of IPMs to the setting of a partially specified feasible set.
- 2. We present Interior Point Methods with Adversarial Networks (IPMAN), an iterative, oracleguided algorithm for learning to generate optimal solutions when given instance information to problems with latent constraints. The classifier trained in this algorithm approximates a  $\delta$ -barrier, the generative model generates ( $\delta, \epsilon$ )-optimal solutions on training instances, and the oracle-guided data augmentation guarantees that the classifier improves in every iteration.
- 3. We prove a generalization bound on the  $(\delta, \epsilon)$ -optimality gap of any model that predicts solutions to a random problem instance. This bound holds for IPMAN, meaning that both in-sample and out-of-sample error from the optimal value can be evaluated.

We apply IPMAN to predict the dose distribution to be delivered to head-and-neck cancer patients. In this context, we model the clinical criteria that must be met before a treatment is accepted as latent constraints to be learned from historically delivered treatment plans. After the classifier is trained, the generative model produces dose distributions that the classifier predicts will satisfy the relevant clinical criteria for that patient. The oracle labels the generated output as correct if the plan satisfied all of the hidden constraints that the oncologist had determined were relevant for the patient. By incorporating the evaluation of feasibility and optimality in training, our approach extends state-of-the-art generative adversarial network (GAN) frameworks for predicting dose distributions. Our final product is a generative model that outputs dose distributions that are guaranteed to be within a neighborhood of optimality (both in- and out-of-sample).

In our numerical experiments, we find that the doses predicted by our model better resemble clinical doses than current state-of-the-art baselines. We then show that once the latent constraints are learned, they can be altered using IPMAN so that dose distributions can be predicted for institutions with different protocols, without collecting a new institution-specific data set. This result has implications for the transfer of automated treatment planning technology between institutions (Wu et al. 2017), as well as closing the global gap in supply of radiation therapy by enabling all clinics to perform automated planning (Atun et al. 2015).

# 2. Background and related work

This paper brings together ideas from several fields. First, the concept of using two learning models, one to evaluate and one to generate solutions, is a standard approach in reinforcement learning (e.g., actor-critic methods (Konda and Tsitsiklis 2000)), and deep learning (e.g., generative adversarial networks (Goodfellow et al. 2014)). The specific practice of training a machine learning model using an oracle is known as imitation learning (Bain and Sammut 1999). Further, our loss function and optimality guarantees are derived using the theory of interior point methods (Nesterov and Nemirovskii 1994), while our learning guarantees extend recent results on Rademacher complexity for data-driven optimization (Bertsimas and Kallus 2019). Finally, the learning performed by our generative model bears a loose resemblance to estimation of distribution algorithms (EDAs), commonly used in evolutionary and black-box optimization (Pelikan et al. 2002). Because our work is most closely tied to interior point methods and the "learning to optimize" literature, we focus specifically on those two areas below.

## 2.1. Interior point methods

Interior point methods are among the most popular techniques for solving constrained optimization problems (Nesterov and Nemirovskii 1994). A constrained problem  $\min_{\mathbf{x}} \{f(\mathbf{x}) \mid \mathbf{x} \in \mathcal{X}\}$  is transformed into an unconstrained problem via a barrier function  $B(\mathbf{x})$ . The barrier satisfies two properties: (i)  $B(\mathbf{x}) = 0$  (i.e.,  $\log B(\mathbf{x}) = -\infty$ ) when  $\mathbf{x}$  is infeasible, and (ii)  $B(\mathbf{x}) > 0$  when  $\mathbf{x}$  is strictly feasible. The resulting problem is  $\min_{\mathbf{x}} \{f(\mathbf{x}) - \lambda \log B(\mathbf{x})\}$  where  $\lambda > 0$  is the dual parameter.

Given a differentiable barrier function and an initial solution  $\mathbf{x}^{(0)}$ , IPMs use the Newton method to iterate over a sequence  $\{(\lambda_j, \mathbf{x}^{(j)})\}_{j=0}^{\infty}$  until convergence to an optimal solution. These methods have found the most success in linear and quadratic optimization where it is possible to give theoretical guarantees on optimality as well as fast empirical convergence rates (Gondzio 2012). However, recent development of a new class of barrier functions known as entropic barriers have renewed interest in interior point methods for challenging (i.e., arbitrary) convex feasible sets (Bubeck and Eldan 2019). The most similar work to ours is by Badenbroek and de Klerk (2018) who solve a sampling-based IPM using a membership oracle for the feasible set and the aforementioned entropic barrier. IPMs have also been adapted for non-convex optimization problems (Vanderbei and Shanno 1999, Benson et al. 2004, Hinder and Ye 2018) although in this setting, efficiency guarantees generally do not exist.

The previous papers all assume access to either explicit constraints or at least a barrier function for the entire feasible set. In this work, we assume access to only a polyhedral relaxation of the true feasible set and learn feasibility via a classification model trained on labelled decisions. Furthermore, the previous papers focus on determining efficiency and complexity analyses of their algorithm for a single given instance of an optimization problem. Because we consider the problem of learning to optimize over a set of problem instances, we focus instead on in-sample and out-of-sample guarantees for the model rather than the efficiency of the IPM algorithm.

### 2.2. Learning to optimize

2.2.1. The operations research perspective. Learning to construct optimal solutions from auxiliary feature data is, in most cases, performed by embedding the output of a machine learning model as parameters in an optimization model (Angalakudati et al. 2014, Ferreira et al. 2015, Elmachtoub and Grigas 2017, Mišić 2019, Liu et al. 2018). This approach is particularly effective when there exists an important parameter in the problem and a clear relationship between it and the features (e.g., the demand in a revenue model (Ferreira et al. 2015)).

Non-parametric methods construct functions that map feature data to the optimal solution or the optimal value of the problem. This is often a stochastic optimization problem with a random objective conditioned on the input. Kao et al. (2009), Ban and Rudin (2018), and Bertsimas and Kallus (2019) consider Empirical Risk Minimization (ERM) where the objective is to learn a function that takes auxiliary data as input and outputs an optimal solution. However, it is often difficult to use advanced, expressive models while preserving computational tractability. Consequently, Hannah et al. (2010), Bertsimas and Kallus (2019), and Bertsimas and McCord (2018) consider a weighted learning framework, where the goal is to obtain a function that determines the weights (i.e., the conditional probability terms) using a sample-average approximation of the stochastic optimization problem.

Our work shares a similar approach to Ban and Rudin (2018) who use ERM to construct a predictor for the optimal solution to a newsvendor problem. As in their paper, we study out-ofsample generalization of the learning model. However, the key difference is that the newsvendor problem is only constrained by the non-negativity of the order quantities. In contrast, our focus on using ERM applies to a more general set of constraints. Our key contribution is to incorporate constraint satisfaction as the output of a binary classification model.

Bertsimas and Kallus (2019) remark on the challenges of constraint satisfaction when using the ERM approach and, consequently, focus on weighted learning. However, they also prove several generalizability results arising from ERM. Our paper further explores the avenue introduced by Bertsimas and Kallus (2019) by extending their generalization bound to our IPM framework.

2.2.2. The deep learning perspective. Recent advances in deep learning have prompted a resurgent interest in using neural networks to solve optimization problems (Bengio et al. 2018). Generally, a model is trained by minimizing a loss function that encourages the model to output optimal solutions to problem instances. While the literature mostly focuses on benchmark combinatorial problems such as the Traveling Salesman Problem (TSP) (Vinyals et al. 2015, Bello et al. 2017, Dai et al. 2017), there is an increasing interest in other operational applications (Donti et al. 2017, Larsen et al. 2018). Methodologically, Vinyals et al. (2015) and Larsen et al. (2018) use supervised learning, where a data set of problem instances and optimal solutions are used to train the model. On the other hand, Bello et al. (2017) and Dai et al. (2017) train a reinforcement learning agent to navigate the space of decisions. Finally, Donti et al. (2017) consider a gradient-descent algorithm that encourages predicting feasible and optimal solutions.

A major challenge in learning to predict optimal solutions is that it is difficult to enforce challenging constraints using a predictive model. The design of the neural network architecture sometimes naturally enforces certain structural constraints. For example, a pointer network is a recurrent neural network that returns permutations of a sequence making it ideal for satisfying tour constraints in a TSP (Vinyals et al. 2015). Alternatively, if the learning process is supervised, simply using a high-quality data set may be empirically sufficient (Larsen et al. 2018). A third approach is to customize the loss function to encourage constraint satisfaction (Donti et al. 2017). Regardless of the approach, a major benefit of learning to predict solutions over optimization is the speed at which solutions are produced. That is, after training has concluded, the model requires a simple function call (e.g., a neural network) to return a solution. Thus, the deep learning approaches produce heuristics that are significantly faster than conventional solvers (Bello et al. 2017, Larsen et al. 2018). The drawback is that they do not admit formal optimality guarantees.

We preserve the efficiency of learning algorithms while addressing the problem of constraint satisfaction. More specifically, we consider constraints that are not formally stated, but rather, are implicitly provided via data and an oracle. Further, we provide optimality guarantees on the solutions generated by the deep learning model as well as characterizing the out-of-sample error.

# 3. Constrained optimization with a partially specified feasible set

We define the problem of solving a constrained optimization problem when the feasible set is specified only by common constraints and auxiliary data. We introduce notation in Section 3.1, before describing the problem and assumptions in Section 3.2. In Section 3.3, we propose a new barrier problem and prove both optimality and feasibility guarantees for our generalization.

# 3.1. Notation

We denote vectors by bold and sets by calligraphic script. The interior, boundary, and closure of a set are denoted  $\operatorname{int}(\mathcal{X})$ ,  $\operatorname{bd}(\mathcal{X})$ , and  $\operatorname{cl}(\mathcal{X})$  respectively. The exclusion of  $\mathcal{X}_1$  from  $\mathcal{X}_2 \supseteq \mathcal{X}_1$  is denoted as  $\mathcal{X}_2 \setminus \mathcal{X}_1$ . We denote probability distributions using  $\mathbb{P}$ . The support of a probability distribution is denoted  $\operatorname{supp}(\mathbb{P})$ . Samples from a random variable  $\mathbf{x} \sim \mathbb{P}$  are accented  $\hat{\mathbf{x}} \sim \mathbb{P}$ .  $\|\cdot\|$  refers to the  $l_2$  norm unless specified otherwise. A function  $f(\mathbf{x})$  is *L*-Lipschitz continuous in  $\mathbf{x}$  if there exists a constant L > 0 such that  $|f(\mathbf{x}_1) - f(\mathbf{x}_2)| \leq L \|\mathbf{x}_1 - \mathbf{x}_2\|$  for all  $\mathbf{x}_1, \mathbf{x}_2$ .

## 3.2. Problem setup

Let  $\mathbf{u} \in \mathcal{U}$  denote auxiliary inputs that describe our optimization problem and  $\mathbb{P}_{\mathbf{u}}$  the probability distribution of inputs. Let  $\mathbf{x} \in \mathbb{R}^n$  denote decisions and for a given  $\mathbf{u}$ , consider the problem

$$\mathbf{OP}(\mathbf{u}): \quad \min_{\mathbf{x}} \left\{ f(\mathbf{x}) \mid \mathbf{x} \in \mathcal{X}(\mathbf{u}) \right\}$$

where  $f(\mathbf{x})$  is an objective function and  $\mathcal{X}(\mathbf{u})$  is a feasible set determined by  $\mathbf{u}$ . We assume  $f(\mathbf{x}) = \mathbf{f}^{\mathsf{T}}\mathbf{x}$  is linear without loss of generality; that is, our results extend to convex objectives with minor modifications. We also make the following assumptions about  $\mathcal{X}(\mathbf{u})$ :

- 1. The space of inputs  $\mathcal{U}$  is compact, continuous, and has a non-empty interior.
- 2. For all  $\mathbf{u} \in \mathcal{U}$ , the feasible set  $\mathcal{X}(\mathbf{u})$  is compact, continuous, and has a non-empty interior. Furthermore, the joint set  $\{(\mathbf{x}, \mathbf{u}) \mid \mathbf{x} \in \mathcal{X}(\mathbf{u}), \mathbf{u} \in \mathcal{U}\}$  is compact.
- 3. We know a full-dimensional polyhedral relaxation  $\mathcal{P} = \{\mathbf{x} \mid \mathbf{a}_m^\mathsf{T} \mathbf{x} \leq b_m, m = 1, \dots, M\}$  such that  $\mathcal{X}(\mathbf{u}) \subset \operatorname{int}(\mathcal{P})$  for all  $\mathbf{u}$ .
- 4. Although we do not know  $\mathcal{X}(\mathbf{u})$  a priori, we have access to:
  - (a) A data set of feasible decision and input pairs  $\mathcal{D} = \{(\hat{\mathbf{x}}_i, \hat{\mathbf{u}}_i)\}_{i=1}^{N_{\mathbf{x}}}$ , where  $\hat{\mathbf{x}}_i \in \mathcal{X}(\hat{\mathbf{u}}_i)$  for all  $i \in \{1, \dots, N_{\mathbf{x}}\}$ . In general, the data set may include multiple feasible decisions per input. This data is sampled i.i.d. from a distribution  $\mathbb{P}_{(\mathbf{x},\mathbf{u})}$ .
  - (b) A feasibility oracle  $\Psi(\mathbf{x}, \mathbf{u})$  where  $\Psi(\mathbf{x}, \mathbf{u}) = 1$  if  $\mathbf{x} \in \mathcal{X}(\mathbf{u})$  and  $\Psi(\mathbf{x}, \mathbf{u}) = 0$  otherwise.

The first assumption ensures that the auxiliary inputs belong to a well-behaved set (i.e., compact, continuous, and non-empty). The second and third assumptions ensure that the instance-specific feasible sets have similar requirements. The set  $\mathcal{P}$  represents the space defined by the common constraints that all instances must satisfy in addition to the instance-specific constraints. Finally,

the fourth assumption describes the available data to learn to construct an optimal solution to  $OP(\mathbf{u})$ . In practice, the data set of feasible solutions may represent previously implemented decisions which are required to learn an approximation of the feasible sets  $\mathcal{X}(\mathbf{u})$ . The oracle  $\Psi(\mathbf{x}, \mathbf{u})$  helps guide the search for an optimal solution by ensuring that any constructed solution is feasible.

# **3.3.** Optimization with a $\delta$ -barrier

Consider an instance  $\mathbf{OP}(\mathbf{u})$  for a fixed  $\mathbf{u}$ . Although  $\mathcal{X}(\mathbf{u})$  is unspecified, we know a relaxation  $\mathcal{P} \supset \mathcal{X}(\mathbf{u})$ . Were the relaxation tight (i.e.,  $\mathcal{P} = \mathcal{X}(\mathbf{u})$ ), then we could consider a canonical barrier function (i.e., defining  $B(\mathbf{x})$  such that  $\log B(\mathbf{x}) = \sum_{m=1}^{M} \log (b_m - \mathbf{a}_m^{\mathsf{T}} \mathbf{x})$ ) and use an IPM to obtain an optimal solution  $\mathbf{x}^*(\mathbf{u})$  to  $\mathbf{OP}(\mathbf{u})$  (Nesterov and Nemirovskii 1994). However, if  $\mathcal{X}(\mathbf{u}) \subset \mathcal{P}$ , then the barrier incorrectly returns  $B(\mathbf{x}) > 0$  for all  $\mathbf{x} \in \mathcal{P} \setminus \mathcal{X}(\mathbf{u})$ .

For this problem, the canonical barrier belongs to a class of functions that are barriers over a relaxation or super-set of  $\mathcal{X}(\mathbf{u})$ . Consider functions  $B_{\delta}(\mathbf{x}, \mathbf{u})$  such that  $B_{\delta}(\mathbf{x}, \mathbf{u}) > 0$  for all  $\mathbf{x} \in \mathcal{X}(\mathbf{u})$ , and  $B_{\delta}(\mathbf{x}, \mathbf{u}) = 0$  for all  $\mathbf{x}$  that are sufficiently far from  $\mathcal{X}(\mathbf{u})$ . We define these functions as  $\delta$ -barriers.

DEFINITION 1. For some  $\delta > 0$ , let  $\mathcal{N}_{\delta}(\mathcal{X}(\mathbf{u})) = \{\mathbf{x} + \boldsymbol{\epsilon} \mid \mathbf{x} \in \mathcal{X}(\mathbf{u}), \|\boldsymbol{\epsilon}\| < \delta\}$  be a  $\delta$ -neighbourhood around  $\mathcal{X}(\mathbf{u})$ . A  $\delta$ -barrier  $B_{\delta}(\mathbf{x}, \mathbf{u}) : \mathbb{R}^n \times \mathcal{U} \to [0, 1)$  is a function that satisfies

$$\mathcal{X}(\mathbf{u}) \subset \{\mathbf{x} \mid B_{\delta}(\mathbf{x}, \mathbf{u}) > 0\} \subseteq \mathcal{N}_{\delta}\left(\mathcal{X}(\mathbf{u})\right)$$

Note that for a given barrier function  $B(\mathbf{x}, \mathbf{u})$  supported over a super-set of  $\mathcal{X}(\mathbf{u})$ ,  $\delta$  is equivalent to the Hausdorff distance between  $\mathcal{X}(\mathbf{u})$  and the support of the function, i.e.,

$$\delta = d_H \left( \mathcal{X}(\mathbf{u}), \{ \mathbf{x} \mid B(\mathbf{x}, \mathbf{u}) > 0 \} \right) = \min_{\xi \ge 0} \left\{ \xi \mid \{ \mathbf{x} \mid B(\mathbf{x}, \mathbf{u}) > 0 \} \subseteq \mathcal{N}_{\xi} \left( \mathcal{X}(\mathbf{u}) \right) \right\}.$$
(1)

REMARK 1. Let  $\Delta(\mathbf{u}) = d_H(\mathcal{X}(\mathbf{u}), \mathcal{P})$ . A canonical barrier for  $\mathcal{P}$  is a  $\Delta(\mathbf{u})$ -barrier for  $\mathcal{X}(\mathbf{u})$ . As  $\mathcal{P}$  is known, we always assume  $\delta \leq \Delta(\mathbf{u})$ .

Given a  $\delta$ -barrier  $B_{\delta}(\mathbf{x}, \mathbf{u})$ , let  $\lambda > 0$  be a constant corresponding to the Lagrangian dual variable. We then define the unconstrained barrier optimization problem

$$\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda): \min \left\{ f(\mathbf{x}) - \lambda \log B_{\delta}(\mathbf{x}, \mathbf{u}) \right\}$$
(2)

REMARK 2. We assume without loss of generality of using a single  $\delta$ -barrier to determine feasibility. Canonical barriers in IPMs generally sum multiple barriers each addressing a single constraint (e.g.,  $\sum_{m=1}^{M} \log (b_m - \mathbf{a}_m^{\mathsf{T}} \mathbf{x})$  for  $\mathcal{P}$ ). Our results also generalize to multiple  $\delta$ -barriers in one problem. The optimal value of  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda)$  is bounded by the optimal value of  $\mathbf{OP}(\mathbf{u})$ .

THEOREM 1. Let  $\mathbf{x}^*(\mathbf{u})$  be an optimal solution to  $\mathbf{OP}(\mathbf{u})$ . For any  $\lambda > 0$ ,  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda)$  is bounded and feasible. An optimal solution  $\mathbf{x}^{\lambda}(\mathbf{u})$  to  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda)$  is  $(\delta, \epsilon)$ -optimal for  $\mathbf{OP}(\mathbf{u})$ :

$$f(\mathbf{x}^{\lambda}(\mathbf{u})) - \epsilon < f(\mathbf{x}^{*}(\mathbf{u})) < f(\mathbf{x}^{\lambda}(\mathbf{u})) + \delta L,$$
(3)

where L is the Lipschitz constant of  $f(\mathbf{x})$  and  $\epsilon = C\lambda$ , where C is a positive constant.

Proof of Theorem 1. We first prove that  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda)$  is bounded and feasible. Note that for any  $\lambda > 0$ , the second term,  $-\lambda \log B_{\delta}(\mathbf{x}, \mathbf{u})$ , is also bounded below since  $B_{\delta}(\mathbf{x}, \mathbf{u}) \in [0, 1)$ . Moreover,  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda)$  is only feasible within  $\{\mathbf{x} | B_{\delta}(\mathbf{x}, \mathbf{u}) > 0\} \subset \mathcal{P}$ . By assumption,  $\mathcal{P}$  is closed and bounded, meaning  $f(\mathbf{x})$ , by assumption of linearity, has a bounded minimum within  $\mathcal{P}$ . Therefore,  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda)$  must have a bounded optimal solution. To show feasibility, note that by the definition of a  $\delta$ -barrier, any solution that is feasible for  $\mathbf{OP}(\mathbf{u})$  is also feasible for  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda)$ .

Suppose we choose  $\epsilon = -\lambda \log B_{\delta}(\mathbf{x}^*(\mathbf{u}), \mathbf{u}))$ . By definition,  $0 < B_{\delta}(\mathbf{x}^*(\mathbf{u}), \mathbf{u}) < 1$ , implying that  $C := -\log B_{\delta}(\mathbf{x}^*(\mathbf{u}), \mathbf{u}) > 0$  is valid. Let  $\mathbf{x}^{\lambda}$  be an optimal solution to  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda)$ . We first prove that  $f(\mathbf{x}^{\lambda}(\mathbf{u})) - \epsilon < f(\mathbf{x}^*(\mathbf{u}))$ :

$$f(\mathbf{x}^{*}(\mathbf{u})) + \epsilon = f(\mathbf{x}^{*}(\mathbf{u})) - \lambda \log B_{\delta}(\mathbf{x}^{*}(\mathbf{u}), \mathbf{u})$$
$$\geq f(\mathbf{x}^{\lambda}(\mathbf{u})) - \lambda \log B_{\delta}(\mathbf{x}^{\lambda}(\mathbf{u}), \mathbf{u})$$
$$> f(\mathbf{x}^{\lambda}(\mathbf{u})).$$

The first inequality follows from the optimality of  $\mathbf{x}^{\lambda}(\mathbf{u})$  for  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda)$  while the second inequality follows from  $\log B_{\delta}(\mathbf{x}^{\lambda}(\mathbf{u}), \mathbf{u}) < 0$ , meaning that  $-\lambda \log B_{\delta}(\mathbf{x}^{\lambda}(\mathbf{u}), \mathbf{u}) > 0$  and can be removed. Moving  $\epsilon$  to the right-hand-side gives the lower bound.

The proof for  $f(\mathbf{x}^*(\mathbf{u})) < f(\mathbf{x}^{\lambda}(\mathbf{u})) + \delta L$  has two cases. If  $\mathbf{x}^{\lambda}(\mathbf{u}) \in \mathcal{X}(\mathbf{u})$ , then by the optimality of  $\mathbf{x}^*(\mathbf{u})$  for  $\mathbf{OP}(\mathbf{u})$ , we have  $f(\mathbf{x}^*(\mathbf{u})) \leq f(\mathbf{x}^{\lambda}(\mathbf{u})) < f(\mathbf{x}^{\lambda}(\mathbf{u})) + \delta L$ . Otherwise if  $\mathbf{x}^{\lambda}(\mathbf{u}) \in \mathcal{N}_{\delta}(\mathcal{X}(\mathbf{u})) \setminus \mathcal{X}(\mathbf{u})$ , then let  $\tilde{\mathbf{x}} \in \arg\min_{\mathbf{x} \in \mathcal{X}(\mathbf{u})} \|\mathbf{x}^{\lambda}(\mathbf{u}) - \mathbf{x}\|$  be the projection of  $\mathbf{x}^{\lambda}(\mathbf{u})$  on  $\mathcal{X}(\mathbf{u})$ . Then,

$$\begin{aligned} f\left(\mathbf{x}^{*}(\mathbf{u})\right) - f\left(\mathbf{x}^{\lambda}(\mathbf{u})\right) &\leq f(\tilde{\mathbf{x}}) - f\left(\mathbf{x}^{\lambda}(\mathbf{u})\right) \\ &\leq \left|f(\tilde{\mathbf{x}}) - f\left(\mathbf{x}^{\lambda}(\mathbf{u})\right)\right| \\ &\leq \left\|\tilde{\mathbf{x}}(\mathbf{u}) - \mathbf{x}^{\lambda}(\mathbf{u})\right\| L \\ &\leq \delta L. \end{aligned}$$

The first inequality follows from the optimality of  $\mathbf{x}^*(\mathbf{u})$  over  $\tilde{\mathbf{x}}$  for  $\mathbf{OP}(\mathbf{u})$ . The third inequality follows from the Lipschitz continuity of  $f(\mathbf{x})$  and the fourth by definition of the  $\delta$ -barrier. Therefore, the upper bound in inequality (3) is proved for both cases.

The  $(\delta, \epsilon)$ -optimality inequality proved in Theorem 1 generalizes the classical  $\epsilon$ -optimality bound of IPMs (Nesterov and Nemirovskii 1994). That is, when  $\delta = 0$ , we obtain  $f(\mathbf{x}^{\lambda}(\mathbf{u})) - \epsilon < f(\mathbf{x}^{*}(\mathbf{u})) < f(\mathbf{x}^{\lambda}(\mathbf{u}))$ . Furthermore, similar to classical IPMs, the  $(\delta, \epsilon)$ -optimality of solutions to  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda)$ can be controlled by tuning  $\lambda$ . Specifically because  $\epsilon = C\lambda$  for a fixed C, as  $\lambda$  goes to 0, so does  $\epsilon$ .

In a classical IPM, the barrier problem is repeatedly solved over a sequence of decreasing  $\lambda$  (Boyd and Vandenberghe 2004). In that setting, a large  $\lambda$  guarantees that the barrier problem yields solutions in the interior of the feasible set, while small  $\lambda$  guarantees solutions close to the true



Figure 1 The bold shape is  $\mathcal{P}$  and the filled region is  $\mathcal{X}(\mathbf{u})$ . We consider a canonical barrier for  $\mathcal{P}$ , where the dotted lines are contours for the barrier. An optimal solution to  $OP(\mathbf{u})$  is  $\mathbf{x}^*(\mathbf{u})$ .

optimum. Here however,  $\epsilon$  may decrease but  $\delta$  is a property of the barrier itself. Therefore, we can only guarantee that for large  $\lambda$ , an optimal solution  $\mathbf{x}^{\lambda}(\mathbf{u})$  is sub-optimal for  $\mathbf{OP}(\mathbf{u})$  but as  $\lambda$  decreases, the optimal solution may have lower objective function value, i.e., become infeasible. Figure 1 shows a sample sequence of decreasing  $\lambda$  and corresponding solutions  $\mathbf{x}^{\lambda}(\mathbf{u})$ . In the Electronic Companion EC.1, we show that this behavior holds more generally and formally prove that the  $\delta$ -barrier IPM inherits many of the same properties as the classical IPM.

# 4. Interior Point Methods with Adversarial Networks

We now develop a learning-based approach to solving instances of  $OP(\mathbf{u})$ . We first introduce our main algorithm that iteratively trains two machine learning models. The algorithm consists of (i) a classifier that learns to distinguish feasibility for  $OP(\mathbf{u})$ , and thus, approximates a  $\delta$ -barrier for any  $\mathbf{u}$ ; and (ii) a generative model that uses the classifier to learn to predict  $(\delta, \epsilon)$ -optimal solutions to  $OP(\mathbf{u})$  for a given  $\mathbf{u}$ .

We require a data set of feasible solutions  $\mathcal{D}$  as well as an oracle  $\Psi(\mathbf{x}, \mathbf{u})$  as introduced in Section 3.2. Furthermore, let  $\hat{\mathcal{U}} = \{\hat{\mathbf{u}}_i\}_{i=1}^{N_{\mathbf{u}}} = \{\hat{\mathbf{u}} \mid \exists \hat{\mathbf{x}} : (\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \mathcal{D}\}$  denote a data set of inputs. In general,  $N_{\mathbf{u}} \leq N_{\mathbf{x}}$  as this data set is obtained by collecting the unique auxiliary inputs from  $\mathcal{D}$ . We also assume an additional data set of *infeasible solutions*,  $\bar{\mathcal{D}} = \{(\hat{\mathbf{x}}_{\bar{i}}, \hat{\mathbf{u}}_{\bar{i}})\}_{\bar{i}=1}^{\bar{N}_{\mathbf{x}}}$  where  $\hat{\mathbf{x}}_{\bar{i}} \in \mathbb{R}^n \setminus \mathcal{X}(\hat{\mathbf{u}}_{\bar{i}})$ . This data set arrives from a distribution  $\bar{\mathbb{P}}_{(\mathbf{x},\mathbf{u})}$ , similar to  $\mathcal{D} \sim \mathbb{P}_{(\mathbf{x},\mathbf{u})}$ . Unlike  $\mathcal{D}$  however,  $\bar{\mathcal{D}}$  is usually not available a priori. Instead as we show in our numerical results,  $\bar{\mathcal{D}}$  is generated by sampling.

# 4.1. Overview of the main algorithm

Let  $\mathcal{F} = \{F : \mathcal{U} \to \mathbb{R}^n\}$  denote a class of generative models. Let  $\mathcal{B} = \{B : \mathbb{R}^n \times \mathcal{U} \to [0, 1]\}$  denote a class of binary classifiers. In each iteration, we first train  $B(\mathbf{x}, \mathbf{u}) \in \mathcal{B}$  to correctly label points in  $\mathcal{D}$ 

and  $\overline{\mathcal{D}}$ . We then re-train  $F(\mathbf{u}) \in \mathcal{F}$  over  $\hat{\mathcal{U}}$  via an IPM that uses the binary classifier as the barrier function. After each step of the IPM, the oracle  $\Psi(\mathbf{x}, \mathbf{u})$  labels the solutions generated by  $F(\mathbf{u})$  as either feasible or infeasible. At the end of the iteration, we add these points to  $\mathcal{D}$  and  $\overline{\mathcal{D}}$ . In this way, the training data for the classifier is augmented so that we iteratively learn a  $\delta$ -barrier that more closely approximates  $\mathcal{X}(\mathbf{u})$  for any  $\mathbf{u}$ . Let k index the iterations, starting from 1. The k-th iteration proceeds as follows:

1. Train the classifier  $B \in \mathcal{B}$  to predict  $(\hat{\mathbf{x}}_i, \hat{\mathbf{u}}_i) \in \mathcal{D}^{(k)}$  as feasible and  $(\hat{\mathbf{x}}_{\bar{i}}, \hat{\mathbf{u}}_{\bar{i}}) \in \bar{\mathcal{D}}^{(k)}$  as infeasible solution-input pairs by solving the Feasibility Classification Problem:

$$\mathbf{FCP}(\mathcal{D}^{(k)}, \bar{\mathcal{D}}^{(k)}) : \sup_{B \in \mathcal{B}} \left\{ \frac{1}{N_{\mathbf{x}}} \sum_{i=1}^{N_{\mathbf{x}}} \log B(\hat{\mathbf{x}}_{i}, \hat{\mathbf{u}}_{i}) + \frac{1}{\bar{N}_{\mathbf{x}}} \sum_{\bar{i}=1}^{\bar{N}_{\mathbf{x}}} \log \left(1 - B(\hat{\mathbf{x}}_{\bar{i}}, \hat{\mathbf{u}}_{\bar{i}})\right) \right\}.$$
(4)

The objective function is known as the cross-entropy loss function in machine learning (Goodfellow et al. 2016). Let  $B^{(k)}$  be an optimal solution to  $\mathbf{FCP}(\mathcal{D}^{(k)}, \bar{\mathcal{D}}^{(k)})$ . The optimal value is 0 and achieved when  $B^{(k)}$  satisfies  $B^{(k)}(\hat{\mathbf{x}}_i, \hat{\mathbf{u}}_i) = 1$  and  $B^{(k)}(\hat{\mathbf{x}}_{\bar{i}}, \hat{\mathbf{u}}_{\bar{i}}) = 0$ .

2. Fix an initial dual parameter  $\lambda_0 > 0$  and a decay rate  $0 < \nu < 1$ . Let M > 0 denote the number of IPM steps and, for  $j \in \{0, ..., M\}$ , let  $\lambda_j = \lambda_0 \nu^j$  denote the dual parameter. We train the generative model to predict optimal solutions to  $\mathbf{BP}(\hat{\mathbf{u}}_i, B^{(k)}, \lambda_j)$  for all  $\hat{\mathbf{u}}_i \in \hat{\mathcal{U}}$ . This task is referred to as the Generative Barrier Problem:

$$\mathbf{GBP}(\hat{\mathcal{U}}, B^{(k)}, \lambda_j) : \min_{F \in \mathcal{F}} \left\{ \frac{1}{N_{\mathbf{u}}} \sum_{i=1}^{N_{\mathbf{u}}} f(F(\hat{\mathbf{u}}_i)) - \lambda_j \log B^{(k)}(F(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i) \right\}.$$
(5)

Let  $F^{(j,k)}$  be an optimal solution to  $\mathbf{GBP}(\hat{\mathcal{U}}, B^{(k)}, \lambda_j)$ . Whereas  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda)$  directly optimizes for a single  $\mathbf{u}$ ,  $\mathbf{GBP}(\hat{\mathcal{U}}, B^{(k)}, \lambda_j)$  is an empirical risk minimization problem that trains  $F^{(j,k)}(\mathbf{u})$  to predict  $\mathbf{x}^{\lambda_j}(\mathbf{u})$ . Furthermore, we now use the classifier  $B^{(k)}(\mathbf{x}, \mathbf{u})$  as the  $\delta$ -barrier.

3. For each pair  $(F^{(j,k)}(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i)$ , use the oracle  $\Psi(\mathbf{x}, \mathbf{u})$  to validate whether the generative model outputs a feasible or infeasible solution. Append the predicted solution to  $\mathcal{D}^{(k)}$  or  $\bar{\mathcal{D}}^{(k)}$ :

$$\mathcal{D}^{(k+1)} = \mathcal{D}^{(k)} \cup \mathcal{Q}, \text{ where } \mathcal{Q} := \left\{ \left( F^{(j,k)}(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i \right) \middle| \Psi(F^{(j,k)}(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i) = 1, \hat{\mathbf{u}}_i \in \hat{\mathcal{U}} \right\}$$
(6)

$$\bar{\mathcal{D}}^{(k+1)} = \bar{\mathcal{D}}^{(k)} \cup \bar{\mathcal{Q}}, \text{ where } \bar{\mathcal{Q}} := \left\{ \left( F^{(j,k)}(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i \right) \middle| \Psi(F^{(j,k)}(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i) = 0, \hat{\mathbf{u}}_i \in \hat{\mathcal{U}} \right\}$$
(7)

Note that for all  $\hat{\mathbf{u}}_i \in \hat{\mathcal{U}}$ , the generative model will always produce solutions that satisfy  $B^{(k)}(F^{(j,k)}(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i) > 0$ . The oracle then checks whether  $B^{(k)}(\mathbf{x}, \mathbf{u})$  is correct for each point. Then, the classifier can correct itself in the k + 1-th iteration.

Figure 2 shows the steps and the outcome for a single iteration k and a single  $\hat{\mathbf{u}}_i$ . In the remainder of this section, we discuss several desirable properties of this algorithm. We first show that the classifier learns to approximate a  $\delta$ -barrier. We then show that the generative model satisfies a  $(\delta, \epsilon)$ optimality guarantee on in-sample instances, albeit not as strong as one that would be obtained



(a)  $B^{(k)}$  learns to classify points in (b)  $F^{(j,k)}$  is trained to predict a (c) After binning, we obtain a tighter  $\mathcal{D}^{(k)}$  and  $\bar{\mathcal{D}}^{(k)}$ . sequence of solutions given  $\lambda_j$ . barrier in the next iteration.

Figure 2 One iteration of IPMAN for a single  $\hat{\mathbf{u}}_i$ .  $\Diamond$  and  $\Box$  represent  $\mathcal{D}$  and  $\overline{\mathcal{D}}$ , respectively. The filled region is  $\mathcal{X}(\mathbf{u})$  and the bold line on the outside is  $\mathcal{P}$ . The inner solid lines show the support of  $B^{(k)}(\mathbf{x}, \mathbf{u})$ .

were we to directly solve  $\mathbf{BP}(\hat{\mathbf{u}}_i, B_{\delta}, \lambda)$ . Finally, we show that the data augmentation procedure shrinks the set of optimal solutions to  $\mathbf{FCP}(\mathcal{D}, \bar{\mathcal{D}})$ . Thus, each iteration learns a  $\delta$ -barrier that more closely approximates  $\mathcal{X}(\mathbf{u})$  for any  $\mathbf{u}$ .

#### 4.2. A data-driven $\delta$ -barrier

A perfect barrier function where  $\delta = 0$  is a perfect classifier of feasibility. That is, it returns positive values if and only if the input **x** is feasible for **OP**(**u**), and zero otherwise. A  $\delta$ -barrier can perfectly classify feasible points, but potentially incorrectly classify infeasible points. By solving this classification problem,  $B(\mathbf{x}, \mathbf{u})$  learns to approximate a  $\delta$ -barrier.

We assume that the classifier is sufficiently parameterized to be able to describe a complex, and potentially non-convex feasible set. A sufficient condition would be that the model class  $\mathcal{B}$  satisfies a Universal Approximation property.

ASSUMPTION 1.  $\mathcal{B}$  satisfies a Universal Approximation Theorem (Hornik 1991). That is, for any continuous function  $B^*(\mathbf{x}, \mathbf{u}) : \mathbb{R}^n \times \mathcal{U} \to [0, 1]$  and degree of accuracy  $\varepsilon > 0$ , there exists  $B \in \mathcal{B}$ such that  $|B(\mathbf{x}, \mathbf{u}) - B^*(\mathbf{x}, \mathbf{u})| < \varepsilon$  for all  $\mathbf{x} \in \mathbb{R}^n$ ,  $\mathbf{u} \in \mathcal{U}$ .

LEMMA 1 (Arjovsky and Bottou (2017)). Let  $\mathcal{B}$  satisfy Assumption 1. If  $\mathcal{D}^{(k)}$  and  $\overline{\mathcal{D}}^{(k)}$  are closed, then  $\mathbf{FCP}(\mathcal{D}^{(k)}, \overline{\mathcal{D}}^{(k)})$  is feasible and has an optimal value equal to 0.

REMARK 3. Arjovsky and Bottou (2017) prove Lemma 1 by showing the cross-entropy loss function yields an optimal value of 0 whenever the two classes to be predicted are supported over compact and disjoint sets. We observe that closedness, rather than compactness, is sufficient.

Proof of Lemma 1. Because  $B(\mathbf{x}, \mathbf{u}) \in [0, 1]$ , the optimal value must be 0 and is attained only when  $B^{(k)}(\mathbf{x}, \mathbf{u}) = 1$  for all  $(\mathbf{x}, \mathbf{u}) \in \mathcal{D}^{(k)}$  and  $B^{(k)}(\mathbf{x}, \mathbf{u}) = 0$  for all  $(\mathbf{x}, \mathbf{u}) \in \overline{\mathcal{D}}^{(k)}$ . Then, Urysohn's Smooth Lemma states that given two closed and disjoint sets  $\mathcal{A}$  and  $\mathcal{A}'$ , there exists a continuous function for which  $f(\mathcal{A}) = 1$  and  $f(\mathcal{A}') = 0$  (Engelking 1977). Now note that for any iteration,

$$\mathcal{D}^{(k)} \subset \left\{ (\mathbf{x}, \mathbf{u}) \mid \mathbf{x} \in \mathcal{X}(\mathbf{u}), \mathbf{u} \in \mathcal{U} \right\}$$
  
 $\bar{\mathcal{D}}^{(k)} \subset \left\{ (\mathbf{x}, \mathbf{u}) \mid \mathbf{x} \in \mathbb{R}^n \setminus \mathcal{X}(\mathbf{u}), \mathbf{u} \in \mathcal{U} \right\},$ 

meaning that they are always disjoint and, by assumption, closed. Therefore, there exists a continuous function that achieves the optimal value to  $\mathbf{FCP}(\mathcal{D}^{(k)}, \bar{\mathcal{D}}^{(k)})$ . Thus, by Assumption 1, there exists  $B(\mathbf{x}, \mathbf{u}) \in \mathcal{B}$  that can approximate the supremum.

Intuitively for  $\mathbf{FCP}(\mathcal{D}^{(k)}, \bar{\mathcal{D}}^{(k)})$ , we require that the two data sets,  $\mathcal{D}^{(k)}$  and  $\bar{\mathcal{D}}^{(k)}$ , be disjoint and sufficiently far from each other to ensure that a classifier can learn a separation between their supports. Our problem naturally provides the disjoint property as the two data sets arise from solutions that are correctly labeled as feasible and infeasible, respectively. Figure 2(a) demonstrates an example of this intuition. We also observe that the optimal solution set to  $\mathbf{FCP}(\mathcal{D}^{(k)}, \bar{\mathcal{D}}^{(k)})$  is large and any function that separates the two sets is optimal.

Given a limited data set, solving  $\mathbf{FCP}(\mathcal{D}^{(k)}, \bar{\mathcal{D}}^{(k)})$  may not yield a *practically useful* classifier. For example, an optimal classifier may over-fit to the data, be unnecessarily complex, or may misclassify regions where data is not available. In order to ensure that the classifier is a  $\delta$ -barrier, we assume sufficient data so as to be able to solve the stochastic optimization variant of  $\mathbf{FCP}(\mathcal{D}, \bar{\mathcal{D}})$ :

$$\mathbf{FCP}(\mathbb{P}_{(\mathbf{x},\mathbf{u})},\bar{\mathbb{P}}_{(\mathbf{x},\mathbf{u})}): \quad \sup_{B\in\mathcal{B}} \left\{ \mathbb{E}_{\mathbf{x},\mathbf{u}\sim\mathbb{P}_{(\mathbf{x},\mathbf{u})}} \left[\log B(\mathbf{x},\mathbf{u})\right] + \mathbb{E}_{\mathbf{x},\mathbf{u}\sim\bar{\mathbb{P}}_{(\mathbf{x},\mathbf{u})}} \left[\log \left(1 - B(\mathbf{x},\mathbf{u})\right)\right] \right\}.$$
(8)

REMARK 4. The key difference between  $\mathbf{FCP}(\mathcal{D}^{(k)}, \bar{\mathcal{D}}^{(k)})$  and  $\mathbf{FCP}(\mathbb{P}_{(\mathbf{x},\mathbf{u})}, \bar{\mathbb{P}}_{(\mathbf{x},\mathbf{u})})$  is that the latter permits arbitrary probability distributions rather than discrete empirical distributions. An optimal solution to  $\mathbf{FCP}(\mathbb{P}_{(\mathbf{x},\mathbf{u})}, \bar{\mathbb{P}}_{(\mathbf{x},\mathbf{u})})$  satisfies  $B^{(k)}(\mathbf{x},\mathbf{u}) = 1$  for all  $(\mathbf{x},\mathbf{u}) \in \mathrm{supp}(\mathbb{P}_{(\mathbf{x},\mathbf{u})})$  and  $B^{(k)}(\mathbf{x},\mathbf{u}) = 0$  for all  $(\mathbf{x},\mathbf{u}) \in \mathrm{supp}(\bar{\mathbb{P}}_{(\mathbf{x},\mathbf{u})})$ , so long as they are closed and disjoint. Further, because  $\bar{\mathcal{D}}^{(k)}$  can be generated by sampling, we always assume access to a distribution  $\bar{\mathbb{P}}_{(\mathbf{x},\mathbf{u})}$  for which  $\mathrm{supp}(\bar{\mathbb{P}}_{(\mathbf{x},\mathbf{u})}) \supseteq \{(\mathbf{x},\mathbf{u}) \mid \mathbf{x} \in \mathbb{R}^n \setminus \mathcal{P}, \mathbf{u} \in \mathcal{U}\}$ . The remaining difference in studying the stochastic versus the data-driven classification problems is that the data set of feasible solutions  $\mathcal{D}^{(k)}$  is sufficiently large.

If the supports of  $\mathbb{P}_{(\mathbf{x},\mathbf{u})}$  and  $\overline{\mathbb{P}}_{(\mathbf{x},\mathbf{u})}$  are over the feasible and infeasible sets respectively, then the Feasibility Classification Problem yields a  $\delta$ -barrier.

COROLLARY 1. Let  $B^{(k)}(\mathbf{x}, \mathbf{u})$  be the optimal solution to  $\mathbf{FCP}(\mathbb{P}_{(\mathbf{x},\mathbf{u})}, \overline{\mathbb{P}}_{(\mathbf{x},\mathbf{u})})$  which achieves an optimal value of 0. At the optimum, the following statements are true:

1. If supp $(\overline{\mathbb{P}}_{(\mathbf{x},\mathbf{u})}) \supseteq \{(\mathbf{x},\mathbf{u}) \mid \mathbf{x} \in \mathbb{R}^n \setminus \mathcal{P}, \mathbf{u} \in \mathcal{U}\}, \text{ then for any } \mathbf{u} \in \mathcal{U}, \{\mathbf{x} \mid B^{(k)}(\mathbf{x},\mathbf{u}) > 0\} \subseteq \mathcal{P}.$ 

2. If  $\operatorname{supp}(\bar{\mathbb{P}}_{(\mathbf{x},\mathbf{u})}) \supseteq \{(\mathbf{x},\mathbf{u}) \mid \mathbf{x} \in \mathbb{R}^n \setminus \mathcal{P}, \mathbf{u} \in \mathcal{U}\}$  and  $\operatorname{supp}(\mathbb{P}_{(\mathbf{x},\mathbf{u})}) = \{(\mathbf{x},\mathbf{u}) \mid \mathbf{x} \in \mathcal{X}(\mathbf{u}), \mathbf{u} \in \mathcal{U}\}$ , then for any  $\zeta \in (0,1)$ , the scaled classifier  $\zeta B^{(k)}(\mathbf{x},\mathbf{u})$  is a  $\delta$ -barrier for some  $\delta \leq \Delta(\mathbf{u})$ .

Proof of Corollary 1.  $\mathbf{FCP}(\mathbb{P}_{(\mathbf{x},\mathbf{u})}, \bar{\mathbb{P}}_{(\mathbf{x},\mathbf{u})})$  achieves an optimal value of 0 if and only if  $B^{(k)}(\mathbf{x},\mathbf{u}) = 1$  for all  $(\mathbf{x},\mathbf{u}) \in \operatorname{supp}(\mathbb{P}_{(\mathbf{x},\mathbf{u})})$  and  $B^{(k)}(\mathbf{x},\mathbf{u}) = 0$  for all  $(\mathbf{x},\mathbf{u}) \in \operatorname{supp}(\bar{\mathbb{P}}_{(\mathbf{x},\mathbf{u})})$ . Thus,

$$\left\{ (\mathbf{x}, \mathbf{u}) \mid \mathbf{x} \in \mathbb{R}^n \setminus \mathcal{P}, \mathbf{u} \in \mathcal{U} \right\} \subseteq \operatorname{supp}(\bar{\mathbb{P}}_{(\mathbf{x}, \mathbf{u})}) \subseteq \left\{ (\mathbf{x}, \mathbf{u}) \mid B^{(k)}(\mathbf{x}, \mathbf{u}) = 0, \mathbf{u} \in \mathcal{U} \right\}.$$

For any fixed **u**, rearranging the set inclusions proves the first statement.

To prove the second statement, note that

$$\operatorname{supp}(\mathbb{P}_{(\mathbf{x},\mathbf{u})}) \subseteq \left\{ (\mathbf{x},\mathbf{u}) \mid B^{(k)}(\mathbf{x},\mathbf{u}) > 0 \right\} \subset \left\{ (\mathbf{x},\mathbf{u}) \mid \mathbf{x} \in \mathcal{P}, \mathbf{u} \in \mathcal{U} \right\},\$$

where the first inclusion follows by definition and the second from Statement 1 of Corollary 1. We fix  $\zeta \in (0,1)$  so that  $\zeta B^{(k)}(\mathbf{x},\mathbf{u}) \in [0,1)$ . Because a canonical barrier for  $\mathcal{P}$  is a  $\Delta(\mathbf{u})$ -barrier (see Remark 1) and  $\zeta B^{(k)}(\mathbf{x},\mathbf{u})$  has smaller support, the classifier is a  $\delta$ -barrier for  $\delta \leq \Delta(\mathbf{u})$ .

Corollary 1 states that given access to two disjoint distributions  $\mathbb{P}_{(\mathbf{x},\mathbf{u})}$  and  $\overline{\mathbb{P}}_{(\mathbf{x},\mathbf{u})}$  with closed supports,  $B^{(k)}(\mathbf{x},\mathbf{u})$  will learn a boundary between them. The first statement ensures that for any  $\mathbf{u}$ , the optimal classifier has a closed and bounded support that is smaller than  $\mathcal{P}$ . The second statement is a sufficient condition for  $B^{(k)}(\mathbf{x},\mathbf{u})$  to be a  $\delta$ -barrier. Furthermore, because  $\overline{\mathcal{D}}$  can be generated by sampling from  $\mathbb{R}^n \setminus \mathcal{P}$ , we therefore assume access to a distribution  $\overline{\mathbb{P}}_{(\mathbf{x},\mathbf{u})}$  supported over  $\{(\mathbf{x},\mathbf{u}) \mid \mathbf{x} \in \mathbb{R}^n \setminus \mathcal{P}, \mathbf{u} \in \mathcal{U}\}$  which implies that the first statement is always satisfied in practice. However, satisfying the second statement is contingent on access to a data set whose support is equal to  $\{(\mathbf{x},\mathbf{u}) \mid \mathbf{x} \in \mathcal{X}(\mathbf{u}), \mathbf{u} \in \mathcal{U}\}$ . In Step 2 of the IPMAN algorithm, the generated points are used to augment the two data sets and produce a better classifier.

### 4.3. In-sample optimality guarantees

We train a generative model  $F(\mathbf{u})$  to solve  $\mathbf{GBP}(\hat{\mathcal{U}}, B^{(k)}, \lambda_j)$  for a decreasing sequence of  $\lambda_j > 0$ . As a result,  $F(\mathbf{u})$  learns to predict optimal solutions to the barrier problem  $\mathbf{x}^{\lambda}(\mathbf{u})$  in an unsupervised fashion, i.e., without using  $\mathbf{x}^{\lambda}(\mathbf{u})$  or  $\mathbf{x}^{*}(\mathbf{u})$ . Moreover, the approximation error of  $F(\hat{\mathbf{u}}_i)$  versus  $\mathbf{x}^{*}(\hat{\mathbf{u}}_i)$  is also bounded, which we characterize below.

THEOREM 2. Fix  $\hat{\mathbf{u}}_i \in \hat{\mathcal{U}}$  and  $\lambda_j > 0$  and consider  $B^{(k)}(\mathbf{x}, \mathbf{u})$  and  $F^{(j,k)}(\mathbf{x}, \mathbf{u})$ . Let  $\mathbf{x}^{\lambda_j}$  be an optimal solution to  $\mathbf{BP}(\hat{\mathbf{u}}_i, B^{(k)}, \lambda_j)$  Then, there exists  $\delta, \epsilon > 0$  such that

$$\left|f\left(F^{(j,k)}(\hat{\mathbf{u}}_{i})\right) - f\left(\mathbf{x}^{*}(\hat{\mathbf{u}}_{i})\right)\right| < \left|f\left(F^{(j,k)}(\hat{\mathbf{u}}_{i})\right) - f\left(\mathbf{x}^{\lambda_{j}}(\hat{\mathbf{u}}_{i})\right)\right| + \max(\delta L, \epsilon)$$

Proof of Theorem 2. By the Triangle inequality,

$$\left|f\left(F^{(j,k)}(\hat{\mathbf{u}}_{i})\right) - f\left(\mathbf{x}^{*}(\hat{\mathbf{u}}_{i})\right)\right| \leq \left|f\left(F^{(j,k)}(\hat{\mathbf{u}}_{i})\right) - f\left(\mathbf{x}^{\lambda_{j}}(\hat{\mathbf{u}}_{i})\right)\right| + \left|f\left(\mathbf{x}^{\lambda_{j}}(\hat{\mathbf{u}}_{i})\right) - f\left(\mathbf{x}^{*}(\hat{\mathbf{u}}_{i})\right)\right|$$

We consider two cases: when  $B^{(k)}(\mathbf{x}^*, \hat{\mathbf{u}}_i) > 0$  and when  $B^{(k)}(\mathbf{x}^*, \hat{\mathbf{u}}_i) = 0$ .

First, if  $B^{(k)}(\mathbf{x}^*, \hat{\mathbf{u}}_i) > 0$ , then let  $\hat{\mathcal{X}}(\hat{\mathbf{u}}_i) = \mathcal{X}(\hat{\mathbf{u}}_i) \cap \{\mathbf{x} \mid B^{(k)}(\mathbf{x}, \hat{\mathbf{u}}_i) > 0\}$  be the correctly classified subset of  $\mathcal{X}(\hat{\mathbf{u}}_i)$ . We need only consider this subset as the feasible set when solving  $\mathbf{OP}(\hat{\mathbf{u}}_i)$ , since  $\mathbf{x}^*(\hat{\mathbf{u}}_i)$  remains feasible. However,  $B^{(k)}(\mathbf{x}, \hat{\mathbf{u}}_i)$  is a  $\delta$ -barrier for the subset. From Theorem 1,  $\mathbf{x}^{\lambda_j}(\hat{\mathbf{u}}_i)$  is  $(\delta, \epsilon)$ -optimal and we bound  $|f(\mathbf{x}^*(\hat{\mathbf{u}}_i)) - f(\mathbf{x}^{\lambda_j}(\hat{\mathbf{u}}_i))| < \max(\delta L, \epsilon)$ .

If  $B^{(k)}(\mathbf{x}^*, \hat{\mathbf{u}}_i) = 0$ , then the classifier is not a  $\delta$ -barrier for  $\mathcal{X}(\hat{\mathbf{u}}_i)$ . Instead, we will construct a "test"  $\delta$ -barrier from  $B^{(k)}(\mathbf{x}, \hat{\mathbf{u}}_i)$  and show that  $\mathbf{x}^{\lambda_j}(\hat{\mathbf{u}}_i)$  is still an optimal solution for this artificial barrier and thus, is  $(\delta, \epsilon)$ -optimal.

Fix a constant parameter  $\varepsilon > 0$ . Let  $\mathbf{x}^{\mathcal{P}} \in \arg\min_{\mathbf{x}} \{f(\mathbf{x}) \mid \mathbf{x} \in \mathcal{P}\}$  and let  $\overline{B}$  be defined as follows:

$$\frac{1}{\bar{B}} = \max\left\{1 + \varepsilon, \exp\left[\left(f\left(\mathbf{x}^{\lambda_{j}}(\hat{\mathbf{u}}_{i})\right) - f\left(\mathbf{x}^{\mathcal{P}}\right)\right)\frac{1}{\lambda_{j}} - \log B^{(k)}\left(\mathbf{x}^{\lambda_{j}}(\hat{\mathbf{u}}_{i}), \hat{\mathbf{u}}_{i}\right)\right]\right\}.$$
(9)

Note that  $\bar{B} \in (0, 1)$ . Then, the following function  $B^{\text{Test}}(\mathbf{x})$  is a  $\delta$ -barrier for  $\mathbf{OP}(\hat{\mathbf{u}}_i)$ :

$$B^{\text{Test}}(\mathbf{x}) = \begin{cases} B^{(k)}(\mathbf{x}, \hat{\mathbf{u}}_i), & \forall \mathbf{x} \in \{\mathbf{x} \mid B^{(k)}(\mathbf{x}, \hat{\mathbf{u}}_i) > 0\} \\ \bar{B}, & \forall \mathbf{x} \in \mathcal{X}(\hat{\mathbf{u}}_i) \cap \{\mathbf{x} \mid B^{(k)}(\mathbf{x}, \hat{\mathbf{u}}_i) = 0\}. \end{cases}$$

We show that  $\mathbf{x}^{\lambda_j}(\hat{\mathbf{u}}_i)$  is an optimal solution to  $\mathbf{BP}(\hat{\mathbf{u}}_i, B^{\text{Test}}, \lambda_j)$ . By definition,  $\mathbf{x}^{\lambda_j}(\hat{\mathbf{u}}_i)$  is optimal in  $\{\mathbf{x} \mid B^{(k)}(\mathbf{x}, \hat{\mathbf{u}}_i) > 0\}$ . To show that it is also optimal in  $\mathcal{X}(\hat{\mathbf{u}}_i) \cap \{\mathbf{x} \mid B^{(k)}(\mathbf{x}, \hat{\mathbf{u}}_i) = 0\}$ , we observe

$$-\log \bar{B} \ge \left( f\left(\mathbf{x}^{\lambda_{j}}(\hat{\mathbf{u}}_{i})\right) - f\left(\mathbf{x}^{\mathcal{P}}\right) \right) \frac{1}{\lambda_{j}} - \log B^{(k)}\left(\mathbf{x}^{\lambda_{j}}(\hat{\mathbf{u}}_{i}), \hat{\mathbf{u}}_{i}\right).$$

The above inequality is obtained by transforming the maximum in (9) to an inequality and taking the logarithm on both sides. Re-arranging this inequality yields

$$f(\mathbf{x}^{\lambda}(\hat{\mathbf{u}}_{i})) - \lambda_{j} \log B^{(k)}(\mathbf{x}^{\lambda_{j}}(\hat{\mathbf{u}}_{i}), \hat{\mathbf{u}}_{i}) \le f(\mathbf{x}^{\mathcal{P}}) - \lambda_{j} \log \bar{B}$$
(10)

$$\leq f(\mathbf{x}) - \lambda_j \log \bar{B}, \quad \forall \mathbf{x} \in \mathcal{X}(\hat{\mathbf{u}}_i).$$
 (11)

We obtain (11) because  $\mathcal{X}(\hat{\mathbf{u}}_i) \subset \mathcal{P}$  and consequently,  $f(\mathbf{x}^{\mathcal{P}}) \leq f(\mathbf{x})$  for all  $\mathbf{x} \in \mathcal{X}(\hat{\mathbf{u}}_i)$ . Therefore,  $\mathbf{x}^{\lambda_j}(\hat{\mathbf{u}}_i) \in \arg\min_{\mathbf{x}} \{f(\mathbf{x}) - \lambda_j \log B^{\text{Test}}(\mathbf{x})\}$ . From Theorem 1,  $\mathbf{x}^{\lambda_j}(\hat{\mathbf{u}}_i)$  is  $(\delta, \epsilon)$ -optimal.

Theorem 2 illustrates the key strengths and challenges with the IPMAN algorithm. Intuitively, the proof considers two cases: one where the classifier  $B^{(k)}(\mathbf{x}, \mathbf{u})$  is a  $\delta$ -barrier for a given  $OP(\mathbf{u})$ , and one where it isn't. If the classifier is a  $\delta$ -barrier, then it is straightforward to bound the quality of the model predictions by directly optimizing  $BP(\mathbf{u}, B^{(k)}, \lambda_j)$  and comparing against the optimal solution. However, the learned classifier may not yet be a  $\delta$ -barrier if there are insufficient points. In this scenario, we can still bound the  $(\delta, \epsilon)$ -optimality of predicted solutions by artificially constructing a  $\delta$ -barrier from the classifier. Therefore, the in-sample performance of  $F^{(j,k)}(\mathbf{u})$  can always be measured at any iteration. Unlike for a  $\delta$ -barrier, however, the  $(\delta, \epsilon)$ -optimality bound using a classifier does not necessarily converge to 0 as we decrease  $\lambda_j$ .

#### 4.4. Data augmentation via the oracle

At every iteration, the oracle evaluates the generative models by labelling the predictions as feasible or infeasible. If  $\Psi(F^{(j,k)}(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i) = 1$  for any  $\hat{\mathbf{u}}_i$ , then  $(F^{(j,k)}(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i)$  is added to the data set of feasible solutions  $\mathcal{D}^{(k)}$ . Otherwise, it is added to  $\overline{\mathcal{D}}^{(k)}$ . Consequently, these two data sets grow after each iteration and the augmented sets are used to train  $B^{(k+1)}(\mathbf{x}, \mathbf{u})$ . This data augmentation procedure implies that the classifier can learn to become a tighter approximation of  $\mathcal{X}(\mathbf{u})$ .

PROPOSITION 1. For any k, let  $\mathcal{B}^{(k)}$  be the optimal solution set of  $\mathbf{FCP}(\mathcal{D}^{(k)}, \overline{\mathcal{D}}^{(k)})$ . If  $\mathcal{B}$  satisfies Assumption 1 and the data sets are closed, then  $\mathcal{B}^{(k+1)} \subset \mathcal{B}^{(k)}$ .

Proof of Proposition 1 From Lemma 1, the optimal value of  $\mathbf{FCP}(\mathcal{D}^{(k)}, \bar{\mathcal{D}}^{(k)})$  is 0, since the data sets are closed and disjoint. The augmentations  $\mathcal{Q}$  and  $\bar{\mathcal{Q}}$  are also closed and disjoint:

$$\mathcal{Q} \subseteq \{(\mathbf{x}, \mathbf{u}) \mid \mathbf{x} \in \mathcal{X}(\mathbf{u}), \mathbf{u} \in \mathcal{U}\}$$
  
 $\bar{\mathcal{Q}} \subset \{(\mathbf{x}, \mathbf{u}) \mid \mathbf{x} \in \mathbb{R}^n \setminus \mathcal{X}(\mathbf{u}), \mathbf{u} \in \mathcal{U}\}$ 

Therefore,  $\mathcal{D}^{(k+1)} = \mathcal{D}^{(k)} \cup \mathcal{Q}$  and  $\bar{\mathcal{D}}^{(k+1)} = \bar{\mathcal{D}}^{(k)} \cup \bar{\mathcal{Q}}$  are both closed and disjoint. From Lemma 1, the optimal value of  $\mathbf{FCP}(\mathcal{D}^{(k+1)}, \bar{\mathcal{D}}^{(k+1)})$  is also 0.

To show  $\mathcal{B}^{(k+1)} \subset \mathcal{B}^{(k)}$ , we first prove  $\mathcal{B}^{(k+1)} \subseteq \mathcal{B}^{(k)}$  and then present a counter-example which disproves equivalence. The objective function of  $\mathbf{FCP}(\mathcal{D}^{(k+1)}, \bar{\mathcal{D}}^{(k+1)})$  is

$$\frac{1}{|\mathcal{D}^{(k+1)}|} \sum_{(\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \mathcal{D}^{(k+1)}} \log B(\hat{\mathbf{x}}, \hat{\mathbf{u}}) + \frac{1}{|\bar{\mathcal{D}}^{(k+1)}|} \sum_{(\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \bar{\mathcal{D}}^{(k+1)}} \log \left(1 - B(\hat{\mathbf{x}}, \hat{\mathbf{u}})\right)$$
$$= \frac{\alpha}{|\mathcal{D}^{(k)}|} \sum_{(\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \mathcal{D}^{(k)}} \log B(\hat{\mathbf{x}}, \hat{\mathbf{u}}) + \frac{1 - \alpha}{|\mathcal{Q}|} \sum_{(\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \mathcal{Q}} \log B(\hat{\mathbf{x}}, \hat{\mathbf{u}})$$
$$+ \frac{\alpha'}{|\bar{\mathcal{D}}^{(k)}|} \sum_{(\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \bar{\mathcal{D}}^{(k)}} \log \left(1 - B(\hat{\mathbf{x}}, \hat{\mathbf{u}})\right) + \frac{1 - \alpha'}{|\bar{\mathcal{Q}}|} \sum_{(\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \bar{\mathcal{Q}}} \log \left(1 - B(\hat{\mathbf{x}}, \hat{\mathbf{u}})\right),$$

where  $\alpha = |\mathcal{D}^{(k)}|/|\mathcal{D}^{(k+1)}|$  and  $\alpha' = |\bar{\mathcal{D}}^{(k)}|/|\bar{\mathcal{D}}^{(k+1)}|$  are the mixture weights defining the ratio of existing to new points in each data set. Because the optimal value of  $\mathbf{FCP}(\mathcal{D}^{(k+1)}, \bar{\mathcal{D}}^{(k+1)})$  is 0 and  $B(\mathbf{x}, \mathbf{u}) \in [0, 1]$ , each of the individual terms must be equal to 0 for an optimal solution. However, the first and third terms define the objective function for  $\mathbf{FCP}(\mathcal{D}^{(k)}, \bar{\mathcal{D}}^{(k)})$ . Thus, any optimal solution  $B^{(k+1)}$  to  $\mathbf{FCP}(\mathcal{D}^{(k+1)}, \bar{\mathcal{D}}^{(k+1)})$  must also be optimal for  $\mathbf{FCP}(\mathcal{D}^{(k)}, \bar{\mathcal{D}}^{(k)})$  implying  $\mathcal{B}^{(k+1)} \subseteq \mathcal{B}^{(k)}$ .

To prove the inclusion is strict, consider the closed and disjoint sets  $\mathcal{D}^{(k)} \cup \bar{\mathcal{Q}}$  and  $\bar{\mathcal{D}}^{(k)}$ . By Lemma 1, there exists a function  $B^*(\mathbf{x}, \mathbf{u})$  such that  $B^*(\mathbf{x}, \mathbf{u}) = 1$  for all  $(\mathbf{x}, \mathbf{u}) \in \mathcal{D}^{(k)} \cup \bar{\mathcal{Q}}$  and  $B^*(\mathbf{x}, \mathbf{u}) = 0$  for all  $(\mathbf{x}, \mathbf{u}) \in \bar{\mathcal{D}}^{(k)}$ , i.e.,  $B^* \in \mathcal{B}^{(k)}$ . However, then  $B^*(\hat{\mathbf{x}}, \hat{\mathbf{u}}) = 1$  for all  $(\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \bar{\mathcal{Q}}$  and  $B^*(\mathbf{x}, \mathbf{u})$  has an infinite objective function value for  $\mathbf{FCP}(\mathcal{D}^{(k+1)}, \bar{\mathcal{D}}^{(k+1)})$ . Thus,  $B^* \notin \mathcal{B}^{(k+1)}$ . After each iteration of the IPMAN algorithm,  $\mathcal{D}^{(k)}$  and  $\bar{\mathcal{D}}^{(k)}$  are augmented. By augmenting  $\bar{\mathcal{D}}^{(k)}$ , we correct regions that the classifier has incorrectly labelled as feasible. On the other hand, augmenting  $\mathcal{D}^{(k)}$  reinforces regions where the classifier has correctly labelled points so that it does not incorrectly mislabel the region in a subsequent iteration.

Because  $\mathcal{X}(\mathbf{u})$  is not known, it is difficult to determine whether the classifier is, in fact, a  $\delta$ -barrier for a given  $\mathbf{u}$ . From the proof of Theorem 2, we know that even if the classifier is not a  $\delta$ -barrier, there exists an equivalent  $\delta$ -barrier for the classifier as well as a fixed  $\lambda > 0$ . It remains, therefore, to estimate the value of  $\delta$  for the classifier in order to fully characterize the optimality bound.

COROLLARY 2. If  $B(\mathbf{x}, \hat{\mathbf{u}}_i)$  is a  $\delta$ -barrier for  $\mathbf{OP}(\hat{\mathbf{u}}_i)$ , then  $\delta \leq d_H(\{\hat{\mathbf{x}}_i \mid (\hat{\mathbf{x}}_i, \hat{\mathbf{u}}_i) \in \mathcal{D}^{(k)}\}, \mathrm{bd}(\mathcal{P}))$ .

Proof of Corollary 2. By definition,  $\delta = d_H(\mathcal{X}(\hat{\mathbf{u}}_i), \{\mathbf{x} \mid B(\mathbf{x}, \hat{\mathbf{u}}_i) > 0\})$ , where  $\mathcal{X}(\hat{\mathbf{u}}_i) \subseteq \{\mathbf{x} \mid B(\mathbf{x}, \hat{\mathbf{u}}_i) > 0\}$ . Then,

$$d_H \left( \mathcal{X}(\hat{\mathbf{u}}_i), \{ \mathbf{x} \mid B(\mathbf{x}, \hat{\mathbf{u}}_i) > 0 \} \right) \le d_H \left( \{ \hat{\mathbf{x}}_i \mid (\hat{\mathbf{x}}_i, \hat{\mathbf{u}}_i) \in \mathcal{D}^{(k)} \}, \{ \mathbf{x} \mid B(\mathbf{x}, \hat{\mathbf{u}}_i) > 0 \} \right)$$
$$\le d_H \left( \{ \hat{\mathbf{x}}_i \mid (\hat{\mathbf{x}}_i, \hat{\mathbf{u}}_i) \in \mathcal{D}^{(k)} \}, \operatorname{bd}(\mathcal{P}) \right).$$

The first inequality follows from  $\{\hat{\mathbf{x}}_i \mid (\hat{\mathbf{x}}_i, \hat{\mathbf{u}}_i) \in \mathcal{D}^{(k)}\} \subset \mathcal{X}(\hat{\mathbf{u}}_i)$ , and the second from  $\{\mathbf{x} \mid B(\mathbf{x}, \hat{\mathbf{u}}_i) > 0\} \subseteq \mathcal{P}$  and that the furthest point in  $\mathcal{P}$  from  $\{\hat{\mathbf{x}}_i \mid (\hat{\mathbf{x}}_i, \hat{\mathbf{u}}_i) \in \mathcal{D}^{(k)}\}$  is on the boundary.  $\Box$ By generating additional feasible points after each iteration, we can obtain a smaller  $\delta$ . Further, to compute the Hausdorff distance, we must calculate the distance of a fixed set of points to each of the facets of the polyhedron. This problem can be solved via a finite number of linear programs.

# 5. Generalization of $(\delta, \epsilon)$ -optimality to unseen instances

In this section, we evaluate the potential for a generative model to predict a  $(\delta, \epsilon)$ -optimal solution when given an unseen out-of-sample problem instance **u**. Our analysis applies on any model and is independent of the IPMAN algorithm itself. However, this analysis, when applied to IPMAN, offers an opportunity to understand the true quality of the trained model at any iteration of the algorithm. Although the optimal solution  $\mathbf{x}^*(\mathbf{u})$  is not known, the objective function error between  $f(F(\mathbf{u}))$  and the optimal value  $f(\mathbf{x}^*(\mathbf{u}))$  can be bounded using the Triangle inequality

$$\left|f\left(F^{*}(\mathbf{u})\right) - f\left(\mathbf{x}^{*}(\mathbf{u})\right)\right| < \left|f\left(F^{*}(\mathbf{u})\right) - f\left(\mathbf{x}^{\lambda}(\mathbf{u})\right)\right| + \left|f\left(\mathbf{x}^{*}(\mathbf{u})\right) - f\left(\mathbf{x}^{\lambda}(\mathbf{u})\right)\right|$$

From Theorem 1, the second term is bounded above by  $\max\{\delta L, \epsilon\}$ . Therefore, it only remains to bound the empirical error of  $F(\mathbf{u})$  versus  $\mathbf{x}^{\lambda}(\mathbf{u})$  (i.e., the first term on the right-hand side).

We use Rademacher complexity theory to obtain a probabilistic bound on the empirical error from an out-of-sample input (Bartlett and Mendelson 2002). Bertsimas and Kallus (2019) develop generalization bounds for predicting decisions to problems with conditional stochastic optimization objectives. We extend their work by providing a probabilistic bound on  $(\delta, \epsilon)$ -optimality when the feasible set is not fully specified. DEFINITION 2 (BERTSIMAS AND KALLUS (2019)). Let  $\mathcal{F} \subset \{F(\mathbf{u}) : \mathcal{U} \to \mathbb{R}^n\}$  be a function class and  $\hat{\mathcal{U}} \sim \mathbb{P}_{\mathbf{u}}$  be an i.i.d. data set. The empirical multivariate Rademacher complexity of  $\mathcal{F}$  is

$$\hat{\mathfrak{R}}_{N_{\mathbf{u}}}(\mathcal{F},\hat{\mathcal{U}}) = \mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}} \left[ \frac{2}{N_{\mathbf{u}}} \sup_{F \in \mathcal{F}} \sum_{i=1}^{N_{\mathbf{u}}} \boldsymbol{\sigma}_{i}^{\mathsf{T}} F(\hat{\mathbf{u}}_{i}) \middle| \hat{\mathcal{U}} = \{\hat{\mathbf{u}}_{i}\}_{i=1}^{N_{\mathbf{u}}} \right],$$

where  $\sigma_i \sim p_{\sigma}$  is an *n*-dimensional vector of i.i.d. Rademacher variables. The multivariate Rademacher complexity of  $\mathcal{F}$  is  $\mathfrak{R}_{N_{\mathbf{u}}}(\mathcal{F}) = \mathbb{E}_{\hat{\mathcal{U}} \sim \mathbb{P}_{\mathbf{u}}} \left[ \hat{\mathfrak{R}}_{N_{\mathbf{u}}}(\mathcal{F}, \hat{\mathcal{U}}) \right].$ 

Rademacher complexities are often used to generate risk bounds in various statistical learning settings (Bartlett and Mendelson 2002). In order to ensure these bounds are practical, it is important to use model classes  $\mathcal{F}$  whose Rademacher complexities are bounded above as a function of the data and parametrization. For generality, we leave the choice of model class open.

REMARK 5. Although the literature mostly focuses on the single-variate Rademacher complexity, Bertsimas and Kallus (2019) and Maurer (2016) prove bounds for several linear multivariate classes (e.g.,  $\mathcal{F}_R = \{\mathbf{Wu} \mid ||\mathbf{W}|| \leq R\}$ ). In general, if  $F(\mathbf{u}) = (F_1(\mathbf{u}), \ldots, F_n(\mathbf{u}))$ , then  $\mathcal{F} \subset$  $\times_{\ell=1}^n \mathcal{F}_\ell$ , where  $\mathcal{F}_\ell = \{F(\mathbf{u})^\mathsf{T} \mathbf{e}_\ell \mid F \in \mathcal{F}\}$  and  $\mathbf{e}_\ell$  is the  $\ell$ -th identity vector. Then,  $\hat{\mathfrak{R}}_{N_{\mathbf{u}}}(\mathcal{F}, \hat{\mathcal{U}}) \leq$  $\sum_{\ell=1}^n \hat{\mathfrak{R}}_{N_{\mathbf{u}}}(\mathcal{F}_\ell, \hat{\mathcal{U}})$  decomposes to a sum of single-variate complexities. We refer to Bartlett and Mendelson (2002) for Rademacher bounds for linear models and decision trees and Bartlett and Mendelson (2002), Neyshabur et al. (2015), Foster et al. (2018) for neural networks.

 $\mathbf{GBP}(\hat{\mathcal{U}}, B_{\delta}, \lambda)$  is trained using a finite data set  $\hat{\mathcal{U}}$ , meaning there is no guarantee whether  $F(\mathbf{u})$ will be feasible or even satisfy  $B_{\delta}(F(\mathbf{u}), \mathbf{u}) > 0$  for an arbitrary  $\mathbf{u}$ . Because  $\mathcal{P}$  is available, however, we can always project any generated solution to the polyhedron.

ASSUMPTION 2. If  $F^{(j,k)}$  is an optimal solution to  $\mathbf{GBP}(\hat{\mathcal{U}}, B_{\delta}, \lambda)$ , then the projected function  $F^*(\mathbf{u}) = \arg\min_{\mathbf{x}} \left\{ \left\| \mathbf{x} - F^{(j,k)}(\mathbf{u}) \right\| \mid \mathbf{x} \in \mathcal{P} \right\}$  is used at test time.

Our generalization bound follows from the bound of Bertsimas and Kallus (2019). While they derive an empirical risk bound on an unconstrained stochastic optimization problem, we focus on  $(\delta, \epsilon)$ -optimality for a constrained continuous optimization problem. We provide the proof in EC.2.

THEOREM 3. Let  $F^*$  satisfy Assumption 2. Let K and  $L_{\infty}$  be sufficiently large positive constants. Let  $\beta \in (0,1)$  be a constant. Then, for any  $\gamma > 0$ , the following inequality holds

$$\begin{split} \mathbb{P}_{\mathbf{u}}\Big\{f\big(F^*(\mathbf{u})\big) - \epsilon - \gamma < f\big(\mathbf{x}^*(\mathbf{u})\big) < f\big(F^*(\mathbf{u})\big) + \delta L + \gamma\Big\} \\ & \geq 1 - \frac{\frac{1}{N_{\mathbf{u}}}\sum_{i=1}^{N_{\mathbf{u}}} \left|f\big(F^*(\hat{\mathbf{u}}_i)\big) - f\big(\mathbf{x}^{\lambda}(\hat{\mathbf{u}}_i)\big)\right| + K\sqrt{\frac{\log(1/\beta)}{2N_{\mathbf{u}}}} + \sqrt{2n}L_{\infty}\mathfrak{R}_{N_{\mathbf{u}}}(\mathcal{F})}{\gamma}, \end{split}$$

with probability at least  $1 - \beta$  with respect to the sampling of  $\hat{\mathcal{U}}$ .

The quality of this bound is dependent on the generative model itself. The first term in the numerator,  $(\sum_{i=1}^{N_{\mathbf{u}}} |f(F^*(\hat{\mathbf{u}}_i)) - f(\mathbf{x}^{\lambda}(\hat{\mathbf{u}}_i))|)/N_{\mathbf{u}}$ , is the empirical error of solving  $\mathbf{GBP}(\hat{\mathcal{U}}, B_{\delta}, \lambda)$  versus  $\mathbf{BP}(\hat{\mathbf{u}}_i)$  for all  $\hat{\mathbf{u}}_i \in \hat{\mathcal{U}}$ . This effectively measures how well the model performs on in-sample data. The second term is dependent on the constants K and  $1/\beta$  and scales with  $O(1/N_{\mathbf{u}})$ . Finally, the third term is dependent on the Rademacher complexity of  $\mathcal{F}$ . Thus, in order to obtain a tight and useful bound, we must balance the trade-off between a model class with high complexity versus obtaining a final model with low empirical error.

Theorem 3 is a bound on the  $(\delta, \epsilon)$ -optimality of a random out-of-sample **u**. Note that we require two distinct probability statements to describe this bound. The explicit statement is calculating, given a  $\gamma$  and  $F^*$ , the probability that the model will predict a  $(\delta + \gamma/L, \epsilon + \gamma)$ -optimal solution for a random  $\mathbf{u}_{N_{\mathbf{u}+1}} \sim \mathbb{P}_{\mathbf{u}}$ . The probability of this event is bounded from below by the right-hand-side of the relation in Theorem 3. However, the implicit statement is that this bound on the probability will only hold with probability at least  $1 - \beta$ .

# 6. Predicting optimal dose in automated radiation therapy treatment planning

We implement IPMAN to predict optimal dose distributions for patients with head-and-neck cancer. Treatment planning in head-and-neck cancer requires balancing the dose delivered to several structures (i.e., organs-at-risk and targets). Clinics generally have strict institution-specific criteria that their RT plans should meet. However, it is often impossible to construct doses distributions that satisfy all of the criteria for a given patient, meaning an oncologist must determine which criteria to prioritize. Such choices of prioritization are effectively hidden constraints known only to an oncologist and obtained from their intuition and expertise rather than formal reasoning. In the current practice however, the choices are facilitated via a surrogate multi-objective optimization problem with parametric weights that are manually tuned.

Automated planning methods aim to predict treatments that achieve the desired dose tradeoffs while reducing oncologist intervention. An automated planning pipeline generally involves two stages. First, a machine learning model uses a computed tomography (CT) image of the patient to predict the dose that should be delivered. While the prediction may estimate which structures to prioritize, the prediction often fails to fully meet the exact clinical criteria or is physically undeliverable by the treatment machine. Consequently, a second stage optimization model is used to correct and transform the prediction into a deliverable and clinically satisfactory treatment plan.

In this work, we recast the task of the machine learning problem from predicting a clinically acceptable dose distribution to constructing an optimal dose distribution for a given patient. The selection of clinical criteria to satisfy for a patient is modeled as a latent choice constraint dependent on the CT image, which is both specific to a patient and the institution providing care. Specifically, a given dose is *feasible* if it satisfies the same set of criteria that the oncologist prescribed for the patient. The objective function minimizes dose to the organs-at-risk (OARs), which describe the healthy tissue in the patient.

A single objective for all patients allows us to form a standard notion of treatment optimality based on minimizing radiation to healthy tissue. We show in our comparisons with benchmark prediction models that treatment plans generated from IPMAN both (i) capture the same clinical trade-offs that oncologists would prescribe after evaluation, and (ii) deliver the same or lower dose on average to healthy tissue. An added benefit of our approach is that IPMAN can be adapted so that it learns institution-specific criteria without training on a new data set of delivered plans. In particular, we use the oracle to learn a constraint that was not present in the original data to demonstrate how IPMAN can be deployed at cancer centers with different clinical criteria.

# 6.1. Data and model

We use a data set of 217 clinical treatment plans for patients with head-and-neck cancer, randomly split into 100, 67, and 50 plans for training, validation, and held-out testing, respectively. Each patient is discretized into a  $128 \times 128 \times 128$  tensor whose elements represent voxels (volumetric pixels) of the patient's geometry. For each patient *i*, the clinical data set contains a CT image  $\hat{\mathbf{u}}_i \in \mathbb{R}^{128 \times 128 \times 128}$  (representing the auxiliary input) and a 3-D dose distribution  $\hat{\mathbf{x}}_i \in \mathbb{R}^{128 \times 128 \times 128}$  (representing the treatment decision) that was prescribed by the oncologist. Each patient contains up to four OARs and three tumor volumes, referred to as planning target volumes (PTVs), that have been contoured and labelled.

Let  $\mathcal{O}$  and  $\mathcal{T}$  index the OARs and PTVs, respectively, and let  $\mathcal{R} := \mathcal{O} \cup \mathcal{T}$  index all structures of interest. For each structure, let  $\mathcal{V}_r$  index the corresponding voxel set (elements of  $\mathbf{x}$  and  $\mathbf{u}$ ). Let  $z_r$ denote the average dose delivered to structure r and  $\mathbf{z} \in \mathbb{R}^7$  denote the vector of  $z_r$ . To illustrate the IPMAN methodology, we formulate an RT optimization problem that minimizes the sum of average doses delivered to the OARs subject to satisfying the relevant clinical criteria and known polyhedral constraints (see Babier et al. 2018b). This model closely approximates the traditional weighted optimization models that are used as a surrogate to the treatment planning problem. Although the objective is simplified for the sake of computational efficiency, the constraints are representative of the realistic clinical problem, which is the focus of our methodology (more discussion on model choice is given in EC.3.1). The clinical criteria for each OAR is an upper bound on either the mean  $z_o$  or maximum dose delivered  $z_o^{\max}$ , while the criteria for each PTV is a lower bound on the minimum dose delivered to the 90-th percentile  $z_t^{90} \in \mathbb{R}^3$ ) of the target structure, a Value-at-Risk (VaR) metric. We formulate each of the clinical criteria as a hidden bound  $\hat{z}_r(\mathbf{u})$  dependent on the patient geometry and oncologist choice. The optimization problem is summarized below:

$$\mathbf{RT}(\mathbf{u}): \min_{\mathbf{x}, \mathbf{z}, \mathbf{z}^{\max}, \mathbf{z}^{90}} \frac{1}{|\mathcal{O}|} \sum_{o \in \mathcal{O}} z_o$$
(12a)

subject to 
$$z_r = \frac{1}{|\mathcal{V}_r|} \sum_{v \in \mathcal{V}_r} x_v$$
  $r \in \mathcal{R}$  (12b)

$$z_o^{\max} \ge x_v$$
  $o \in \mathcal{O}, v \in \mathcal{V}_o$  (12c)

$$z_t^{90} = \operatorname{VaR}_{90}(\{x_v \mid v \in \mathcal{V}_t\}) \quad t \in \mathcal{T}$$
(12d)

$$\underline{z}_r \le z_r \le \overline{z}_r \qquad \qquad r \in \mathcal{R} \tag{12e}$$

$$z_o \leq \hat{z}_o(\mathbf{u})$$
  $o \in \{\text{Right Parotid, Left Parotid, Larynx}\}$  (12f)

$$z_0^{\max} \le \hat{z}_o(\mathbf{u}) \qquad o \in \{\text{Mandible}\}$$
 (12g)

$$z_t^{90} \ge \hat{z}_t(\mathbf{u}) \qquad \qquad t \in \mathcal{T}. \tag{12h}$$

Constraints (12b)–(12d) define the dose summary statistics of the mean, maximum, and VaR. In (12e), we mandate a fixed set of polyhedral constraints on  $\mathbf{z}$  obtained by calculating the maximum and minimum mean doses over all patients in the ground truth data set; this constitutes the polyhedral relaxation  $\mathcal{P}$ . Finally, (12f)–(12h) define the hidden patient-specific constraints, i.e., the clinical criteria that must be learned. Whereas (12f) and (12g) ensure that the dose delivered to each OAR is below a threshold, (12h) ensures that PTVs receive a sufficiently high dose of radiation. Because it is not often possible to simultaneously satisfy all clinical criteria, these hidden constraints are conditional. If the ground truth plan from the data set satisfied a hidden constraint (i.e., an oncologist deemed it necessary), we require that a generated plan must satisfy it as well. In other words for any patient  $\hat{\mathbf{u}}_i$  in our clinical data set,

$$\hat{z}_r(\hat{\mathbf{u}}_i) = \begin{cases} \hat{z}_r & \text{if the ground truth dose for } \hat{\mathbf{u}}_i \text{ satisfies the bound in Table 1} \\ +\infty & \text{otherwise for } r \in \mathcal{O} \\ 0 & \text{otherwise for } r \in \mathcal{T} \end{cases}$$

The hidden nature of these constraints arises from the fact that a planner does not know a priori if the constraint is needed. Note that the VaR clinical criteria for each PTV is a non-convex constraint and thus, the model would be difficult to solve directly even if the hidden constraints were known. The values for the bounds  $\hat{z}_r$  are given in Table 1 (column 2).

# 6.2. Methods

We use two state-of-the-art benchmark models to analyze the quality of the predictions produced by IPMAN: a U-net convolutional neural network (CNN) (Nguyen et al. 2019) and a generative adversarial network (GAN) (Mahmood et al. 2018). Nguyen et al. (2019) implement a CNN for predicting dose distributions from 3-D CT images and show its effectiveness in the prediction stage of automated planning. The CNN is trained via supervised learning by minimizing an  $l_2$  norm of predicted doses from a ground truth clinical data set. Mahmood et al. (2018) implement a conditional generative adversarial network to predict dose distributions from CT images. A GAN is composed of two networks that are trained adversarially. The first is a generator  $G(\mathbf{u})$  that creates sample dose distributions from CT images, while the second is a discriminator  $D(\mathbf{x}, \mathbf{u})$  that predicts whether a given sample was generated or belongs to a real data set. This discriminator is used to help train the generator in producing more "realistic" doses.

As the generator and classifier in IPMAN play similar roles to the generator and discriminator in a GAN, we use a slightly modified architecture from Mahmood et al. (2018) to create  $F(\mathbf{u})$  and  $B(\mathbf{x}, \mathbf{u})$  (details are provided in EC.3.2). Specifically for the experiments in Section 6.3, we include an  $l_1$  regularization term  $||F(\mathbf{u}) - \mathbf{u}||_1$  to the loss function of  $\mathbf{GBP}(\hat{\mathcal{U}}, B, \lambda_j)$ , which is commonly used for model stability in Style Transfer GANs (e.g., Isola et al. 2017). All models are trained using the Adam optimizer with  $(\beta_1, \beta_2) = (0.5, 0.999)$ . We train the classifier with a learning rate of  $1 \times 10^{-3}$  and the generator with a learning rate of  $2 \times 10^{-5}$ . Initially, the data set of feasible decisions  $\mathcal{D} = \{(\hat{\mathbf{u}}_i, \hat{\mathbf{x}}_i)\}_{i=1}^{N_{\mathbf{u}}}$  consists solely of the 100 clinical plans used in training, the data set of parameters  $\hat{\mathcal{U}} = \{\hat{\mathbf{u}}_i\}_{i=1}^{N_{\mathbf{u}}}$  contains their corresponding patient CT images, and the data set of infeasible decisions is empty (i.e.,  $\bar{\mathcal{D}} = \emptyset$ ).

Using the training set of patients, we train the generator and classifier iteratively with IPMAN. At the end of each iteration, we evaluate the predictions made by the generator on the validation set of patients. After 11 iterations, constraint satisfaction on the validation set stabilizes and we use the held-out test set to assess the performance of our models.

We train four generative models corresponding to  $\lambda \in \{256, 64, 16, 4\}$ . In each iteration, the generator predicts solutions to the corresponding barrier problem, meaning that training over a range of  $\lambda$  ensures that in every iteration, the oracle labels and augments the data sets with predictions lying in a diverse set of areas in and around the feasible set. Similarly, we train each baseline model for 200 epochs (approximately 24 hours). This is roughly the same duration of time required to train 11 iterations of IPMAN, thus maintaining a fair comparison. The data generation, implementation, and training details are provided in EC.3.3. Below, we highlight the initialization steps and refinements to the IPMAN algorithm made for our computational experiments.

6.2.1. Pre-training to initial feasible decisions. Classical IPMs generally require an initial point  $\mathbf{x}^{(0)}$  that is strictly feasible. Analogously in IPMAN, ensuring that  $F(\mathbf{u})$  is initialized at a stage where it usually predicts feasible decisions implies that the training loss is not extremely high at early stages and helps to stabilize training. Thus, before starting the IPMAN algorithm, we pre-train  $F(\mathbf{u})$  as the generator in a GAN and save the weights.

Structure	Criteria (Gy)	Baselines		IPMAN $(\lambda)$			
		GAN	CNN	256	64	16	4
Right Parotid	$z_o \leq 26$	85.7	85.7	86.2	90.0	93.3	100
Left Parotid	$z_o \leq 26$	70.0	60.0	70.0	90.0	90.0	100
Larynx	$z_o \leq 45$	93.3	83.3	89.7	89.7	93.3	100
Mandible	$z_o^{\max} \leq 73.5$	100	100	100	100	100	100
PTV70	$z_t^{90} \ge 70$	97.6	97.6	97.6	97.6	95.2	92.8
PTV63	$z_t^{90} \ge 63$	96.3	96.3	96.3	96.3	96.3	96.3
PTV56	$z_t^{90} \ge 56$	100	100	100	100	100	100
All hidden constraints		86.0	78.0	82.0	88.0	88.0	94.0
All polyhedral constraints		92.0	90.0	94.0	92.0	90.0	94.0
Objective function value		40.3	41.0	41.0	41.0	40.0	37.8

Table 1The percentage of predicted decisions on the held-out test set that satisfy each hidden constraint to 1Gy relaxation. The best performing models on the summary statistics are highlighted.

6.2.2. Generating an initial  $\overline{\mathcal{D}}$ . Before training the classifier in the first iteration of IPMAN, we require an initial data set of infeasible decisions  $\overline{\mathcal{D}}$ . During the pre-training stage, the generator of the GAN creates a set of candidate decisions. We label the generated decisions and assign them to the appropriate sets  $\mathcal{D}$  and  $\overline{\mathcal{D}}$  in order to initialize IPMAN with an augmented data set.

6.2.3. Learning multi-label feasibility. A feasible dose distribution must satisfy multiple polyhedral and hidden constraints corresponding to different PTVs and OARs. Learning to classify a decision as feasible is challenging due to the granularity of constraint satisfaction and the variety of constraints that are present. Consequently, we separate the classification problem into one for each of the four OARs and three PTVs of the patient. That is, for each structure of interest, we train a separate classifier. The  $\delta$ -barrier optimization problem is then the sum of the different classifier outputs; this is equivalent to the classical barrier problem. That is, let  $B_r(\mathbf{x}, \mathbf{u})$  be a classifier for  $r \in \mathcal{R}$  that assesses whether the polyhedral and hidden constraints for that structure are satisfied. Then, the barrier problem is min<sub>**x**</sub>{ $f(\mathbf{x}) - \lambda \sum_{r \in \mathcal{R}} B_r(\mathbf{x}, \mathbf{u})$ }.

## 6.3. Learning to predict feasible and optimal dose distributions

In practice, oncologists evaluate plans by assessing how many of the institutionally mandated clinical criteria are satisfied and to what degree. We use a set of clinical criteria (see column 2 of Table 1) to evaluate our models. As the relevancy of criteria (i.e., feasibility) for each patient is determined by an oncologist, we evaluate plans by first identifying the hidden constraints satisfied by the ground truth dose distribution. If the clinical dose satisfies a given criterion, we then evaluate whether the predicted dose also satisfies that criterion. Satisfaction is defined as meeting the dose bound to within a 1 Gy relaxation. In the clinical literature, dose predictions are commonly



(a) Objective function and fraction of feasible plans with respect to the hidden and polyhedral constraints.



(b) Average difference from the hidden constraint bound  $\hat{z}_r(\mathbf{u})$ . The dashed line is 0 Gy. Above 0 suggests plans satisfy the constraints on average.

Figure 3 Statistics on the validation set obtained during training on criteria from our institution.

evaluated on a voxel-level to within 3% of the maximum prescribed dose, i.e., 2.1 Gy for our problem (Low et al. 1998). In our analysis, we consider constraints rather than direct voxels and tighten the tolerance to 1 Gy. The relaxation can also be interpreted as the  $\delta$  for a  $\delta$ -barrier; if all decisions satisfy a given constraint, the the corresponding classifier is a  $\delta$ -barrier with  $\delta = 1$  Gy.

Figure 3(a) displays the average objective function value and the fraction of plans that satisfy the hidden and polyhedral constraints over training iterations. We find that in training, the objective function value improves as a function of the number of iterations, while hidden constraint satisfaction also increases. For example, in the first iteration, 89% of predictions in the validation set satisfy all of their hidden constraints, whereas this fraction increases to 97% by iteration 11. Polyhedral constraint satisfaction also increases from 88% to 95%. This suggests that the IPMAN algorithm trains the model to generate fewer infeasible doses. In other words, the classifier is learning to produce a tighter characterization of the feasible set (see Proposition 1).

Figure 3(b) shows the average difference from the boundary of the hidden constraint for each structure. If the difference is positive, doses on average satisfy the hidden criteria. We observe

two important phenomena. First, the four leftmost plots are associated with OAR constraints. By minimizing the objective, the associated OAR constraints see progressively better adherence as expected. Note that the Mandible and PTV70 structures often overlap, meaning their constraints conflict with each other, preventing improvement for this organ. Second, the PTV constraints show small but sustained improvement as the number of training iterations increase. This is because they are solely associated with feasibility and are not part of the objective function. In particular, the PTV70 constraint is typically the hardest to satisfy in practice; IPMAN learns this difficulty and makes predictions that lie close to the boundary of the feasible set.

Table 1 shows performance on the held-out test set for IPMAN at iteration 11 against the baseline models. In general, IPMAN models with  $\lambda \leq 64$  satisfy the hidden constraints better than the baselines. IPMAN with  $\lambda = 4$  dominates all other models, including the baselines, in hidden and polyhedral constraint satisfaction, as well as objective function value. That is, this model predicts dose distributions that deliver lower dose to healthy tissue while better satisfying the clinical criteria. We conclude that training via IPMAN yields prediction models that produce feasible decisions more often and with a lower objective function value than existing state-of-the-art methods. We also observe that with higher values of  $\lambda$ , constraint satisfaction comes with a price; the objective function value is higher.

Recall that optimal solutions to the barrier problem satisfy a  $(\delta, \epsilon)$ -optimality guarantee. At high  $\lambda$ , this translates to a non-trivial upper bound with respect to the optimal value of the true problem, while at low  $\lambda$ , the guarantee translates to a non-trivial lower bound (see Theorem 1 and also EC.1). Because the generator learns via empirical risk minimization of the barrier problem, the corresponding predictions should also satisfy similar upper and lower bounds (see Theorem 2 and 3). While we may consequently expect setting a low  $\lambda$  to yield predictions that are infeasible (i.e., satisfying a non-trivial lower bound with respect to the optimal value), we find that  $\lambda = 4$ yields the best performance on out-of-sample data. In the next experiment, where the data is not perfectly indicative of the constraints, we observe the benefits of using a higher  $\lambda$  value to predict decisions that are more likely to be feasible.

### 6.4. Adapting to the clinical constraints of a different institution

The previous experiments were constructed using the clinical criteria from one institution under which the ground truth plans were developed. However, different clinics often use different criteria (Wu et al. 2017). Further, small clinics may have limited patient volume which may not be sufficient to properly train institution-specific models using existing methods (Boutilier et al. 2016). In this experiment, we show how IPMAN can be trained using the original data set to learn to predict feasible and optimal treatments for new clinical constraints.



(a) Objective function and fraction of feasible plans with respect to the hidden and polyhedral constraints.



(b) Average difference from the hidden constraint bound  $\hat{z}_r(\mathbf{u})$ . The dashed line is 0 Gy. Above 0 suggests plans satisfy the constraints on average.  $\lambda = 4$  is omitted to preserve scale.

Figure 4 Statistics on the validation set obtained during training on criteria from Geretschläger et al. (2015).

We use clinical criteria obtained from Geretschläger et al. (2015) who pursue a more aggressive treatment policy for head-and-neck cancer. They prescribe tumors to receive 72 Gy, 66 Gy, and 54 Gy to their three target sites, respectively, which we re-label in our data sets as PTV72, PTV66, and PTV54. Note that relative to the previous criteria in Section 6.3, two of the criteria have become stricter while the third is now easier to satisfy. Although we do not know the exact preferences of oncologists in determining when a criteria is necessary, we assume that any patient in our data set who was prescribed dose that satisfied the PTV hidden lower bound constraint from our institution would be prescribed dose at the corresponding new level at this new institution.

We train the generator and classifier using IPMAN for 11 iterations using the same settings as the previous section with one difference: we omit the  $l_1$  regularization term. While regularization can be useful to ensure that predictions are not vastly different from clinical data (see EC.3.3.4 for details), in this experiment, the clinical doses tend to be infeasible under the new criteria because they were generated using criteria from the original institution. For example, no plans in our data set received more than 72 Gy of dose to PTV70. Including the  $l_1$  term would, therefore,

ingingineu.							
Structure	Criteria (Gy)	Baseline	IPMAN $(\lambda)$				
		GAN	256	64	16	4	
Right Parotid	$z_o \leq 26$	83.3	85.7	85.7	100	100	
Left Parotid	$z_o \leq 26$	70.0	50.0	60.0	100	100	
Larynx	$z_o \leq 45$	93.3	86.7	76.8	100	100	
Mandible	$z_o^{\max} \leq 73.5$	100	81.0	90.5	100	100	
PTV72	$z_t^{90} \ge 72$	7.31	95.2	95.2	14.3	0	
PTV66	$z_t^{90} \ge 66$	77.8	96.3	96.3	85.2	0	
PTV54	$z_t^{90} \ge 54$	100	100	100	96.0	0	
All hidden con	straints	18.3	64.0	66.0	26.0	0	
All polyhedral constraints		93.8	94.0	90.0	88.0	0	
Objective function value		40.3	42.3	42.3	34.3	10.0	

Table 2The percentage of predicted decisions on the held-out test set that satisfy each hidden constraintof Geretschläger et al. (2015) to 1 Gy relaxation. The best performing models on the summary statistics arebighlighted

inappropriately guide the model to generate doses that tried to match the old criteria, rather than learn the new criteria.

Figure 4(a) displays the objective function value and the fraction of plans that satisfied the hidden and polyhedral constraints. The models trained for  $\lambda \leq 16$  decrease in objective function value and constraint satisfaction as the algorithm progresses. In the early stages, the classifier (which is initially trained mainly using the clinical data and doses sampled from the same distribution) has not yet observed a sufficient and diverse number of feasible plans. Therefore, the classifier is not yet a sufficient  $\delta$ -barrier, allowing the generator to leave the feasible set.

Generally at high  $\lambda$ , optimal solutions to the barrier problem are less aggressive in terms of minimizing the objective and instead lie well in the interior of the feasible set (see EC.1). As the classifier improves, particularly after iteration 6, the generative models for  $\lambda \geq 64$  quickly learn to predict solutions that are more likely to satisfy the hidden constraints. In particular, the distance from the PTV72 boundary in Figure 4(b) begins to increase from the 6th iteration and passes 0 (i.e., satisfy the hidden constraint) by the 10th iteration. This result demonstrates that our model is learning a new constraint, the PTV72 hidden criteria, which cannot be learned by naively using the training data. Recall that no clinical dose in the original data set reached 72 Gy on the PTV70. Thus, learning the PTV72 constraint is entirely attributable to the IPMAN procedure.

Table 2 shows the performance on the held-out test set for the generator at iteration 11. The baseline for comparison is the previous GAN which was trained on data from the original institution. Since no new data exists, the GAN cannot be re-trained to recognize the updated clinical criteria. As a result, few plans (18.3%) produced by the GAN satisfy the hidden constraints with only 77.8%

and 7.31% of the PTV66 and PTV72 criteria being satisfied, respectively. In contrast, IPMAN is able to learn the new hidden constraints. Overall, hidden constraint satisfaction is 64.0% with 96.3% and 95.2% of plans satisfy the PTV66 and PTV72, constraints, respectively with  $\lambda = 256$  or 64. Nevertheless, as a result of learning higher doses to the PTVs, constraint satisfaction on the OARs slightly degrades, which is noted in the slightly higher objective function values.

# 7. Conclusion

Conventional optimization techniques generally require well-structured problem formulations and make limited account of auxiliary data present in problems where different instances must be regularly solved. We propose Interior Point Methods with Adversarial Networks, a learning-based approach for generating solutions to optimization problems whose feasible sets are determined by instance-specific auxiliary information. We develop an unconstrained barrier problem where the barrier is replaced by a classifier trained on historical instances to predict feasibility. Because a classifier is not perfectly accurate, we extend the theory of interior point methods to the setting where only a relaxation of the feasible set is known and develop a corresponding optimality guarantee. Our main algorithm iteratively trains the classifier as well as a generative model via empirical risk minimization of the barrier problem. We demonstrate that the classifier learns to better approximate an effective barrier and the generative model learns to predict solutions with an optimality guarantee for both in-sample and out-of-sample instances. Ultimately, we obtain a deep learning model that can predict optimal solutions to problems in a fraction of the time that it would take a conventional optimization solver. Furthermore, our predictions account for instance-specific variations in the feasible set that conventional optimization would fail to permit.

To illustrate the application of our algorithm, we use it to predict dose distributions for radiation therapy as part of an automated planning pipeline. We find that our method learns to predict doses that better satisfy hidden clinical constraints and minimize objective function values as compared to state-of-the-art baseline learning methods. Furthermore, we show that our approach is adaptable in learning clinical criteria that are different from those that were used to generate the ground truth doses. This result suggests that an institution without a sufficient data set for training a dose prediction model could apply our methodology using data from another clinic; our approach would learn to produce appropriate doses tailored to the unique clinical criteria of the new institution while ensuring all solutions are certifiably optimal. As the global demand for radiation therapy grows and new clinics open in rural and developing areas, such adaptable automated planning methodologies have the potential to close the supply-demand gap in treatment planning capacity.

# References

- Angalakudati M, S B, Calzada J, Chatterjee B, Perakis G, Raad N, Uichanco J (2014) Business analytics for flexible resource allocation under random emergencies. *Management Science* 60(6):1552–1573.
- Arjovsky M, Bottou L (2017) Towards principled methods for training generative adversarial networks. arXivpreprint arXiv:1701.04862.
- Atun R, Jaffray DA, Barton MB, Bray F, Baumann M, Vikram B, Hanna TP, Knaul FM, Lievens Y, Lui TYM, Milosevic M, O'Sullivan B, Rodin DL, Rosenblatt E, Van Dyk J, Yap ML, Zubizarreta E, Gospodarowicz M (2015) Expanding global access to radiotherapy. *Lancet Oncology* 16(10):1153–86.
- Babier A, Boutilier JJ, McNiven AL, Chan TCY (2018a) Knowledge-based automated planning for oropharyngeal cancer. *Medical Physics* 45(7):2875–2883.
- Babier A, Boutilier JJ, Sharpe MB, McNiven AL, Chan TCY (2018b) Inverse optimization of objective function weights for treatment planning using clinical dose-volume histograms. *Physics in Medicine & Biology* 63(10):105004.
- Babier A, Mahmood R, McNiven AL, Diamant A, Chan TCY (2019) Knowledge-based automated planning with three-dimensional generative adversarial networks. *forthcoming in Medical Physics*.
- Badenbroek R, de Klerk E (2018) Complexity analysis of a sampling-based interior point method for convex optimization. *arXiv preprint arXiv:1811.07677*.
- Bain M, Sammut C (1999) A framework for behavioural cloning. Machine Intelligence 15, Intelligent Agents, 103–129.
- Ban GY, Rudin C (2018) The big data newsvendor: Practical insights from machine learning. *Operations Research* 1(67):90–108.
- Bartlett PL, Mendelson S (2002) Rademacher and gaussian complexities: Risk bounds and structural results. Journal of Machine Learning Research 3(Nov):463–482.
- Bello I, Pham H, Le QV, Norouzi M, Bengio S (2017) Neural combinatorial optimization.
- Bengio Y, Lodi A, Prouvost A (2018) Machine learning for combinatorial optimization: a methodological tour d'horizon. arXiv preprint arXiv:1811.06128.
- Benson HY, Shanno DF, Vanderbei RJ (2004) Interior-point methods for nonconvex nonlinear programming: jamming and numerical testing. *Mathematical Programming* 99(1):35–48.
- Bertsimas D, Kallus N (2019) From predictive to prescriptive analytics. Management Science 0(0).
- Bertsimas D, McCord C (2018) Optimization over continuous and multi-dimensional decisions with observational data. Advances in Neural Information Processing Systems, 2966–2974,.
- Boutilier JJ, Craig T, Sharpe MB, Chan TCY (2016) Sample size requirements for knowledge-based treatment planning. *Medical Physics* 43(3):1212–21.
- Boyd S, Vandenberghe L (2004) Convex Optimization (Cambridge University Press).

- Bubeck S, Eldan R (2019) The entropic barrier: Exponential families, log-concave geometry, and selfconcordance. *Mathematics of Operations Research* 44(1):264–276.
- Dai H, Khalil EB, Zhang Y, Dilkina B, Song L (2017) Learning combinatorial optimization algorithms over graphs. Advances in Neural Information Processing Systems,, 6348–6358.
- Delaney G, Jacob S, Featherstone C, Barton M (2005) The role of radiotherapy in cancer treatment. *Cancer* 104(6):1129–1137.
- Donti P, Amos B, Kolter JZ (2017) Task-based end-to-end model learning in stochastic optimization. Advances in Neural Information Processing Systems, 5484–5494.
- Elmachtoub AN, Grigas P (2017) Smart "predict, then optimize". arXiv preprint arXiv:1710.08005.
- Engelking R (1977) General Topology (Polish Scientific Publishers).
- Ferreira KJ, Lee BHA, Simchi-Levi D (2015) Analytics for an online retailer: Demand forecasting and price optimization. Manufacturing & Service Operations Management 18(1):69–88.
- Foster DJ, Sekhari A, Sridharan K (2018) Uniform convergence of gradients for non-convex learning and optimization. Advances in Neural Information Processing Systems, 8759–8770.
- Geretschläger A, Bojaxhiu B, Dal Pra A, Leiser D, Schmücking M, Arnold A, Ghadjar P, Aebersold DM (2015) Definitive intensity modulated radiotherapy in locally advanced hypopharygeal and laryngeal squamous cell carcinoma: mature treatment results and patterns of locoregional failure. *Radiation* Oncolology 10:20.
- Gondzio J (2012) Interior point methods 25 years later. European Journal of Operational Research 218(3):587–601.
- Goodfellow I, Bengio Y, Courville A (2016) Deep Learning, volume 1 (MIT press Cambridge).
- Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y (2014) Generative adversarial nets. Advances in Neural Information Processing Systems, 2672–2680.
- Hannah L, Powell W, Blei DM (2010) Nonparametric density estimation for stochastic optimization with an observable state variable. Advances in Neural Information Processing Systems, 820–828.
- Hinder O, Ye Y (2018) A one-phase interior point method for nonconvex optimization.  $arXiv \ preprint$ arXiv:1801.03072.
- Hopfield JJ, Tank DW (1985) "neural" computation of decisions in optimization problems. *Biological Cyber*netics 52(3):141–152.
- Hornik K (1991) Approximation capabilities of multilayer feedforward networks. *Neural Networks* 4(2):251–257.
- Isola P, Zhu JY, Zhou T, Efros AA (2017) Image-to-image translation with conditional adversarial networks. The IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

- Kao YH, Roy BV, Yan X (2009) Directed regression. Advances in Neural Information Processing Systems, 889–897.
- Konda VR, Tsitsiklis JN (2000) Actor-critic algorithms. Advances in Neural Information Processing Systems, 1008–1014.
- Larsen E, Lachapelle S, Bengio Y, Frejinger E, Lacoste-Julien S, Lodi A (2018) Predicting solution summaries to integer linear programs under imperfect information with machine learning. arXiv preprint arxiv:1807.11876.
- Liu S, He L, Shen ZJ (2018) Data-driven order assignment for last mile delivery. Available at SSRN 3179994
- Low DA, Harms WB, Mutic S, Purdy JA (1998) A technique for the quantitative evaluation of dose distributions. *Medical Physics* 25(5):656–661.
- Mahmood R, Babier A, McNiven A, Diamant A, Chan TCY (2018) Automated treatment planning in radiation therapy using generative adversarial networks. of Machine Learning Research P, ed., Machine Learning for Health Care, volume 85.
- Maurer A (2016) A vector-contraction inequality for rademacher complexities. International Conference on Algorithmic Learning Theory, 3–17 (Springer).
- McIntosh C, Purdie TG (2017) Voxel-based dose prediction with multi-patient atlas selection for automated radiotherapy treatment planning. *Physics in Medicine & Biology* 62(2):415–431.
- McIntosh C, Welch M, McNiven A, Jaffray DA, Purdie TG (2017) Fully automated treatment planning for head and neck radiotherapy using a voxel-based dose prediction and dose mimicking method. *Physics* in Medicine & Biology 62(15):5926–5944.
- Mišić VV (2019) Optimization of tree ensembles. forthcoming in Operations Research.
- Nesterov Y, Nemirovskii A (1994) Interior-point polynomial algorithms in convex programming, volume 13 (SIAM).
- Neyshabur B, Tomioka R, Srebro N (2015) Norm-based capacity control in neural networks. *Conference on Learning Theory*, 1376–1401.
- Nguyen D, Jia X, Sher D, Lin MH, Iqbal Z, Liu H, Jiang S (2019) 3d radiotherapy dose prediction on head and neck cancer patients with a hierarchically densely connected u-net deep learning architecture. *Physics in Medicine & Biology* 64(6):065020.
- Pelikan M, Goldberg DE, Lobo FG (2002) A survey of optimization by building and using probabilistic models. Computational Optimization and Applications 21(1):5–20.
- Shiraishi S, Tan J, Olsen LA, Moore KL (2015) Knowledge-based prediction of plan quality metrics in intracranial stereotactic radiosurgery. *Medical Physics* 42(2):908.
- Vanderbei RJ, Shanno DF (1999) An interior-point algorithm for nonconvex nonlinear programming. Computational Optimization and Applications 13(1-3):231–252.

- Vinyals O, Fortunato M, Jaitly N (2015) Pointer networks. Advances in Neural Information Processing Systems, 2692–2700.
- Wu B, Kusters M, Kunze-Busch M, Dijkema T, McNutt T, Sanguineti G, Bzdusek K, Dritschilo A, Pang D (2017) Cross-institutional knowledge-based planning (KBP) implementation and its performance comparison to auto-planning engine (APE). Journal of the European Society for Therapeutic Radiology and Oncology 123(1):57–62.

This page is intentionally blank. Proper e-companion title page, with INFORMS branding and exact metadata of the main paper, will be produced by the INFORMS office when the issue is being assembled.

# **Electronic Companion**

# EC.1. Structural properties of $(\delta, \epsilon)$ -optimality for the barrier problem

The IPMAN algorithm simultaneously trains a classifier and a generative model to learn feasibility and predictive optimal solutions respectively. Alternatively, if we are already given a  $\delta$ -barrier  $B_{\delta}(\mathbf{x}, \mathbf{u})$ , we may consider directly optimizing  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda)$ . In this section, we show how tuning the  $\lambda$  parameter can yield feasible or infeasible solutions of different qualities.

Under a mild regularity assumption, for a sufficiently large  $\lambda$ , an optimal solution  $\mathbf{x}^{\lambda}(\mathbf{u})$  to  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda)$  is guaranteed to lie inside  $\mathcal{X}(\mathbf{u})$ . Once  $\lambda$  is sufficiently small, the optimal solutions then enter  $\mathcal{N}_{\delta}(\mathcal{X}(\mathbf{u})) \setminus \mathcal{X}(\mathbf{u})$ . We first state this assumption before characterizing the trajectory of the sequence of points obtained via an IPM.

Assumption EC.1 (Regularity of the  $\delta$ -barrier).

- 1. There exist  $\tilde{\mathbf{x}} \in \operatorname{int}(\mathcal{X}(\mathbf{u}))$  such that  $B_{\delta}(\tilde{\mathbf{x}}, \mathbf{u}) > B_{\delta}(\mathbf{x}, \mathbf{u})$  for all  $\mathbf{x} \in \operatorname{cl}(\mathcal{N}_{\delta}(\mathcal{X}(\mathbf{u})) \setminus \mathcal{X}(\mathbf{u}))$ .
- 2. There exist  $\tilde{\mathbf{x}}' \in \mathcal{N}_{\delta}(\mathcal{X}(\mathbf{u})) \setminus \mathcal{X}(\mathbf{u})$  such that  $f(\tilde{\mathbf{x}}') < f(\mathbf{x}^*(\mathbf{u}))$  and  $0 < B_{\delta}(\tilde{\mathbf{x}}', \mathbf{u}) < B_{\delta}(\mathbf{x}, \mathbf{u})$  for all  $\mathbf{x} \in \mathcal{X}(\mathbf{u})$ .

The first statement implies that there exists a point inside  $\mathcal{X}(\mathbf{u})$  for which  $B_{\delta}(\mathbf{x}, \mathbf{u})$  is greater than any point outside of  $\mathcal{X}(\mathbf{u})$ . Similarly, the second statement implies that there exists a point outside of  $\mathcal{X}(\mathbf{u})$  for which  $B_{\delta}(\mathbf{x}, \mathbf{u})$  is lower than any point inside  $\mathcal{X}(\mathbf{u})$ . Intuitively, the barrier yields higher values for points inside  $\mathcal{X}(\mathbf{u})$  rather than outside. Furthermore, the existence of  $\tilde{\mathbf{x}}'$  for which  $f(\tilde{\mathbf{x}}) > f(\mathbf{x}^*(\mathbf{u})) > f(\tilde{\mathbf{x}}')$  is a direct consequence of the linear objective. Figure EC.1 shows an example of such points for a feasible set where the  $\delta$ -barrier is a canonical barrier for  $\mathcal{P}$ . Given a barrier function satisfying Assumption EC.1,  $\lambda$  controls the feasibility of  $\mathbf{x}^{\lambda}(\mathbf{u})$  for **OP**( $\mathbf{u}$ ).

LEMMA EC.1. If Assumption EC.1 is satisfied, then there exists  $\tilde{\lambda}$  such that for all  $\lambda \geq \tilde{\lambda}$ , the optimal solution to  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda)$  is feasible for  $\mathbf{OP}(\mathbf{u})$ , i.e.,  $\mathbf{x}^{\lambda}(\mathbf{u}) \in \mathcal{X}(\mathbf{u})$ .

Proof of Lemma EC.1. Let  $\mathbf{x}^+ \in \arg \sup_{\mathbf{x}} \{B_{\delta}(\mathbf{x}, \mathbf{u}) \mid \mathbf{x} \in \mathcal{N}_{\delta}(\mathcal{X}(\mathbf{u})) \setminus \mathcal{X}(\mathbf{u})\}$  and  $\mathbf{x}^- \in \arg \inf_{\mathbf{x}} \{f(\mathbf{x}) \mid B_{\delta}(\mathbf{x}, \mathbf{u}) > 0\}$ . Then, for  $\tilde{\mathbf{x}}$  satisfying Assumption EC.1 Statement 1, we set

$$\tilde{\lambda} = \frac{f(\tilde{\mathbf{x}}) - f(\mathbf{x}^{-})}{\log B_{\delta}(\tilde{\mathbf{x}}, \mathbf{u}) - \log B_{\delta}(\mathbf{x}^{+}, \mathbf{u})}.$$
(EC.1)

From the optimality of  $\mathbf{x}^-$ , we have  $f(\tilde{\mathbf{x}}) > f(\mathbf{x}^-)$ . Then, Assumption EC.1 implies that the denominator is positive, and therefore  $\tilde{\lambda} > 0$ . Rearranging (EC.1) yields

$$f(\tilde{\mathbf{x}}) - \tilde{\lambda} \log B_{\delta}(\tilde{\mathbf{x}}, \mathbf{u}) = f(\mathbf{x}^{-}) - \tilde{\lambda} \log B_{\delta}(\mathbf{x}^{+}, \mathbf{u}).$$

By optimality of  $\mathbf{x}^+$  and  $\mathbf{x}^-$ , we have  $f(\mathbf{x}) \ge f(\mathbf{x}^-)$  and  $\log B_{\delta}(\mathbf{x}, \mathbf{u}) \le \log B_{\delta}(\mathbf{x}^+, \mathbf{u})$  respectively, for all  $\mathbf{x} \in \mathcal{N}_{\delta}(\mathcal{X}(\mathbf{u})) \setminus \mathcal{X}(\mathbf{u})$ . Therefore,  $f(\tilde{\mathbf{x}}) - \tilde{\lambda} \log B_{\delta}(\tilde{\mathbf{x}}, \mathbf{u}) \le f(\mathbf{x}) - \tilde{\lambda} \log B_{\delta}(\mathbf{x}, \mathbf{u})$  for all  $\mathbf{x} \in \mathcal{N}_{\delta}(\mathcal{X}(\mathbf{u})) \setminus \mathcal{X}(\mathbf{u})$ , concluding that the optimal solution to  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \tilde{\lambda})$  must satisfy  $\mathbf{x}^{\tilde{\lambda}}(\mathbf{u}) \in \mathcal{X}(\mathbf{u})$ . Now for any  $\varepsilon > 0$ , observe that

$$\begin{aligned} f(\tilde{\mathbf{x}}) - (\tilde{\lambda} + \varepsilon) \log B_{\delta}(\tilde{\mathbf{x}}, \mathbf{u}) &\leq f(\mathbf{x}) - \tilde{\lambda} \log B_{\delta}(\mathbf{x}, \mathbf{u}) - \varepsilon \log B_{\delta}(\tilde{\mathbf{x}}, \mathbf{u}), \qquad \forall \mathbf{x} \in \mathcal{N}_{\delta}\left(\mathcal{X}(\mathbf{u})\right) \setminus \mathcal{X}(\mathbf{u}) \\ &< f(\mathbf{x}) - \tilde{\lambda} \log B_{\delta}(\mathbf{x}, \mathbf{u}) - \varepsilon \log B_{\delta}(\mathbf{x}, \mathbf{u}), \qquad \forall \mathbf{x} \in \mathcal{N}_{\delta}\left(\mathcal{X}(\mathbf{u})\right) \setminus \mathcal{X}(\mathbf{u}). \end{aligned}$$

The first line is obtained by adding  $\varepsilon \log B_{\delta}(\tilde{\mathbf{x}}, \mathbf{u})$  to both sides, and the second from  $B_{\delta}(\tilde{\mathbf{x}}, \mathbf{u}) > B_{\delta}(\mathbf{x}, \mathbf{u})$  for  $\mathbf{x} \in \mathcal{N}_{\delta}(\mathcal{X}(\mathbf{u})) \setminus \mathcal{X}(\mathbf{u})$ . Thus,  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \tilde{\lambda} + \varepsilon)$  yields feasible solutions to  $\mathbf{OP}(\mathbf{u})$ .  $\Box$ 

LEMMA EC.2. If Assumption EC.1 is satisfied, then there exists  $\tilde{\lambda}'$  such that for all  $\lambda \leq \tilde{\lambda}'$ , the optimal solution to  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda)$  is infeasible for  $\mathbf{OP}(\mathbf{u})$ , i.e.,  $\mathbf{x}^{\lambda}(\mathbf{u}) \in \mathcal{N}_{\delta}(\mathcal{X}(\mathbf{u})) \setminus \mathcal{X}(\mathbf{u})$ .

Proof of Lemma EC.2. Let  $\mathbf{x}^{\dagger} \in \operatorname{arg\,max}_{\mathbf{x}} \{B_{\delta}(\mathbf{x}, \mathbf{u}) \mid \mathbf{x} \in \mathcal{X}(\mathbf{u})\}$ . Then, for  $\tilde{\mathbf{x}}'$  satisfying Assumption EC.1 Statement 2, let

$$\tilde{\lambda}' = \frac{f\left(\mathbf{x}^{*}(\mathbf{u})\right) - f(\tilde{\mathbf{x}}')}{\log B_{\delta}(\mathbf{x}^{\dagger}, \mathbf{u}) - \log B_{\delta}(\tilde{\mathbf{x}}', \mathbf{u})}.$$
(EC.2)

Assumption EC.1 Statement 2 ensures  $f(\mathbf{x}^*(\mathbf{u})) > f(\tilde{\mathbf{x}}')$  and  $\log B_{\delta}(\mathbf{x}^{\dagger}, \mathbf{u}) > \log B_{\delta}(\tilde{\mathbf{x}}', \mathbf{u})$ . Therefore,  $\tilde{\lambda}' > 0$ . Rearranging (EC.2) gives us

$$f(\tilde{\mathbf{x}}') - \tilde{\lambda}' \log B_{\delta}(\tilde{\mathbf{x}}', \mathbf{u}) = f(\mathbf{x}^*(\mathbf{u})) - \tilde{\lambda}' \log B_{\delta}(\mathbf{x}^{\dagger}, \mathbf{u}).$$

By optimality of  $\mathbf{x}^*(\mathbf{u})$  and  $\mathbf{x}^\dagger$ , we have  $f(\mathbf{x}) \ge f(\mathbf{x}^*(\mathbf{u}))$  and  $\log B_{\delta}(\mathbf{x}, \mathbf{u}) \le \log B_{\delta}(\mathbf{x}^\dagger, \mathbf{u})$  respectively, for all  $\mathbf{x} \in \mathcal{X}(\mathbf{u})$ . Therefore  $f(\tilde{\mathbf{x}}') - \tilde{\lambda}' \log B_{\delta}(\tilde{\mathbf{x}}', \mathbf{u}) \le f(\mathbf{x}) - \tilde{\lambda}' \log B_{\delta}(\mathbf{x}, \mathbf{u})$  for all  $\mathbf{x} \in \mathcal{X}(\mathbf{u})$ , concluding that the optimal solution to  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \tilde{\lambda}')$  must satisfy  $\mathbf{x}^{\tilde{\lambda}'}(\mathbf{u}) \in \mathcal{N}_{\delta}(\mathcal{X}(\mathbf{u})) \setminus \mathcal{X}(\mathbf{u})$ .

Now for any  $\varepsilon > 0$ , observe that

$$\begin{aligned} f(\tilde{\mathbf{x}}') - (\tilde{\lambda}' - \varepsilon) \log B_{\delta}(\tilde{\mathbf{x}}', \mathbf{u}) &\leq f(\mathbf{x}) - \tilde{\lambda}' \log B_{\delta}(\mathbf{x}, \mathbf{u}) + \varepsilon \log B_{\delta}(\tilde{\mathbf{x}}', \mathbf{u}), & \forall \mathbf{x} \in \mathcal{X}(\mathbf{u}) \\ &< f(\mathbf{x}) - \tilde{\lambda}' \log B_{\delta}(\mathbf{x}, \mathbf{u}) + \varepsilon \log B_{\delta}(\mathbf{x}, \mathbf{u}), & \forall \mathbf{x} \in \mathcal{X}(\mathbf{u}). \end{aligned}$$

The first line is obtained by subtracting  $\varepsilon \log B_{\delta}(\tilde{\mathbf{x}}', \mathbf{u})$  to both sides, and the second from  $B_{\delta}(\tilde{\mathbf{x}}', \mathbf{u}) < B_{\delta}(\mathbf{x}, \mathbf{u})$  for all  $\mathbf{x} \in \mathcal{X}(\mathbf{u})$ . Thus,  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \tilde{\lambda}' - \varepsilon)$  yields infeasible solutions to  $\mathbf{OP}(\mathbf{u})$ .

Lemma EC.2 and the second statement of Assumption EC.1 give the condition where the barrier problem produces undesirable results. Otherwise, if  $f(\tilde{\mathbf{x}}') > f(\mathbf{x}^*(\mathbf{u}))$  and  $B_{\delta}(\tilde{\mathbf{x}}', \mathbf{u}) \ge B_{\delta}(\mathbf{x}, \mathbf{u})$  for all  $\tilde{\mathbf{x}}' \in \mathcal{N}_{\delta}(\mathcal{X}(\mathbf{u})) \setminus \mathcal{X}(\mathbf{u})$  and  $\mathbf{x} \in \mathcal{X}(\mathbf{u})$ , **OP**( $\mathbf{u}$ ) could be solved by classical IPMs.

Lemmas EC.1 and EC.2 state that when  $\lambda$  is set sufficiently high (or low), the corresponding optimal solution  $\mathbf{x}^{\lambda}(\mathbf{u})$  is a certifiably feasible (or infeasible) solution to  $\mathbf{OP}(\mathbf{u})$ . Furthermore, there exists a trajectory, i.e., feasibility (or infeasibility) is guaranteed for all  $\lambda$  sufficiently high (or low). Assuming access to an oracle  $\Psi(\mathbf{x}, \mathbf{u})$ , we conconstruct a simple IPM to obtain optimal solutions to  $\mathbf{OP}(\mathbf{u})$ . We initialize with a large  $\lambda_0$  that satisfies Lemma EC.1. We define a decay



Figure EC.1 The dotted lines are level sets.  $\mathbf{x}^*(\mathbf{u})$  is optimal for  $OP(\mathbf{u})$  while  $\tilde{\mathbf{x}}$  and  $\tilde{\mathbf{x}}'$  satisfy Lemmas EC.1 and EC.2 respectively.

rate  $\nu < 1$  and a number of iterations  $j \in 0, ..., M$ . Then, for each j, we simply let  $\lambda_j = \lambda_0 \nu^j$  and solve **BP**( $\mathbf{u}, B_{\delta}, \lambda_j$ ) to obtain a new ( $\delta, \epsilon$ )-optimal solution in each iteration. At the end of each iteration, the oracle checks if the solution is still feasible, and terminates when the solution exits the feasible set. We prove the structure of this approach below.

PROPOSITION EC.1. Suppose that  $\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2 \in \mathcal{X}(\mathbf{u})$  and  $\tilde{\mathbf{x}}'_1, \tilde{\mathbf{x}}'_2 \in \mathcal{N}_{\delta}(\mathcal{X}(\mathbf{u})) \setminus \mathcal{X}(\mathbf{u})$  satisfy Statements 1 and 2 of Assumption EC.1, respectively. Assume without loss of generality  $B_{\delta}(\tilde{\mathbf{x}}_1, \mathbf{u}) > B_{\delta}(\tilde{\mathbf{x}}_2, \mathbf{u})$  and  $f(\tilde{\mathbf{x}}'_1) > f(\tilde{\mathbf{x}}'_2)$ . Let  $\mathbf{x}^{\mathcal{P}} \in \operatorname{arg\,min}_{\mathbf{x}} \{f(\mathbf{x}) \mid \mathbf{x} \in \mathcal{P}\}$ . For M > 0 and  $j \in \{0, \ldots, M\}$ , consider

$$\lambda_0 = \frac{f(\tilde{\mathbf{x}}_1) - f(\mathbf{x}^{\mathcal{P}})}{\log B_{\delta}(\tilde{\mathbf{x}}_1, \mathbf{u}) - \log B_{\delta}(\tilde{\mathbf{x}}_2, \mathbf{u})}, \quad \nu = \left(\frac{f(\tilde{\mathbf{x}}_1') - f(\tilde{\mathbf{x}}_2')}{-\lambda_0 \log B_{\delta}(\tilde{\mathbf{x}}_1', \mathbf{u})}\right)^{1/M}, \quad \lambda_j = \lambda_0 \nu^{j_j}$$

Then, the following statements are true:

- 1. An optimal solution  $\mathbf{x}^{\lambda_0}(\mathbf{u})$  to  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda_0)$  is a feasible solution for  $\mathbf{OP}(\mathbf{u})$ .
- 2. There exists  $1 \leq j^* \leq M$  such that for all  $j < j^*$ , an optimal solution  $\mathbf{x}^{\lambda_j}(\mathbf{u})$  to  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda_j)$  is feasible for  $\mathbf{OP}(\mathbf{u})$  and for all  $j \geq j^*$ ,  $\mathbf{x}^{\lambda_j}(\mathbf{u})$  is infeasible for  $\mathbf{OP}(\mathbf{u})$ .
- 3. For any  $j < j^*$ , an optimal solution  $\mathbf{x}^{\lambda_j}(\mathbf{u})$  is  $(0, \epsilon_j)$ -optimal for  $\mathbf{OP}(\mathbf{u})$  where

$$\epsilon_j = (f(\tilde{\mathbf{x}}_1') - f(\tilde{\mathbf{x}}_2'))\nu^{j-M}$$

Furthermore for any  $j \ge j^*$ ,  $\mathbf{x}^{\lambda_j}(\mathbf{u})$  is  $(\Delta(\mathbf{u}), \epsilon_j)$ -optimal for  $\mathbf{OP}(\mathbf{u})$ .

Proof of Proposition EC.1. We first make several observations about the parameters. Note that because  $\mathcal{X}(\mathbf{u}) \subset \mathcal{P}$  relaxes the feasible set, we have  $f(\mathbf{x}^{\mathcal{P}}) \leq f(\mathbf{x}^*(\mathbf{u}))$ . Next for all  $j \leq M$ ,  $\lambda_j = \lambda_0 \nu^j$ and specifically  $\lambda_M = \lambda_0 \nu^M = -(f(\tilde{\mathbf{x}}_1') - f(\tilde{\mathbf{x}}_2')) / \log B_{\delta}(\tilde{\mathbf{x}}_1', \mathbf{u})$ . To prove the first statement, we show that  $\lambda_0 > \tilde{\lambda}$  where  $\tilde{\lambda}$  is defined as in (EC.1) and constructed using  $\tilde{\mathbf{x}}_1$ . Note that  $f(\mathbf{x}^{\mathcal{P}}) \leq f(\mathbf{x}^-)$  and by Assumption EC.1,  $\log B_{\delta}(\tilde{\mathbf{x}}_2, \mathbf{u}) > \log B_{\delta}(\mathbf{x}^+, \mathbf{u})$ . We substitute  $f(\mathbf{x}^{\mathcal{P}})$  and  $\log B_{\delta}(\tilde{\mathbf{x}}_2, \mathbf{u})$  in  $\lambda_0$  and prove  $\lambda_0 > \tilde{\lambda}$ . By Lemma EC.1, Statement 1 must hold.

We use a similar argument to show  $\lambda_M < \tilde{\lambda}'$  as defined in (EC.2) using  $\tilde{\mathbf{x}}'_1$ . By Lemma EC.2, an optimal solution  $\mathbf{x}^{\lambda_M}$  must be infeasible for  $\mathbf{OP}(\mathbf{u})$ . Given that  $\lambda_j$  decreases every iteration and using the first statement, there must exist a cutoff point  $1 \leq j^* \leq M$  for which  $\lambda_{j^*} < \tilde{\lambda}'$  and  $\lambda_{j^*-1} \geq \tilde{\lambda}'$ . Therefore, Statement 2 must also hold.

In order to prove the third statement, recall that we assume  $\delta \leq \Delta(\mathbf{u})$  for all j. We first prove  $(\Delta(\mathbf{u}), \epsilon_j)$ -optimality when j = M, and then prove for j < M. Let  $\epsilon_M = f(\tilde{\mathbf{x}}'_1) - f(\tilde{\mathbf{x}}'_2)$ . Note that

$$\lambda_M = \frac{f(\tilde{\mathbf{x}}_1') - f(\tilde{\mathbf{x}}_2')}{-\log B_{\delta}(\tilde{\mathbf{x}}_1', \mathbf{u})} = \frac{\epsilon_M}{-\log B_{\delta}(\tilde{\mathbf{x}}_1', \mathbf{u})} < \frac{\epsilon_M}{-\log B_{\delta}(\mathbf{x}^*, \mathbf{u})}.$$

The second equality follows from substituting the value of  $\epsilon_M$  and the inequality from  $B_{\delta}(\tilde{\mathbf{x}}'_1, \mathbf{u}) < B_{\delta}(\mathbf{x}^*(\mathbf{u}), \mathbf{u})$  (i.e., Assumption EC.1). We next show that  $\mathbf{x}^{\lambda_M}$  satisfies  $(\Delta(\mathbf{u}), \epsilon_M)$ -optimality,

$$\begin{split} f\left(\mathbf{x}^{*}(\mathbf{u})\right) + \epsilon_{M} &> f\left(\mathbf{x}^{*}(\mathbf{u})\right) - \lambda_{M} \log B_{\delta}\left(\mathbf{x}^{*}(\mathbf{u}), \mathbf{u}\right) \\ &\geq f(\mathbf{x}^{\lambda_{M}}(\mathbf{u})) - \lambda_{M} \log B_{\delta}\left(\mathbf{x}^{\lambda_{M}}(\mathbf{u}), \mathbf{u}\right) \\ &> f\left(\mathbf{x}^{\lambda_{M}}(\mathbf{u})\right). \end{split}$$

The first line follows from substituting the value of  $\epsilon_M$  and the second from the optimality of  $\mathbf{x}^{\lambda_M}(\mathbf{u})$  for  $\mathbf{BP}(\mathbf{u}, B_{\delta}, \lambda_M)$ . The third line follows from the fact that  $-\lambda_M \log B_{\delta}(\mathbf{x}^{\lambda_M}(\mathbf{u}), \mathbf{u}) > 0$ .

For each j < M, we have  $\lambda_j = \lambda_M \nu^{j-M}$ . Then, we write  $\epsilon_j = (f(\tilde{\mathbf{x}}'_1) - f(\tilde{\mathbf{x}}'_2))\nu^{j-M}$ . The same steps used for the j = M case are repeated to obtain  $(\Delta(\mathbf{u}), \epsilon_j)$ -optimality certificates. Finally, note that from Statement 2, for all  $j < j^*$ , the optimal solutions  $\mathbf{x}^{\lambda_j}(\mathbf{u})$  are feasible for  $\mathbf{OP}(\mathbf{u})$ . By optimality of  $\mathbf{x}^*(\mathbf{u})$  for  $\mathbf{OP}(\mathbf{u})$ , we have  $\delta = 0$  for all  $j < j^*$ .

Proposition EC.1 first provides parameters  $\lambda_0 > \tilde{\lambda}$  and  $\lambda_M < \tilde{\lambda}'$  for which the optimal solutions to **BP**( $\mathbf{u}, B_{\delta}, \lambda_0$ ) and **BP**( $\mathbf{u}, B_{\delta}, \lambda_M$ ) lie inside and outside of  $\mathcal{X}(\mathbf{u})$ , respectively. Next, it shows that the sequence of  $\lambda_j$  produces a sequence of optimal solutions { $\mathbf{x}^{\lambda_j}(\mathbf{u})$ } that start within the feasible set  $\mathcal{X}(\mathbf{u})$  and proceed to move outside. Finally, it derives a sequence of corresponding { $\epsilon_j$ } such that the sequence of solutions are ( $\Delta(\mathbf{u}), \epsilon_j$ )-optimal for **OP**( $\mathbf{u}$ ). This implies the final solution is  $(0, (f(\tilde{\mathbf{x}}'_1) - f(\tilde{\mathbf{x}}'_2)\nu^{j^*-1-M}))$ -optimal for **OP**( $\mathbf{u}$ ).

The above proposition summarizes an IPM for solving  $OP(\mathbf{u})$  when given a  $\delta$ -barrier and  $\Psi(\mathbf{x}, \mathbf{u})$ . The IPM behaves in a desirable and predictable fashion; by initializing with large  $\lambda$ , we ensure that we obtain feasible solutions, but by decreasing  $\lambda$ , we know that the solution will ultimately be infeasible. An oracle could identify the point of termination immediately before the IPM leaves the feasible set. We can from here obtain a tight bound on the  $(\delta, \epsilon)$ -optimality of the final solution. While direct optimization is desirable for its structural properties, this IPM approach is reliant on access to a  $\delta$ -barrier. On the other hand, IPMAN learns a classifier that approximates a  $\delta$ -barrier after several iterations. Therefore, unless we are given an a priori  $\delta$ -barrier (e.g., a canonical barrier for  $\mathcal{P}$ ), this IPM approach is not necessarily feasible from the onset. A potential fix would be to first train IPMAN until a  $\delta$ -barrier is obtained and then use the  $\delta$ -barrier IPM to solve subsequent problems. This ties to the second difference between the two approaches; IPMAN is ultimately a predictive model and is therefore subject to prediction error. On the other hand, prediction from a trained model is much faster than direct optimization. Therefore, in cases where the problem is large and an IPM would be difficult to solve or require numerous queries from an oracle, the predictive power of IPMAN yields more practical benefits.

# EC.2. Proof of the generalization bound (Theorem 3)

The proof of the generalization bound uses a Generalization Lemma of Bertsimas and Kallus (2019) to bound the error in objective function value of  $F^*(\mathbf{u})$  versus  $\mathbf{x}^{\lambda}(\mathbf{u})$  and Markov's inequality to translate this bound to a probabilistic  $(\delta, \epsilon)$ -optimality certificate. However, in order to use the lemma in this way, we first require an auxiliary result to relate  $F^*$  with  $\mathfrak{R}_{N_{\mathbf{u}}}(\mathcal{F})$ .

Assumption 2 states that the generative model  $F^*$  is a composition; we project the optimal solution  $F^{(k)}$  of  $\mathbf{GBP}(\hat{\mathcal{U}}, B_{\delta}, \lambda)$  to  $\mathcal{P}$  whenever  $F^{(k)}(\mathbf{u}) \notin \mathcal{P}$ . Although  $F^{(k)} \in \mathcal{F}$ , the final model  $F^*(\mathbf{u}) := \operatorname{proj}(F(\mathbf{u})) = \operatorname{arg\,min}_{\mathbf{x}} \{ \|\mathbf{x} - F(\mathbf{u})\| \mid \mathbf{x} \in \mathcal{P} \}$  is not a member of  $\mathcal{F}$ . We first bound the Rademacher complexity of models composed from projection below.

LEMMA EC.3. Let  $\mathcal{F} = \{F : \mathcal{U} \to \mathbb{R}^n\}$  be a model class and  $\operatorname{proj}(\mathcal{F}) = \{\operatorname{proj}(F) \mid F \in \mathcal{F}\}$ be the class of models composed by a projection to a polyhedron  $\mathcal{P}$ . Then for any  $\hat{\mathcal{U}} \sim \mathbb{P}_{\mathbf{u}}$ ,  $\hat{\mathfrak{R}}_{N_{\mathbf{u}}}(\operatorname{proj}(\mathcal{F}), \hat{\mathcal{U}}) \leq \sqrt{2n} \hat{\mathfrak{R}}_{N_{\mathbf{u}}}(\mathcal{F}, \hat{\mathcal{U}}).$ 

Proof of Lemma EC.3. We want to show for any fixed  $\hat{\mathcal{U}}$  that

$$\mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}} \left[ \frac{2}{N_{\mathbf{u}}} \sup_{F \in \mathcal{F}} \sum_{i=1}^{N_{\mathbf{u}}} \boldsymbol{\sigma}_{i}^{\mathsf{T}} \operatorname{proj}(F(\hat{\mathbf{u}}_{i})) \right] \leq \sqrt{2n} \mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}} \left[ \frac{2}{N_{\mathbf{u}}} \sup_{F \in \mathcal{F}} \sum_{i=1}^{N_{\mathbf{u}}} \boldsymbol{\sigma}_{i}^{\mathsf{T}} F(\hat{\mathbf{u}}_{i}) \right].$$
(EC.3)

By conditioning and iterating, it suffices to prove the following inequality for any function  $\Xi(F)$ :  $\mathcal{F} \to \mathbb{R}$ ,

$$\mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}} \left[ \sup_{F \in \mathcal{F}} \boldsymbol{\sigma}^{\mathsf{T}} \operatorname{proj}(F) + \Xi(F) \right] \leq \mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}} \left[ \sup_{F \in \mathcal{F}} \sqrt{2n} \boldsymbol{\sigma}^{\mathsf{T}} F + \Xi(F) \right].$$
(EC.4)

We first prove inequality (EC.4), before returning to the main lemma.

As  $\sigma \sim p_{\sigma}$  is a random vector of i.i.d. Rademacher variables, it is supported over the (ordered) set  $\{(-1, \ldots, -1, -1), (-1, \ldots, -1, 1), \ldots, (1, \ldots, 1, 1)\}$  all with equal probability. Let  $\hat{\sigma}_{\ell}$  denote the

 $\ell$ -th element of this set. By iterating over all values, we expand the left-hand-side of (EC.4) out to:

$$\mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}} \left[ \sup_{F \in \mathcal{F}} \boldsymbol{\sigma}^{\mathsf{T}} \operatorname{proj}(F) + \Xi(F) \right] = \frac{1}{2^n} \sum_{\ell=1}^{2^n} \left( \sup_{F \in \mathcal{F}} \hat{\boldsymbol{\sigma}}_{\ell}^{\mathsf{T}} \operatorname{proj}(F) + \Xi(F) \right)$$
(EC.5)

$$= \frac{1}{2^n} \sum_{\ell=1}^{2^{n-1}} \left( \sup_{F \in \mathcal{F}} \left\{ \hat{\boldsymbol{\sigma}}_{\ell}^{\mathsf{T}} \operatorname{proj}(F) + \Xi(F) \right\} + \sup_{F \in \mathcal{F}} \left\{ -\hat{\boldsymbol{\sigma}}_{\ell}^{\mathsf{T}} \operatorname{proj}(F) + \Xi(F) \right\} \right)$$
(EC.6)

$$= \frac{1}{2^n} \sum_{\ell=1}^{2^{n-1}} \left( \sup_{F_1, F_2 \in \mathcal{F}} \hat{\sigma}_{\ell}^{\mathsf{T}} \left( \operatorname{proj}(F_1) - \operatorname{proj}(F_2) \right) + \Xi(F_1) + \Xi(F_2) \right).$$
(EC.7)

Equation (EC.5) follows by letting  $\hat{\boldsymbol{\sigma}}_{\ell}$  iterate over the support of the distribution. Equation (EC.6) follows from the symmetry of the Rademacher distribution. That is, for every  $\hat{\boldsymbol{\sigma}}_{\ell}$ , there exists  $-\hat{\boldsymbol{\sigma}}_{\ell}$  with equal probability, and we need to only characterize half of the elements in the support. (EC.7) merges the suprema.

By the Obtuse Angle Criterion, projection to a convex set is a non-expansive operation (i.e.,  $\|\operatorname{proj}(F_1) - \operatorname{proj}(F_2)\| \leq \|F_1 - F_2\|$ ). We use the Cauchy-Schwarz inequality and the nonexpansiveness property (in (EC.8) and (EC.9) below, respectively) to remove the dependency on the projection operator:

RHS (EC.7) 
$$\leq \frac{1}{2^n} \sum_{\ell=1}^{2^{n-1}} \left( \sup_{F_1, F_2 \in \mathcal{F}} \| \hat{\sigma}_\ell \| \| \operatorname{proj}(F_1) - \operatorname{proj}(F_2) \| + \Xi(F_1) + \Xi(F_2) \right)$$
 (EC.8)

$$\leq \frac{1}{2^{n}} \sum_{\ell=1}^{2^{n-1}} \left( \sup_{F_{1}, F_{2} \in \mathcal{F}} \| \hat{\boldsymbol{\sigma}}_{\ell} \| \| F_{1} - F_{2} \| + \Xi(F_{1}) + \Xi(F_{2}) \right)$$
(EC.9)

$$\leq \frac{1}{2^{n}} \sum_{\ell=1}^{2^{n-1}} \left( \sup_{F_{1}, F_{2} \in \mathcal{F}} \sqrt{n} \|F_{1} - F_{2}\| + \Xi(F_{1}) + \Xi(F_{2}) \right)$$
(EC.10)

$$\leq \frac{1}{2} \left( \sup_{F_1, F_2 \in \mathcal{F}} \sqrt{n} \|F_1 - F_2\| + \Xi(F_1) + \Xi(F_2) \right)$$
(EC.11)

$$\leq \frac{1}{2} \left( \sqrt{n} \| F_1^* - F_2^* \| + \Xi(F_1^*) + \Xi(F_2^*) \right).$$
(EC.12)

Inequality (EC.10) follows by noting  $\|\sigma\| \leq \sqrt{n}$  for all  $\sigma \sim p_{\sigma}$  and (EC.11) from the fact that the dependency on  $\hat{\sigma}_{\ell}$  has been removed. We obtain (EC.12) by letting  $F_1^*$  and  $F_2^*$  be the two values that attain the supremum.

Finally, we use the Khintchine inequality to bound  $||F_1^* - F_2^*|| \leq \sqrt{2}\mathbb{E}_{\sigma \sim p_{\sigma}} \left[ \left| \boldsymbol{\sigma}^{\mathsf{T}} (F_1^* - F_2^*) \right| \right]$ . We then rearrange the terms as follows:

RHS (EC.12) 
$$\leq \frac{1}{2} \left( \sqrt{2n} \mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}} \left[ \left| \boldsymbol{\sigma}^{\mathsf{T}} (F_1^* - F_2^*) \right| \right] + \Xi(F_1^*) + \Xi(F_2^*) \right)$$
 (EC.13)

$$= \frac{1}{2} \left( \mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}} \left[ \sqrt{2n} \left| \boldsymbol{\sigma}^{\mathsf{T}} (F_1^* - F_2^*) \right| + \Xi(F_1^*) + \Xi(F_2^*) \right] \right)$$
(EC.14)

$$\leq \frac{1}{2} \left( \mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}} \left| \sup_{F_1, F_2 \in \mathcal{F}} \sqrt{2n} \left| \boldsymbol{\sigma}^{\mathsf{T}}(F_1 - F_2) \right| + \Xi(F_1) + \Xi(F_2) \right| \right)$$
(EC.15)

$$= \frac{1}{2} \left( \mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}} \left[ \sup_{F \in \mathcal{F}} \left\{ \sqrt{2n} \boldsymbol{\sigma}^{\mathsf{T}} F + \Xi(F) \right\} + \sup_{F \in \mathcal{F}} \left\{ -\sqrt{2n} \boldsymbol{\sigma}^{\mathsf{T}} F + \Xi(F) \right\} \right] \right)$$
(EC.16)

$$= \mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}} \left[ \sup_{F \in \mathcal{F}} \sqrt{2n} \boldsymbol{\sigma}^{\mathsf{T}} F + \Xi(F) \right].$$
(EC.17)

Inequality (EC.14) brings all of the terms inside the expectation. (EC.15) upper bounds by the supremum. Because  $\Xi(F_1) + \Xi(F_2)$  is invariant under the exchange of  $F_1$  and  $F_2$ , the supremum will be obtained when  $\sigma^{\mathsf{T}}(F_1 - F_2)$  is positive, meaning we can remove the absolute value and separate the supremum in (EC.16). Finally, the symmetry of the random variable  $\sigma$  implies that the two suprema are equal, thereby giving (EC.17).

To complete the proof, we use a standard conditioning argument (see Maurer (2016)) to show (EC.3) decomposes to (EC.4). For any  $0 \le m \le N_{\mathbf{u}}$ , we prove the following by induction:

$$\mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}}\left[\sup_{F \in \mathcal{F}} \sum_{i=1}^{N_{\mathbf{u}}} \boldsymbol{\sigma}_{i}^{\mathsf{T}} \operatorname{proj}(F(\hat{\mathbf{u}}_{i}))\right] \leq \mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}}\left[\sup_{F \in \mathcal{F}} \sum_{i=1}^{m} \sqrt{2n} \boldsymbol{\sigma}_{i}^{\mathsf{T}} F(\hat{\mathbf{u}}_{i}) + \sum_{i=m+1}^{N_{\mathbf{u}}} \boldsymbol{\sigma}_{i}^{\mathsf{T}} \operatorname{proj}(F(\hat{\mathbf{u}}_{i}))\right].$$

The case for m = 0 is an identity. Now for fixed values of  $\hat{\sigma}_i, \forall i \neq m$ , let

$$\Xi(F) = \sum_{i=1}^{m-1} \sqrt{2n} \hat{\boldsymbol{\sigma}}_i^{\mathsf{T}} F(\hat{\mathbf{u}}_i) + \sum_{i=m+1}^{N_{\mathbf{u}}} \hat{\boldsymbol{\sigma}}_i^{\mathsf{T}} \operatorname{proj} (F(\hat{\mathbf{u}}_i))$$

Then, assuming the inequality holds for m-1, we show

$$\mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}} \left[ \sup_{F \in \mathcal{F}} \sum_{i=1}^{N_{\mathbf{u}}} \boldsymbol{\sigma}_{i}^{\mathsf{T}} \operatorname{proj}(F(\hat{\mathbf{u}}_{i})) \right] \leq \mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}} \left[ \sup_{F \in \mathcal{F}} \sum_{i=1}^{m-1} \sqrt{2n} \boldsymbol{\sigma}_{i}^{\mathsf{T}} F(\hat{\mathbf{u}}_{i}) + \sum_{i=m}^{N_{\mathbf{u}}} \boldsymbol{\sigma}_{i}^{\mathsf{T}} \operatorname{proj}(F(\hat{\mathbf{u}}_{i})) \right] \\ = \mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}} \left[ \mathbb{E}_{\boldsymbol{\sigma}_{m} \sim p_{\boldsymbol{\sigma}}} \left[ \sup_{F \in \mathcal{F}} \boldsymbol{\sigma}_{m}^{\mathsf{T}} \operatorname{proj}(F(\hat{\mathbf{u}}_{m})) + \Xi(F) \middle| \{\hat{\boldsymbol{\sigma}}_{i}, \forall i \neq m\} \right] \right] \\ \leq \mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}} \left[ \mathbb{E}_{\boldsymbol{\sigma}_{m} \sim p_{\boldsymbol{\sigma}}} \left[ \sup_{F \in \mathcal{F}} \sqrt{2n} \boldsymbol{\sigma}_{m}^{\mathsf{T}} F(\hat{\mathbf{u}}_{m}) + \Xi(F) \middle| \{\hat{\boldsymbol{\sigma}}_{i}, \forall i \neq m\} \right] \right] \\ = \mathbb{E}_{\boldsymbol{\sigma} \sim p_{\boldsymbol{\sigma}}} \left[ \sup_{F \in \mathcal{F}} \sum_{i=1}^{m} \sqrt{2n} \boldsymbol{\sigma}_{i}^{\mathsf{T}} F(\hat{\mathbf{u}}_{i}) + \sum_{i=m+1}^{N_{\mathbf{u}}} \boldsymbol{\sigma}_{i}^{\mathsf{T}} \operatorname{proj}(F(\hat{\mathbf{u}}_{i})) \right].$$

The second inequality comes from substituting (EC.4). When  $m = N_{\mathbf{u}}$ , the proof is complete.  $\Box$ 

Lemma EC.3 can be seen as an extension of the main theorem of Maurer (2016) and is proved using a similar sequence of steps. There, the authors showed that composition of a Lipschitz scalarvalued vector function onto a vector-valued model class bounds the Rademacher complexity of the composed class by  $\sqrt{2L}$ . In the above, we compose the projection operator, a vector-valued function, to the vector-valued model class and bound the Rademacher complexity by  $\sqrt{2n}$ . Although we

=

only specifically consider the projection operator, the proof easily extends to any vector-valued function, so long as it is L-Lipschitz, whereupon we would reintroduce L back into the bound.

Before proving Theorem 3, we re-state the Generalization Lemma of Bertsimas and Kallus (2019).

LEMMA EC.4 (Bertsimas and Kallus (2019)). Consider a function  $z(\mathbf{x}, \mathbf{u}) : \mathcal{P} \times \mathcal{U} \to \mathbb{R}$  that is bounded and  $L_{\infty}$ -Lipschitz continuous in  $\mathbf{x}$  using the  $\|\cdot\|_{\infty}$  norm,

$$\sup_{\mathbf{x}\in\mathcal{P},\mathbf{u}\in\mathcal{U}} z(\mathbf{x},\mathbf{u}) \leq K, \quad \sup_{\mathbf{x}_1\neq\mathbf{x}_2\in\mathcal{P},\mathbf{u}\in\mathcal{U}} \frac{z(\mathbf{x}_1,\mathbf{u}) - z(\mathbf{x}_2,\mathbf{u})}{\|\mathbf{x}_1 - \mathbf{x}_2\|_{\infty}} \leq L_{\infty}.$$

For any  $\beta > 0$ , with probability at least  $1 - \beta$  with respect to the sampling of  $\hat{\mathcal{U}}$ ,

$$\mathbb{E}_{\mathbf{u}\sim\mathbb{P}_{\mathbf{u}}}\Big[z\big(F(\mathbf{u}),\mathbf{u}\big)\Big] \leq \frac{1}{N_{\mathbf{u}}}\sum_{i=1}^{N_{\mathbf{u}}} z\big(F(\hat{\mathbf{u}}_{i}),\hat{\mathbf{u}}_{i}\big) + K\sqrt{\frac{\log(1/\beta)}{2N_{\mathbf{u}}}} + L_{\infty}\mathfrak{R}_{N_{\mathbf{u}}}\big(\operatorname{proj}(\mathcal{F})\big), \quad \forall F \in \operatorname{proj}(\mathcal{F}).$$

We are now ready to prove Theorem 3.

Proof of Theorem 3. The proof follows by first applying Lemma EC.4, before applying Markov's inequality. We let  $z(\mathbf{x}, \mathbf{u}) = |f(\mathbf{x}) - f(\mathbf{x}^{\lambda}(\mathbf{u}))|$ , as a function of  $\mathbf{x} \in \mathcal{P}$  and  $\mathbf{u} \in \mathcal{U}$ , and show it is bounded from above

$$\sup_{\mathbf{x}\in\mathcal{P},\mathbf{u}\in\mathcal{U}} z(\mathbf{x},\mathbf{u}) = \sup_{\mathbf{x}\in\mathcal{P},\mathbf{u}\in\mathcal{U}} \left| f(\mathbf{x}) - f(\mathbf{x}^{\lambda}(\mathbf{u})) \right|$$
(EC.18)

$$\leq \max_{\mathbf{x}\in\mathcal{P}} f(\mathbf{x}) - \min_{\mathbf{x}\in\mathcal{P}} f(\mathbf{x}) = K.$$
(EC.19)

Because  $\mathcal{P}$  is a closed and bounded set and  $f(\mathbf{x})$  is linear, (EC.19) is bounded. We define K to be equal to RHS (EC.19).

We next show  $L_{\infty}$ -Lipschitz continuity,

$$\sup_{\mathbf{x}_1 \neq \mathbf{x}_2 \in \mathcal{P}, \mathbf{u} \in \mathcal{U}} \frac{z(\mathbf{x}_1, \mathbf{u}) - z(\mathbf{x}_2, \mathbf{u})}{\|\mathbf{x}_1 - \mathbf{x}_2\|_{\infty}} = \sup_{\mathbf{x}_1 \neq \mathbf{x}_2 \in \mathcal{P}, \mathbf{u} \in \mathcal{U}} \frac{\left|f(\mathbf{x}_1) - f(\mathbf{x}^{\lambda}(\mathbf{u}))\right| - \left|f(\mathbf{x}_2) - f(\mathbf{x}^{\lambda}(\mathbf{u}))\right|}{\|\mathbf{x}_1 - \mathbf{x}_2\|_{\infty}} \quad (\text{EC.20})$$

$$\leq \sup_{\mathbf{x}_1 \neq \mathbf{x}_2 \in \mathcal{P}, \mathbf{u} \in \mathcal{U}} \frac{\left| f(\mathbf{x}_1) - f(\mathbf{x}^{\lambda}(\mathbf{u})) - f(\mathbf{x}_2) + f(\mathbf{x}^{\lambda}(\mathbf{u})) \right|}{\left\| \mathbf{x}_1 - \mathbf{x}_2 \right\|_{\infty}}$$
(EC.21)

$$= \sup_{\mathbf{x}_1 \neq \mathbf{x}_2 \in \mathcal{P}} \frac{|f(\mathbf{x}_1) - f(\mathbf{x}_2)|}{\|\mathbf{x}_1 - \mathbf{x}_2\|_{\infty}} = L_{\infty}$$
(EC.22)

Inequality (EC.21) follows from the Reverse Triangle Inequality. (EC.22) follows from the fact that  $f(\mathbf{x})$  is linear and therefore, Lipschitz continuous using the  $\|\cdot\|_{\infty}$  norm. We let  $L_{\infty}$  be the Lipschitz constant of  $f(\mathbf{x})$ .

Because  $z(\mathbf{x}, \mathbf{u})$  satisfies the bounded and Lipschitz continuity assumptions, we apply Lemma EC.4 to obtain

$$\mathbb{E}_{\mathbf{u}\sim\mathbb{P}_{\mathbf{u}}}\Big[z\big(F(\mathbf{u}),\mathbf{u}\big)\Big] \leq \frac{1}{N_{\mathbf{u}}}\sum_{i=1}^{N_{\mathbf{u}}} z\big(F(\hat{\mathbf{u}}_{i}),\hat{\mathbf{u}}_{i}\big) + K\sqrt{\frac{\log(1/\beta)}{2N_{\mathbf{u}}}} + L_{\infty}\mathfrak{R}_{N_{\mathbf{u}}}\big(\operatorname{proj}(\mathcal{F})\big), \quad \forall F \in \operatorname{proj}(\mathcal{F}).$$

Specifically, this bound holds for  $F^* \in \operatorname{proj}(\mathcal{F})$ . By Lemma EC.3, we can bound  $\mathfrak{R}_{N_{\mathbf{u}}}(\operatorname{proj}(\mathcal{F})) \leq \sqrt{2n}\mathfrak{R}_{N_{\mathbf{u}}}(\mathcal{F})$ .

The remainder of the proof follows from Markov's inequality. For  $\gamma > 0$ ,

$$\mathbb{P}_{\mathbf{u}}\left\{z\left(F^{*}(\mathbf{u}),\mathbf{u}\right)>\gamma\right\} = \mathbb{P}_{\mathbf{u}}\left\{\left|f\left(F^{*}(\mathbf{u})\right) - f\left(\mathbf{x}^{\lambda}(\mathbf{u})\right)\right|>\gamma\right\}\right.$$
$$\leq \frac{\mathbb{E}_{\mathbf{u}\sim\mathbb{P}_{\mathbf{u}}}\left[\left|f\left(F^{*}(\hat{\mathbf{u}}_{i})\right) - f\left(\mathbf{x}^{\lambda}(\hat{\mathbf{u}}_{i})\right)\right|\right]}{\gamma}.$$

From the Law of Total Probability, we obtain

$$\begin{aligned} \mathbb{P}_{\mathbf{u}} \Big\{ \left| f\left(F^{*}(\mathbf{u})\right) - f\left(\mathbf{x}^{\lambda}(\mathbf{u})\right) \right| &\leq \gamma \Big\} &= 1 - \mathbb{P}_{\mathbf{u}} \Big\{ \left| f\left(F^{*}(\mathbf{u})\right) - f\left(\mathbf{x}^{\lambda}(\mathbf{u})\right) \right| > \gamma \Big\} \\ &\geq 1 - \frac{\mathbb{E}_{\mathbf{u} \sim \mathbb{P}_{\mathbf{u}}} \Big[ \left| f\left(F^{*}(\hat{\mathbf{u}}_{i})\right) - f\left(\mathbf{x}^{\lambda}(\hat{\mathbf{u}}_{i})\right) \right| \Big]}{\gamma}, \\ &\geq 1 - \frac{\frac{1}{N_{\mathbf{u}}} \sum_{i=1}^{N_{\mathbf{u}}} \left| f\left(F^{*}(\hat{\mathbf{u}}_{i})\right) - f\left(\mathbf{x}^{\lambda}(\hat{\mathbf{u}}_{i})\right) \right| + K \sqrt{\frac{\log(1/\beta)}{2N_{\mathbf{u}}}} + \sqrt{2n} L_{\infty} \mathfrak{R}_{N_{\mathbf{u}}}(\mathcal{F})}{\gamma} \end{aligned}$$

with probability  $1 - \beta$ . The second and third line follow from Markov's inequality and substituting the bound from Lemma EC.4, respectively. Given that we have a probabilistic bound for the error of  $F^*(\mathbf{u})$  from  $\mathbf{x}^{\lambda}(\mathbf{u})$ , we bound the error to  $\mathbf{x}^*(\mathbf{u})$ . Recall that  $f(\mathbf{x}^{\lambda}, \mathbf{u})$  is  $(\delta, \epsilon)$ -optimal. There are two cases to consider. First, if  $f(\mathbf{x}^{\lambda}(\mathbf{u})) \leq f(F(\mathbf{u})) \leq f(\mathbf{x}^{\lambda}(\mathbf{u})) + \gamma$ , then by substitution,

$$f\left(F^*(\mathbf{u})\right) - \epsilon - \gamma < f\left(\mathbf{x}^*(\mathbf{u})\right) < f\left(F^*(\mathbf{u})\right) + \delta L$$

Alternatively, if  $f(F^*(\mathbf{u})) \leq f(\mathbf{x}^{\lambda}(\mathbf{u})) \leq f(F^*(\mathbf{u})) + \gamma$ , then by substitution,

$$f(F^*(\mathbf{u})) - \epsilon < f(\mathbf{x}^*(\mathbf{u})) < f(F^*(\mathbf{u})) + \delta L + \gamma.$$

Note that both of these events can be covered by adding and subtracting  $\gamma$  to both the upper and lower bounds respectively. Then,

$$\mathbb{P}_{\mathbf{u}}\Big\{f\big(F^*(\mathbf{u})\big) - \epsilon - \gamma < f\big(\mathbf{x}^*(\mathbf{u})\big) < f\big(F^*(\mathbf{u})\big) + \delta L + \gamma\Big\} \ge \mathbb{P}_{\mathbf{u}}\Big\{\Big|f\big(F^*(\mathbf{u})\big) - f\big(\mathbf{x}^{\lambda}(\mathbf{u})\big)\Big| \le \gamma\Big\},$$

completing the proof.

# EC.3. Implementation details for predicting optimal dose distributions EC.3.1. Problem formulation

RT treatment is delivered by a linear accelerator (LINAC) that projects high-energy X-rays from different angles to a patient's tumor. The patient's body is discretized into voxels (i.e.,  $5 \text{mm} \times 5 \text{mm} \times 2 \text{mm}$  volumetric pixels) and the dose delivered to each of these voxels is used to assess

the quality of a treatment. The design of an RT treatment plan is typically done by mathematical optimization where the decision variable has two components (i) beamlet intensity and (ii) dose delivered (in Gy). Most automated planning systems involve predicting doses that are clinically desirable for a given patient. In practical implementations, the appropriate beamlets that can reconstruct a predicted dose distribution can be obtained via a secondary optimization process (Babier et al. 2018b). In this work, we focus on constructing clinically desirable dose distributions.

Each patient contains seven organs-at-risk (OARs) (i.e., brainstem, spinal cord, right parotid, left parotid, larynx, esophagus, and mandible) to which we minimize the average dose. Each patient also contains up to three planning target volumes (PTVs) with different prescription doses (i.e., PTV56, PTV63, and PTV70 with 56 Gy, 63 Gy, and 70Gy as prescription doses, respectively). We remark that constraints to the brainstem, spinal cord, and esophagus are generally easily satisfied by all predictions. Consequently, we focus specifically on the right parotid, left parotid, larynx, mandible, PTV56, PTV63, and PTV70.

In the RT optimization problem, each of the OARs and targets require polyhedral upper and lower bound constraints to the mean dose delivered to that structure. Furthermore, there exists a hidden "clinical criteria" constraint for each OAR and target that must be satisfied at the discretion of an oncologist. That is, if the ground truth treatment plan for a given patient from the data set satisfies a hidden constraint, then any generated plan for that patient must also satisfy that constraint. The hidden constraint for each OAR is an upper bound on either the mean or maximum dose delivered to that structure, while the hidden constraint for each target is a lower bound on the value-at-risk, i.e., minimum dose delivered to 90-th percentile of the target structure. The oracle  $\Psi(\mathbf{x}, \mathbf{u})$  is a look-up table that compares the dose generated by our model with the ground truth (i.e., what was actually delivered). In particular, for each structure,  $\Psi(\mathbf{x}, \mathbf{u})$  checks whether the input dose satisfies all the constraints (i.e., two polyhedral and one hidden constraints). We expand on the classification of feasibility in the sections below.

### EC.3.2. Neural network architecture

We use a modified version of the generative adversarial network (GAN) of (Mahmood et al. 2018), where two networks learn to predict dose distributions. The architectures for  $F(\mathbf{u})$  and  $B(\mathbf{x}, \mathbf{u})$  are described in Tables EC.1 and EC.2, respectively.

The generator takes as input a tensor  $\mathbf{u} \in \mathbb{R}^{128 \times 128 \times 128 \times 128 \times 8}$ , where the first three dimensions correspond to a voxel in the patient's geometry. The fourth dimension is a concatenation of the CT image greyscale and a one-hot encoded vector in  $\{0,1\}^7$  whose elements label whether the voxel belongs to one of the seven contoured structures. The generator then outputs a tensor  $\mathbf{x} \in \mathbb{R}^{128 \times 128 \times 128 \times 128}$  whose elements specify the dose to be delivered to each voxel of the patient.

		aropouti		
Layer	Concatenate with	Input shape	Block	Activation
1		$128 \times 128 \times 128 \times 8$	conv3d	BN-LR
2		$64\times 64\times 64\times 64$	conv3d	BN-LR
3		$32\times32\times32\times128$	conv3d	BN-LR
4		$16\times16\times16\times256$	conv3d	BN-LR
5		$8\times8\times8\times512$	conv3d	BN-LR
6		$4\times 4\times 4\times 512$	conv3d	BN-LR
7		$2\times 2\times 2\times 512$	deconv3d	LR
8	layer 5 output	$4\times 4\times 4\times 1024$	deconv3d	BN-R
9	layer 4 output	$8\times8\times8\times1024$	deconv3d	BN-D-R
10	layer 3 output	$16\times16\times16\times512$	deconv3d	BN-D-R
11	layer 2 output	$32\times32\times32\times256$	deconv3d	BN-R
12	layer 1 output	$64 \times 64 \times 64 \times 128$	deconv3d	AP-tanh
Output		$128 \times 128 \times 128 \times 1$		

Table EC.1Overview of the generator architecture. BN refers to batch normalization; LR, R, and tanh refer to<br/>Leaky ReLU (0.2 slope), ReLU, and Tanh activations, respectively; AP refers to a mean pool; and D refers to

 Table EC.2
 Overview of the classifier architecture. BN refers to batch normalization; LR, R, and sigmoid refer to Leaky ReLU (0.2 slope), ReLU, and Sigmoid activations.

-		_	
Layer	Input size	Block	Activation
1	$128 \times 128 \times 128 \times 9$	conv3d	LR
2	$64\times 64\times 64\times 64$	conv3d	BN-LR
3	$32\times32\times64\times128$	conv3d	BN-LR
4	$16\times16\times16\times256$	conv3d	BN-LR
5	$8\times8\times8\times512$	conv3d	$\operatorname{sigmoid}$
Output	7		

The classifier is trained to predict whether a given dose distribution satisfies all of the constraints (both hidden and polyhedral) for each structure of the patient. This network takes as input the concatenated tensor  $(\mathbf{x}, \mathbf{u})$  and outputs a vector in  $[0, 1]^7$ , whose elements each indicate the classifier's belief of whether the given dose distribution has satisfied all of the constraints for each specific structure. Consequently, learning feasibility becomes a multi-label classification problem and the classifier acts as seven separate classifiers each predicting feasibility with respect to an individual structure, but whose model parameters are shared with each other. For any structure, in order to classify a dose distribution as satisfying the relevant constraints, the classifier must: (i) first determine from the dose whether the polyhedral constraint are satisfied, (ii) determine from the dose whether the patient requires a hidden constraint to be satisfied, and (iii) determine from the dose whether the hidden constraint is satisfied if this constraint is required for the patient. Overall, a dose distribution is feasible only if all constraints are satisfied.

# EC.3.3. Implementation of the IPMAN algorithm

# Algorithm 1 Generator pre-training and data augmentation

**Input:** Feasible and input data sets  $\mathcal{D} = \{(\hat{\mathbf{x}}_i, \hat{\mathbf{u}}_i)\}_{i=1}^{N_{\mathbf{u}}}, \hat{\mathcal{U}} = \{\hat{\mathbf{u}}_i\}_{i=1}^{N_{\mathbf{u}}}, \text{ infeasible data set } \bar{\mathcal{D}} = \emptyset, \text{ Pretraining number of epochs } E_{ST}.$ 

**Output:** Pre-trained generative model  $F^{(0,0)}$ , Feasible and infeasible data sets  $\mathcal{D}, \overline{\mathcal{D}}$ .

- 1: Initialize generator and discriminator F, D
- 2: for e = 1 to  $E_{ST}$  do
- 3: Update generator and discriminator  $F^*, D^* \leftarrow \operatorname{Adam}(\nabla L_{ST})$ .
- 4: for all  $\hat{\mathbf{u}}_i \in \hat{\mathcal{U}}$  do

5: Append 
$$\mathcal{D} \leftarrow \mathcal{D} \cup (F^*(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i)$$
 if  $\Psi(F^*(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i) = 1$  else  $\bar{\mathcal{D}} \leftarrow \bar{\mathcal{D}} \cup (F^*(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i)$ .

- 6: **end for**
- 7: end for
- 8: return  $F^{(0,0)} \leftarrow F^*, \mathcal{D}, \overline{\mathcal{D}}$ .

In this subsection, we describe the exact implementation of the IPMAN algorithm used in our experiments. We summarize the steps in Algorithm 2. As our generative and classification models are neural networks, we remark on several improvements that can be made to the algorithm.

**EC.3.3.1. Pre-training as a GAN** Just as classical IPMs require a good initial point (i.e., lying within the feasible set) in order to construct a trajectory of points leading to an optimal solution, IPMAN can be made more efficient by ensuring that the generative model is initialized to predict points that are likely to be feasible. This initialization can greatly improve the training time and stability of the algorithm. Consequently, we first pre-train the generative model and subsequently apply transfer learning at the beginning of the algorithm (Goodfellow et al. 2016).

Pre-training amounts to training the generative model first as a Style Transfer GAN to learn to predict dose distributions from CT images as in Mahmood et al. (2018). The steps are summarized in Algorithm 1. In order to pre-train our generative model, we introduce a discriminator network  $D(\mathbf{x}, \mathbf{u}) : \mathbb{R}^{128 \times 128 \times 128} \times \mathbb{R}^{128 \times 128 \times 128 \times 3} \rightarrow \mathbb{R}$ . As in a Style Transfer GAN, we train the discriminator using  $\mathcal{D}$  to classify whether a given dose and CT image pair belongs to the data distribution and a generative model to predict dose distributions that fool the discriminator. Specifically,

$$\min_{F} \max_{D} \left\{ L_{ST} := \frac{1}{N_{\mathbf{u}}} \sum_{(\hat{\mathbf{u}}_i, \hat{\mathbf{x}}_i) \in \mathcal{D}} \log D(\hat{\mathbf{x}}_i, \hat{\mathbf{u}}_i) + \log \left(1 - D(F(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i)\right) + \lambda_{ST} \left\|F(\hat{\mathbf{u}}_i) - \hat{\mathbf{x}}_i\right\|_1 \right\}$$

The  $l_1$ -loss is a regularization term that ensures that the generator predicts dose distributions that resemble the ground truth and  $\lambda_{ST}$  is the regularization parameter. GANs are trained by iterative gradient descent between  $F(\mathbf{u})$  and  $D(\mathbf{x}, \mathbf{u})$ . In our implementation, we set  $\lambda_{ST} = 90$  and train the GAN for 50 epochs, following the practice from Mahmood et al. (2018). At the end of pre-training, we discard  $D(\mathbf{x}, \mathbf{u})$  and let  $F^{(0,0)}(\mathbf{u})$  denote the trained generative model.

# Algorithm 2 IPMAN

**Input:** Data sets of decisions  $\mathcal{D}, \bar{\mathcal{D}}$ , and inputs  $\hat{\mathcal{U}} = \{\hat{\mathbf{u}}_i\}_{i=1}^{N_{\mathbf{u}}}$ , Set of dual variables  $\{\lambda_j\}_{j=0}^M$ , Number of iterations K, Number of epochs  $E_B, E_F$ , Subset sampling rate s**Output:** Final generative models  $F^{(j,K)}$  for  $j \in \{0, \dots, M\}$ 

1: Pre-train generator using Algorithm 1.

- 2: Initialize generator  $F^{(j,0)} \leftarrow F^*$  for  $j \in \{0, \dots, M\}$ , classifier B.
- 3: for k = 1 to K do
- 4: Sample subsets to train  $\mathcal{D}^{(k)} = \sigma(\mathcal{D}; s), \ \bar{\mathcal{D}}^{(k)} = \sigma(\bar{\mathcal{D}}; s|\mathcal{D}|/|\bar{\mathcal{D}}|).$
- 5: for e = 0 to  $E_B$  do
- 6: Update classifier  $B^{(k)} \leftarrow \operatorname{Adam}(\nabla L_B)$ .
- 7: end for
- 8: for j = 0 to M do
- 9: for e = 0 to  $E_F$  do
- 10: Update generator  $F^{(j,k)} \leftarrow \operatorname{Adam}(\nabla L_F)$ .
- 11: **end for**

```
12: for all \hat{\mathbf{u}}_i \in \hat{\mathcal{U}} do
```

- 13: Append  $\mathcal{D} \leftarrow \mathcal{D} \cup (F^{(j,k)}(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i)$  if  $\Psi(F^{(j,k)}(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i) = 1$  else  $\bar{\mathcal{D}} \leftarrow \bar{\mathcal{D}} \cup (F^{(j,k)}(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i)$ .
- 14: **end for**
- 15: **end for**
- 16: **end for**
- 17: **return**  $F^{(j,K)}$  for  $j \in \{0, ..., M\}$

EC.3.3.2. Sampling an infeasible data set of decisions  $\overline{\mathcal{D}}$  Training IPMAN requires an initial data set of infeasible decisions  $\overline{\mathcal{D}}$ . In practice, a data set of infeasible decisions would not be available a priori and, instead, is generated by sampling. Note, however, that in every epoch of the pre-training step, the generative model generates candidate solutions  $F(\hat{\mathbf{u}}_i)$  to attempt to fool the discriminator. We save the generated decisions during pre-training and label them afterwards as feasible or infeasible using the oracle. By training using 100 patients for 50 epochs, we generate a total of 5000 dose distributions that are labelled as feasible or infeasible and then binned in the appropriate  $\mathcal{D}$  or  $\overline{\mathcal{D}}$ , respectively. The steps are summarized in Algorithm 1.

EC.3.3.3. Learning multi-label feasibility with sub-sampled data sets Training IPMAN for multiple iterations can produce a large quantity of generated data points. Furthermore, because we consider feasibility for each structure separately, training the classifier quickly becomes prohibitively expensive. In order to reduce training time, we do not use the entire data sets  $\mathcal{D}$ 

and  $\bar{\mathcal{D}}$  but rather smaller sampled subsets. Let  $\sigma(\cdot; s)$  be a random sampling operator (without replacement) where s is the fraction of points to sample. For example,  $\sigma(\mathcal{D}; 0.5)$  denotes a randomly sampled subset of size  $0.5|\mathcal{D}|$ . In our implementation, we set s = 0.3 and trained the classifier using  $\mathcal{D}^{(k)} = \sigma(\mathcal{D}; 0.3)$  and  $\bar{\mathcal{D}}^{(k)} = \sigma(\bar{\mathcal{D}}; 0.3|\mathcal{D}|/|\bar{\mathcal{D}}|)$ ; this reduced the training time to 24 hours.

We next define the multi-label classification problem. For any  $(\hat{\mathbf{u}}_i, \hat{\mathbf{x}}_i)$  in  $\mathcal{D}$  or  $\mathcal{D}$ , let  $\psi_{i,r}$  be a label determining whether the dose distribution had satisfied the polyhedral and (conditional) hidden constraints for structure r. That is, if the ground truth dose for  $\hat{\mathbf{u}}_i$  satisfied the hidden constraints,  $\psi_{i,r} = 1$  if the polyhedral and hidden constraints were satisfied and zero otherwise. If the clinical dose did not satisfy the hidden constraint, then  $\psi_{i,r} = 1$  if only the polyhedral constraints were satisfied; here, the hidden constraint is inactive for this patient. Then, let  $[B(\mathbf{x}, \mathbf{u})]_r$  denote the r-th element of the classifier output. The classifier problem is

$$\max_{B \in \mathcal{B}} \left\{ L_B := \frac{1}{N_{\mathbf{x}} + \bar{N}_{\mathbf{x}}} \sum_{\substack{(\hat{\mathbf{x}}_i, \hat{\mathbf{u}}_i) \in \\ \mathcal{D}^{(k)} \cup \bar{\mathcal{D}}^{(k)}}} \sum_{r \in \mathcal{R}} \left( \psi_{i,r} \log \left[ B(\hat{\mathbf{x}}_i, \hat{\mathbf{u}}_i) \right]_r + (1 - \psi_{i,r}) \log \left( 1 - \left[ B(\hat{\mathbf{x}}_i, \hat{\mathbf{u}}_i) \right]_r \right) \right) \right\}.$$

The above problem specializes to  $\mathbf{FCP}(\mathcal{D}^{(k)}, \bar{\mathcal{D}}^{(k)})$  in the single-class setting (i.e.,  $|\mathcal{R}| = 1$ ). For a dose distribution to be classified feasible,  $B(\mathbf{x}, \mathbf{u})_r = 1$  for all  $r \in \mathcal{R}$ . This approach of separating the constraint satisfaction along all structures individually is equivalent to modeling the optimization problem via a barrier function for each structure. Furthermore, the barriers are approximated by a neural network classifier with shared weights except in the last layer. As we describe later below, the objective of the barrier optimization problem is obtained by summing all of the separate barriers, i.e.,  $f(\mathbf{x}) - \lambda \sum_{r \in \mathcal{R}} \log[B(\mathbf{x}, \mathbf{u})]_r$ . Finally, we minimize  $L_B$  using the Adam optimizer for  $E_B = 10$  epochs in every iteration. Note that it is essential to ensure that the classifier accurately predicts feasibility in order to be able to approximate a  $\delta$ -barrier.

EC.3.3.4. Regularized barrier optimization problem We include an  $l_1$  regularization term in training. This term is equivalent to the one used in the pre-training stage and is useful to ensure that predicted dose distributions do not deviate too far from the ground truth. Note that we only use this regularization in the first set of experiments (Section 6.3) and remove it in the second set of experiments (Section 6.4). There, the ground truth plans may not be feasible, meaning that it would be incorrect to replicate ground truth behavior. However, a consequence of removing a regularization term is that certain models may become unstable and deviate significantly if the classifier is not a complete  $\delta$ -barrier. We observe this behavior in  $\lambda = 4$  where the model minimizes dose while ignoring feasibility. With slight abuse of notation, let  $\mathbf{z}(F(\mathbf{u}))$  denote the vector of average doses to each structure as constructed by the generative model. Then, the generative barrier problem in this setting is

$$\min_{F \in \mathcal{F}} \left\{ L_F := \frac{1}{N_{\mathbf{u}}} \sum_{\hat{\mathbf{u}}_i \in \hat{\mathcal{U}}} \left( \frac{1}{\lambda_j} f\left( \mathbf{z} \left( F(\hat{\mathbf{u}}_i) \right) \right) + \lambda_{ST} \| F(\hat{\mathbf{u}}_i) - \hat{\mathbf{x}}_i \|_1 - \sum_{r \in \mathcal{R}} \left[ B^{(k)} \left( F(\hat{\mathbf{u}}_i), \hat{\mathbf{u}}_i \right) \right]_r \right) \right\}.$$

We minimize  $L_F$  using the Adam optimizer for  $E_F = 1$  epoch in every iteration. It is important to ensure that the classifier is trained close to optimality to ensure that it approximates a  $\delta$ -barrier. Furthermore, given the nature of training the classifier, it is often the case that the classifier's support is uneven and may have areas of local optimality that the generator may abuse. A standard practice in the GAN literature is to control the training duration of the two networks; we employ a similar strategy by training the generative model for a shorter duration than the classifier in order to ensure that the generator does not overfit and abuse local optima caused by the classifier. As previously mentioned, we set  $\lambda_{ST} = 50$  for the first set of experiments but require that  $\lambda_{ST} = 0$  for the second set.