

IMPROVING FEASIBILITY VIA FAST AUTOENCODER-BASED PROJECTIONS

Anonymous authors

Paper under double-blind review

ABSTRACT

Enforcing complex (e.g., nonconvex) operational constraints is a critical challenge in real-world learning and control systems. However, existing methods struggle to efficiently and reliably enforce general classes of constraints. To address this, we propose a novel data-driven amortized approach that uses a trained autoencoder as an approximate projector to provide fast corrections to infeasible predictions. Specifically, we train an autoencoder using an adversarial objective to learn a structured, convex latent representation of the feasible set, enabling rapid correction of neural network outputs by projecting them onto a simple convex shape before decoding into the original feasible set. We test our approach on a diverse suite of constrained optimization and reinforcement learning problems with challenging nonconvex constraints. Results show that our method effectively improves constraint satisfaction at a low computational cost, offering a practical alternative to expensive feasibility correction techniques based on traditional solvers.

1 INTRODUCTION

Many learning and control systems, across areas such as robotics, energy systems, and industrial automation, must produce outputs that respect difficult-to-satisfy (often nonconvex) constraints. Enforcing these constraints reliably and efficiently is crucial for both safety and real-world usability. A number of approaches have been proposed to address this challenge within learning-based systems, such as penalty methods (Fioretto et al. (2021); Stooke et al. (2020)), differentiable projection and correction methods (Chen et al. (2021); Pan et al. (2022)), and post-hoc repair algorithms (Zamzam & Baker (2020); Nocedal & Wright (2006)). However, each of these approaches comes with distinct trade-offs in terms of reliability, generality, computational cost, and solution quality – for instance, failing to satisfy constraints in practice, not handling general classes of constraints, being prohibitively slow to deploy in real-world systems, and/or degrading end-to-end model performance.

To address this challenge, we propose a data-driven amortized alternative to traditional constraint enforcement algorithms: a trained autoencoder that acts as an approximate projector to provide fast corrections to infeasible predictions. Specifically, we train an autoencoder to learn a structured, convex latent representation of the feasible set, in a way that enables rapid mapping from infeasible to feasible points. This autoencoder can then be leveraged as a plug-and-play “attachment” to standard neural networks. While not aiming to supersede strict projections in regimes demanding hard feasibility, this approach is designed to improve feasibility in low-latency settings or in settings where moving predictions closer to the feasible set suffices to guarantee downstream system performance.

Our key contributions are as follows:

- **Framework for data-driven feasibility improvement.** We pose an approach for learning feasibility improvement mappings that can be appended to neural networks, as an alternative to expensive, albeit

047 exact, constraint enforcement approaches. We propose a particular instantiation of this approach, FAB,
048 which employs an autoencoder as the basis of the learned mapping.

- 049 • **Structured latent representation learning.** We propose a mechanism to learn faithful, feasibility-
050 preserving latent representations of the feasible set, via adversarial training.
- 051 • **Empirical validation.** We test our method on a range of constrained optimization and reinforcement
052 learning (RL) problems. For constrained optimization problems, the method consistently provides an
053 efficient approximate projection, learning to map solutions to the feasible set close to 100% of the time
054 in a fraction of a millisecond (faster than any other method). For RL settings, the method consistently
055 provides safer actions than methods like proximal policy optimization (PPO) and trust-region policy
056 optimization (TRPO), as well as their constrained variants. Overall, this demonstrates the promise of
057 our approach in providing fast, reliable feasibility improvements across a wide range of settings.

059 2 RELATED WORK

061 **Amortized optimization with constraints.** Also known as *learning to optimize*, amortized optimization
062 is a paradigm designed to accelerate the process of solving optimization problems by using machine learn-
063 ing models as fast function approximators for optimization and control problems (Amos, 2023). Models
064 are trained either via supervised learning on a dataset of known solutions (Chen et al., 2022), or through
065 unsupervised/self-supervised methods, where the model’s loss function incorporates the optimization prob-
066 lem’s objectives and constraints (Van Hentenryck, 2025). A primary challenge in amortized optimization is
067 ensuring that the model’s output satisfies constraints. Several lines of work have emerged to address this,
068 including penalty methods, differentiable projection and/or correction methods, and post-hoc “repair tech-
069 niques.” Penalty methods turn constrained optimization problems into unconstrained ones by incorporating
070 penalty terms for constraint violations in the objective terms (Fioretto et al., 2021; Stooke et al., 2020; Raissi
071 et al., 2019); however, while they incentivize feasibility, they do not guarantee it, and may produce highly
072 infeasible solutions in practice. In contrast, differentiable projection and/or correction methods embed exact
073 solvers directly as layers in neural networks (Pham et al., 2018; Chen et al., 2021; Pan et al., 2022; Nguyen &
074 Donti, 2025; Donti et al., 2021). While these methods do provide feasibility guarantees, they are often either
075 expensive to run, or highly specialized to certain classes of constraints (Min & Azizian, 2024; Tordesillas
076 et al., 2023). Likewise, post-hoc “repair” methods (Boyd et al., 2011; Douglas & Rachford, 1956; Zamzam
077 & Baker, 2020) are often either expensive or highly specialized, and the inherent train-test mismatch can
078 further degrade overall performance.

079 In this work, we propose an alternative, data-driven approach that can learn an inexpensive, non-iterative
080 transformation from a latent convex set to any arbitrary constraint set, thereby speeding up feasibility im-
081 provement in practice (albeit at the expense of provable guarantees). The work closest in the literature to
082 ours is (Liang et al., 2024; 2023), proposing a homeomorphic projection approach which learns an invertible
083 mapping from a hypersphere to a topologically-equivalent constraint set, and then uses a bisection proce-
084 dure to recover feasibility. However, this procedure is designed for post-hoc feasibility correction rather
085 than end-to-end training, and is limited to ball-homeomorphic constraint sets. Our work presents a method
086 that can learn a mapping between a latent hypersphere and any continuous constraint set, and is compatible
087 with end-to-end training and inference. **In addition, while homeomorphic projections rely on an iterative
bisection procedure, FAB projections are one-shot, leading to significant speed gains in practice.**

088 **Safe reinforcement learning.** Safe reinforcement learning (RL) is an extension of standard RL, where
089 an agent learns to make a sequence of decisions in an environment to maximize a cumulative reward signal
090 while also aiming to minimize the costs associated with constraint violations (García & Fernández, 2015; Gu
091 et al., 2024). While standard RL algorithms, such as Proximal Policy Optimization (PPO, Schulman et al.,
092 2017) and Trust Region Policy Optimization (TRPO, Schulman et al., 2015) excel in unconstrained set-
093 tings, they are not equipped to handle constraints. Prominent approaches to address this include Lagrangian

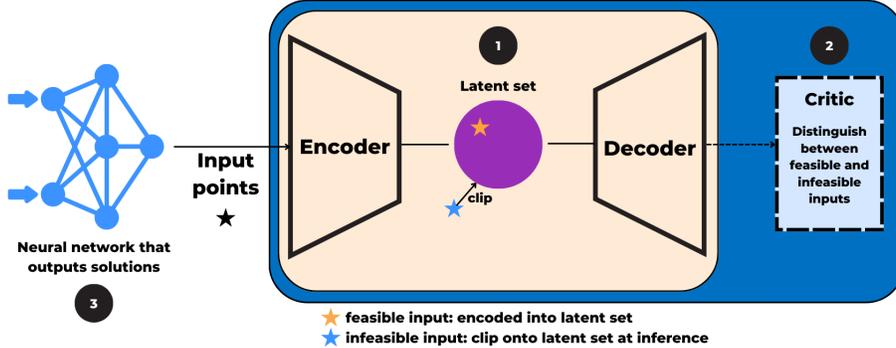


Figure 1: A schematic of FAB approximate projections. (1) Phase 1 of autoencoder training aims to enable reconstructions of the feasible set. (2) Phase 2 of autoencoder training introduces a discriminator to enable further structuring and refinement of the latent representation. (3) The trained autoencoder can be utilized as a plug-and-play attachment to another neural network model.

relaxation (Stooke et al., 2020), safe exploration (Hans et al., 2008), and differentiable projection (Pham et al., 2018; Chen et al., 2021). However, mirroring the discussion on constraint enforcement in amortized optimization, these approaches come with distinct trade-offs in terms of reliability and computational cost.

Adversarial training. Adversarial training, widely known for its application in Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) and adversarially robust deep learning (Bai et al., 2021), offers a general framework for learning in the face of complex data distributions. In the canonical GAN setup, a *generator* network learns to produce realistic data samples from a noise vector, while a *discriminator* network is trained to distinguish between real and generated samples. In particular, minimax training of the generator and the discriminator drives the generator to produce increasingly plausible outputs. Inspired by this approach, our work leverages an adversarial training paradigm to learn feasible sets, where the “discriminator” is trained to distinguish between feasible and infeasible outputs produced by another model, and that output-producing model is thereby forced to learn more robust decision boundaries between feasible and infeasible points.

3 FAST AUTOENCODER-BASED (FAB) FEASIBILITY IMPROVEMENT

In this paper, we consider the task of repeatedly solving parametric optimization problems of the form:

$$y^*(x) \in \arg \min_{y \in \mathcal{C}(x)} f(y; x), \quad (1)$$

where $x \in \mathcal{X} \subseteq \mathbb{R}^m$ are the problem parameters, $y \in \mathcal{Y} \subseteq \mathbb{R}^n$ are decision variables, $f : \mathcal{Y} \times \mathcal{X} \rightarrow \mathbb{R}$ is the objective function, $\mathcal{C}(x) \subseteq \mathcal{Y}$ is a (parameter-dependent) constraint set, and $y^*(x) \in \mathcal{C}(x)$ is a solution. Because the optimal solutions of the problems change as the parameters x change, this suggests a mapping between parameters x and optimal solutions. In addition to offline optimization, this formulation also captures RL settings, where \mathcal{X} and \mathcal{Y} are the state and action spaces, respectively (Amos, 2023). Our aim is to use neural networks to approximate a mapping from x to $y^*(x)$ such that the resultant solution has low objective value, satisfies constraints, and is fast to compute at inference time.

We specifically consider neural networks of the form:

$$\hat{y}_\theta := \phi_x \circ N_\theta, \quad (2)$$

where $N_\theta : \mathcal{X} \rightarrow \mathbb{R}^d$ is a standard feedforward network with parameters θ , and $\phi_x : \mathbb{R}^d \rightarrow \mathcal{C}(x)$ is a differentiable feasibility improvement procedure that aims to map to a point in the constraint set. Prior work

has considered the design of mappings ϕ_x that can ensure exact feasibility, e.g., by solving a differentiable orthogonal projection or via exact iterative procedures (see §2). While such approaches are indeed useful for, e.g., safety-critical settings where provable constraint satisfaction is a must, they tend to be relatively expensive for general classes of constraints. Instead, we propose to learn a fast, data-driven, *approximate* mapping for use in settings where (near-)feasibility is important, but where the benefits of reduced latency outweigh the need for strict constraint satisfaction.

The particular choice of ϕ_x that we propose in this work is motivated by the observation that while $\mathcal{C}(x)$ may be expensive to enforce directly, we can potentially learn a transformation between this set and some specially-structured set that is much cheaper to work with. In particular, let $E_\gamma : \mathcal{I} \rightarrow \mathcal{Z}$ be an encoder with parameters γ , where $\mathcal{I} := \mathcal{Y} \times \mathcal{X}$ and $\mathcal{Z} \subseteq \mathbb{R}^k$, and let $R_\psi : \mathcal{Z} \rightarrow \mathcal{Y}$ be a decoder (“reconstructor”) with parameters ψ , where both E_γ and R_ψ are standard feedforward neural networks.¹ In addition, let $\mathcal{S} \subseteq \mathcal{Z}$ be a set that is by construction simple to map onto (e.g., a simplex or a ball), and let $\phi^{\mathcal{S}} : \mathcal{Z} \rightarrow \mathcal{S}$ be a cheap exact mapping to points in \mathcal{S} (e.g., a softmax operation or closed-form projection). We then define our feasibility improvement procedure as

$$\phi_x := R_\psi \circ \phi^{\mathcal{S}} \circ E_\gamma. \quad (3)$$

In other words, given an output from N_θ alongside its corresponding input, we feed these into an encoder, map the encoded latent point into \mathcal{S} , and then decode the resultant point. The idea is that once E_γ and R_ψ are trained, each step of ϕ_x is cheap to execute, leading to fast overall inference.

The key challenge with this approach is that we must now design the autoencoder parameters (γ, ψ) such that the output of ϕ_x is actually feasible, i.e., actually lies within $\mathcal{C}(x)$. To do this, we propose a two-stage autoencoder training procedure aimed at structuring the latent space \mathcal{Z} such that, to the extent possible, latent points decode to our feasible set $\mathcal{C}(x)$ if and only if they lie within $\mathcal{S} \subseteq \mathcal{Z}$. A schematic of our overall approach is provided in Figure 2. We now provide additional detail on the two-stage training procedure.

3.1 TWO-PHASE AUTOENCODER TRAINING

We train the encoder-decoder pair (E_γ, R_ψ) in two phases: (1) a *constraint set reconstruction phase* leveraging standard autoencoder training, and (2) a *latent space structuring phase* involving adversarial training with a discriminator. Together, these aim to enable the autoencoder to faithfully reconstruct the feasible set, and encourage latent points from \mathcal{S} to decode to feasible points. Pseudocode for both training phases is given in Algorithm 1

Phase 1: Constraint set reconstruction. We first train the autoencoder to reconstruct the feasible set. Specifically, we train on a dataset of exclusively feasible points, $\mathcal{T}_{\text{feas}} := \{(y^{(i)}, x^{(i)}) \mid y^{(i)} \in \mathcal{C}(x^{(i)})\}$. As standard in autoencoder training, the encoder E_γ and the decoder R_ψ are jointly trained to minimize a standard L_2 reconstruction loss to prioritize the fidelity of decoded points:

$$\mathcal{L}_{\text{recon}}(y, x) = \|y - R_\psi(E_\gamma(y, x))\|_2^2. \quad (4)$$

Phase 2: Latent space structuring. The second autoencoder training phase aims to structure the latent space such that any point sampled from $\mathcal{S} \subseteq \mathcal{Z}$, when decoded, results in a feasible point. To do this, we draw loose inspiration from generative adversarial network (GAN) training, and train our autoencoder alternatingly with a discriminator $D_\xi : \mathcal{I} \rightarrow [0, 1]$ whose role is to distinguish between feasible and infeasible points. Specifically, the discriminator maps from inputs in \mathcal{I} to an estimated probability that the input is feasible. We let D_ξ be a standard feedforward neural network.

¹We focus our discussion in the main text on the case of input-varying constraints $\mathcal{C}(x)$, but note that the formulations for input-independent constraints are similar. In the latter case, $\mathcal{I} := \mathcal{Y}$; our autoencoder is simply a standard, rather than conditional, autoencoder; and all mentions of x in §3.1 can be dropped.

Algorithm 1 Two-Phase Autoencoder Training for FAB Feasibility Improvement

```

188 1: input: Dataset of feasible points  $\mathcal{T}_{\text{feas}} = \{(y^{(i)}, x^{(i)}) \mid y^{(i)} \in \mathcal{C}(x^{(i)})\}$ 
189
190 2: input: Discriminator dataset  $\mathcal{T}_{\text{disc}} := \{(y^{(i)}, x^{(i)}), c^{(i)} \mid c^{(i)} = 1 \text{ if } y^{(i)} \in \mathcal{C}(x^{(i)}), 0 \text{ otherwise}\}$ 
191
192 3:
193 4: procedure PHASE1( $\mathcal{T}_{\text{feas}}$ ) // Constraint set reconstruction
194 5:   init encoder  $E_\gamma$ , decoder  $R_\psi$ 
195 6:   while not converged do for batch of  $(y, x) \in \mathcal{T}_{\text{feas}}$ 
196 7:     compute  $\mathcal{L}_{\text{recon}}(y, x)$  via Eq. 4
197 8:     update  $\gamma, \psi$  using  $\nabla_\gamma \mathcal{L}_{\text{recon}}, \nabla_\psi \mathcal{L}_{\text{recon}}$ 
198 9:   end while
199 10:  return  $E_\gamma, R_\psi$ 
200 11: end procedure
201
202 12:
203 13: procedure PHASE2( $\mathcal{T}_{\text{disc}}, E_\gamma, R_\psi$ ) // Latent set structuring
204 14:  init discriminator  $D_\xi$ 
205 15:  while not converged do for batch of  $((y, x), c) \in \mathcal{T}_{\text{disc}}$ 
206 16:    // Update discriminator  $D_\xi$ 
207 17:    compute  $\mathcal{L}_{\text{disc}}(y, x, c)$  via Eq. 5
208 18:    update  $\xi$  using  $\nabla_\xi \mathcal{L}_{\text{disc}}$ 
209 19:    // Update autoencoder ( $E_\gamma, R_\psi$ )
210 20:    compute  $\mathcal{L}_{\text{recon}}(y, x), \mathcal{L}_{\text{hinge}}(y, x, c)$  via Eq. 4, 7
211 21:    sample batch of  $z \sim \mathcal{S}$ 
212 22:    compute  $\mathcal{L}_{\text{latent}}(z), \mathcal{L}_{\text{geom}}(\text{batch of } z)$  via Eq. 8–9
213 23:    compute  $\mathcal{L}_{\text{struc}}$  from  $\mathcal{L}_{\text{recon}}, \mathcal{L}_{\text{hinge}}, \mathcal{L}_{\text{latent}}, \mathcal{L}_{\text{geom}}$  via Eq. 6
214 24:    update  $\gamma, \psi$  using  $\nabla_\gamma \mathcal{L}_{\text{struc}}, \nabla_\psi \mathcal{L}_{\text{struc}}$ 
215 25:  end while
216 26:  return  $E_\gamma, R_\psi$ 
217 27: end procedure

```

We leverage two datasets during this training phase: (a) a labeled dataset of feasible *and* infeasible points, $\mathcal{T}_{\text{disc}} := \{(y^{(i)}, x^{(i)}), c^{(i)} \mid c^{(i)} = 1 \text{ if } y^{(i)} \in \mathcal{C}(x^{(i)}), 0 \text{ otherwise}\}$, and (b) samples $z^{(j)} \sim \mathcal{S}$ from our latent subset. Unlike in GANs, where the generator and discriminator are trained via a minimax loss, here, the autoencoder and discriminator are trained using distinct loss functions.

Specifically, the discriminator D_ξ is trained to distinguish between feasible and infeasible points, by minimizing the negative log likelihood on $\mathcal{T}_{\text{disc}}$:

$$\mathcal{L}_{\text{disc}}(y, x, c) = -(c \log D_\xi(y, x) + (1 - c) \log (1 - D_\xi(y, x))). \quad (5)$$

The autoencoder is trained to minimize a composite loss function leveraging both $\mathcal{T}_{\text{disc}}$ and $z^{(j)} \sim \mathcal{S}$, defined as:

$$\begin{aligned} \mathcal{L}_{\text{struc}} = & \frac{1}{N} \sum_{i=1}^N \left[\lambda_{\text{recon}} \mathcal{L}_{\text{recon}}(y^{(i)}, x^{(i)}) \right] + \frac{1}{N} \sum_{i=1}^N \lambda_{\text{hinge}} \left[\mathcal{L}_{\text{hinge}}(y^{(i)}, x^{(i)}, c^{(i)}) \right] \\ & + \frac{1}{M} \sum_{j=1}^M \left[\lambda_{\text{latent}} \mathcal{L}_{\text{latent}}(z^{(j)}) \right] + \lambda_{\text{geom}} \mathcal{L}_{\text{geom}}(\{z^{(1)}, \dots, z^{(M)}\}), \end{aligned} \quad (6)$$

where $\lambda_{\text{recon}}, \lambda_{\text{latent}}, \lambda_{\text{hinge}}, \lambda_{\text{geom}} \in \mathbb{R}$, where $\mathcal{L}_{\text{recon}}$ is defined similarly as above, and where the remaining loss terms are defined as follows:

- The hinge loss $\mathcal{L}_{\text{hinge}}$ structures the latent space by mapping feasible points to the interior of a hypersphere and infeasible points to its exterior. Here, we write the hinge loss for the specific choice of $\mathcal{S} := \{z : \|z\|_2 \leq r\}$ for some radius r :

$$\mathcal{L}_{\text{hinge}}(y, x, c) = c \text{ReLU}(\|E_\gamma(y, x)\|_2 - r) + (1 - c) \text{ReLU}(r - \|E_\gamma(y, x)\|_2). \quad (7)$$

- The latent loss $\mathcal{L}_{\text{latent}}$ encourages outputs decoded from latent points $z \in \mathcal{S}$ to be feasible, via supervision from the discriminator:

$$\mathcal{L}_{\text{latent}}(z) = -\log D_\xi(R_\psi(z)). \quad (8)$$

- The Jacobian regularization term $\mathcal{L}_{\text{geom}}$ encourages uniform coverage of the feasible set by the decoder (Nazari et al., 2023), and is computed over a set of latent points $\hat{\mathcal{S}} \subseteq \mathcal{S}$:

$$\mathcal{L}_{\text{geom}}(\hat{\mathcal{S}}) = \text{Variance}_{z \sim \hat{\mathcal{S}}} [\log \det(J_z J_z^\top + \varepsilon I_k)], \quad (9)$$

where $J_z = \nabla_z R_\psi(z)$ is the Jacobian of the decoder with respect to z , I_k is the $k \times k$ identity matrix, and ε is a small scalar value. This loss term measures how the decoder’s output changes under small changes in the latent code (εI_k). Specifically, it calculates the determinant of the Gram matrix $J_z J_z^\top$ and measures how much the decoder locally stretches or shrinks the space.

3.2 LEVERAGING THE TRAINED FEASIBILITY MAPPING

After autoencoder training, we fix the weights of the encoder and decoder and use them to construct the feasibility mapping $\phi_x = R_\psi \circ \phi^{\mathcal{S}} \circ E_\gamma$, as also given by Equation 3. We then append the feasibility mapping to our base neural network as $\hat{y}_\theta = \phi_x \circ N_\theta$, as given by Equation 2. The resultant neural network \hat{y}_θ (with learnable parameters θ) can then be trained end-to-end as usual to address the amortized optimization problem at hand, i.e., to [aim to learn optimal solutions to Problem 1](#). In particular, N_θ [learns to adapt to the autoencoder’s mapping procedure as part of the end-to-end training, and can account for this to pick neural network parameters aimed at improving the optimality of the end-to-end solution](#). (While in principle the encoder and decoder parameters could be further adjusted end-to-end alongside the base neural network parameters θ , rather than being fixed, we do not consider that case here.) Overall, our learned mapping serves as a plug-and-play attachment to standard deep learning models that be used during both training and inference to improve model feasibility, while also being fast to run. In the following sections, we demonstrate the performance of \hat{y}_θ in learning feasible solutions to amortized optimization problems across both offline optimization and RL settings.

4 EXPERIMENTS ON CONSTRAINED OPTIMIZATION PROBLEMS

Problem classes. We test our approach on 4 nonconvex constraint families (Fig. 2) and 3 objective types: linear, quadratic, and distance minimization. The nonconvex constraint sets are defined analytically; we consider input-independent constraint sets $\mathcal{C} \subseteq \mathcal{Y}$. The problems are formulated as:

Linear:	Quadratic:	Distance minimization:
$\min_y a^\top y$	$\min_y y^\top Q y + a^\top y$	$\min_y \ y - t\ ^2$
s.t. $y \in \mathcal{C}$	s.t. $y \in \mathcal{C}$	s.t. $y \in \mathcal{C}$,

where $y \in \mathcal{Y} := \mathbb{R}^n$ is the decision variable, and the problem parameters $a \in \mathbb{R}^n$, $Q \in \mathbb{R}^{n \times n}$, $t \in \mathbb{R}^n$ vary between instances. Our goal is to learn a neural network approximator \hat{y}_θ for each problem class that minimizes the objective function while satisfying constraints.

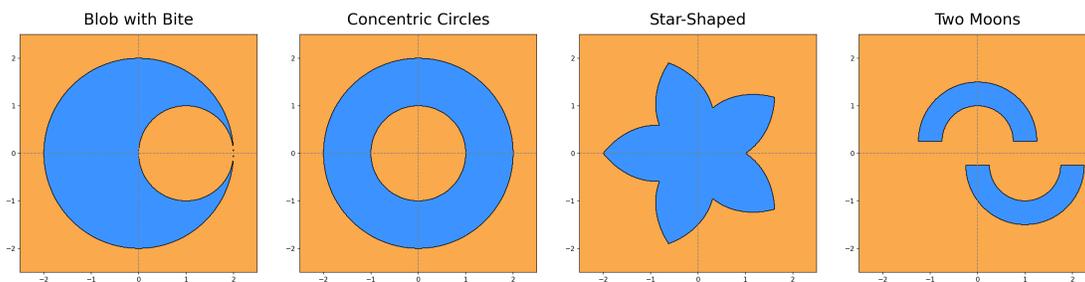


Figure 2: The nonconvex constraint sets tested in our constrained optimization settings.

Baselines. We compare FAB feasibility improvement against:

- **Projected Gradient Descent.** A classical optimization algorithm that iteratively takes a gradient descent step and then projects the current iterate onto the feasible set to enforce constraints.
- **Penalty.** A classical optimization approach that solves an unconstrained version of the problem, $\text{minimize}_x J(x) = f(x) + \mu P(x)$, where μ is a large penalty coefficient.
- **Augmented-Lagrangian.** A classical optimization method that combines penalty terms and Lagrange multipliers to handle constraints. The objective is to minimize $J(x, \lambda) = f(x) + \lambda^T c(x) + \frac{\rho}{2} \|c(x)\|^2$. We alternate between updating the primal variables x using estimated gradients and updating the dual variables λ .
- **Interior Point** (a.k.a. barrier method). A classical optimization algorithm that ensures feasibility by strictly penalizing the optimizer as it approaches constraint boundaries. Iterates are initialized within the feasible set, and a logarithmic barrier function $-\mu \sum \log(\hat{d}(x))$ is added to the objective, where $\hat{d}(x)$ is a smoothed proximity score (derived from local sampling). The barrier coefficient μ is annealed over time to guide the solution toward the optimum.
- **FSNet** (Nguyen & Donti, 2025). A learning-based method that combines NNs with an exact solution of an iterative constraint violation-minimization optimization problem.
- **Homeomorphic Projection** (Liang & Chen, 2025). A method that learns an invertible mapping between a hypersphere and a ball-homeomorphic constraint set, using a [post-hoc](#) bisection procedure to recover feasibility of neural network outputs.

Implementation details. We run experiments over 5 seeds with 300 problems per seed (trained 500 epochs with a batch size of 32, tested on 1,500 problems), totalling 18,000 optimization problems across the entire testing suite. All methods were run on a workstation equipped with one NVIDIA RTX 5090 GPU. We choose \mathcal{S} to be a 0.5-radius hypersphere. Hyperparameters are given in Appendix A.1.

Evaluation metrics. We evaluate the methods by their feasibility rates, optimality gaps, and inference times over the 4 constraint types and 3 objective types.

Results

Method	Quadratic			Linear			Dist. Min.			
	↑ Feas (%)	↓ Time (ms)	↓ Gap	↑ Feas (%)	↓ Time (ms)	↓ Gap	↑ Feas (%)	↓ Time (ms)	↓ Gap	
Classical	Projected Gradient	100.0 $\sigma=0.0$	38.73 $\sigma=28.68$	0.60 $\sigma=0.88$	100.0 $\sigma=0.0$	19.80 $\sigma=31.73$	1.04 $\sigma=1.11$	100.0 $\sigma=0.0$	32.34 $\sigma=49.68$	0.95 $\sigma=2.42$
	Penalty Method	59.7 $\sigma=49.0$	64.76 $\sigma=47.86$	0.97 $\sigma=1.39$	43.7 $\sigma=49.6$	38.72 $\sigma=52.87$	1.29 $\sigma=1.28$	55.1 $\sigma=49.7$	55.01 $\sigma=156.36$	5.37 $\sigma=6.47$
	Augmented Lagrangian	58.5 $\sigma=49.3$	71.64 $\sigma=64.04$	0.98 $\sigma=1.42$	44.5 $\sigma=49.7$	42.68 $\sigma=49.49$	1.33 $\sigma=1.29$	56.5 $\sigma=49.6$	45.49 $\sigma=33.47$	5.54 $\sigma=6.73$
	Interior Point	95.3 $\sigma=21.2$	61.82 $\sigma=35.51$	2.50 $\sigma=2.61$	92.6 $\sigma=26.2$	45.66 $\sigma=56.51$	1.63 $\sigma=1.66$	94.5 $\sigma=22.7$	40.50 $\sigma=41.82$	11.89 $\sigma=11.35$
ML-Based	FSNet	99.5 $\sigma=7.3$	2.66 $\sigma=11.09$	1.15 $\sigma=1.67$	98.7 $\sigma=11.2$	2.49 $\sigma=1.43$	2.09 $\sigma=1.98$	99.8 $\sigma=4.5$	2.31 $\sigma=1.33$	13.63 $\sigma=13.30$
	Homeomorphic Projection	100.0 $\sigma=0.0$	196.631 $\sigma=193.4$	4.37 $\sigma=4.28$	100.0 $\sigma=0.0$	196.247 $\sigma=195.8$	2.58 $\sigma=2.30$	100.0 $\sigma=0.0$	196.021 $\sigma=195.5$	13.36 $\sigma=12.73$
	FAB: 1 Decoder	100.0 $\sigma=0.0$	0.69 $\sigma=1.44$	0.53 $\sigma=0.71$	100.0 $\sigma=0.0$	0.50 $\sigma=0.48$	1.14 $\sigma=0.96$	100.0 $\sigma=0.0$	0.39 $\sigma=0.17$	2.89 $\sigma=4.26$
	FAB: 2 Decoders	62.0 $\sigma=48.5$	1.14 $\sigma=4.08$	0.52 $\sigma=0.69$	55.9 $\sigma=49.7$	0.59 $\sigma=0.43$	1.09 $\sigma=0.95$	86.0 $\sigma=34.7$	0.45 $\sigma=0.15$	2.59 $\sigma=4.23$
	FAB: 3 Decoders	95.6 $\sigma=20.5$	1.15 $\sigma=3.18$	0.53 $\sigma=0.70$	95.3 $\sigma=21.2$	0.60 $\sigma=0.35$	1.11 $\sigma=0.92$	94.6 $\sigma=22.6$	0.55 $\sigma=0.42$	3.59 $\sigma=4.49$
	FAB: Only phase 1	66.7 $\sigma=47.1$	0.77 $\sigma=0.34$	0.51 $\sigma=0.62$	64.1 $\sigma=48.0$	1.16 $\sigma=2.88$	1.06 $\sigma=0.83$	76.7 $\sigma=42.3$	0.56 $\sigma=0.13$	1.32 $\sigma=1.48$
	FAB: No hinge	100.0 $\sigma=0.0$	0.94 $\sigma=1.79$	0.63 $\sigma=0.79$	100.0 $\sigma=0.0$	1.15 $\sigma=1.58$	1.29 $\sigma=1.11$	100.0 $\sigma=0.0$	0.61 $\sigma=0.40$	6.47 $\sigma=5.28$
	FAB: No D_ξ	80.8 $\sigma=39.4$	0.74 $\sigma=0.50$	0.52 $\sigma=0.69$	80.6 $\sigma=39.5$	0.82 $\sigma=0.47$	1.11 $\sigma=0.92$	77.9 $\sigma=41.5$	0.73 $\sigma=0.58$	4.79 $\sigma=4.34$

Table 1: Mean and std.dev. for feasibility (%), time (ms), and optimality gap for each method: Blob w/ Bite.

As presented in Table 1 (as well as the tables in Appendix B, our methods achieve a substantial reduction in computation time, with sub-millisecond inference times. In addition, they also do well at finding feasible solutions even with disjoint or non-ball-homeomorphic sets. As shown in the 2 and 3 decoder ablations, their performance improves for sets with multiple components, as each decoder "specializes" in a distinct feasible region.

5 EXPERIMENTS ON SAFETY GYM (SAFE RL)

Safe RL problem formulation. Safe RL can be framed within our parametric optimization setup from Equation 1. We consider a discrete-time dynamical system where the state evolves according to:

$$s_{k+1} = f(s_k, u_k), \quad (10)$$

where $s_k \in \mathcal{X} \subseteq \mathbb{R}^m$ is the current state and $u_k \in \mathcal{Y} \subseteq \mathbb{R}^n$ is the control action at timestep k . In this context, the actions u_k are the decision variables, and the states s_k correspond to the problem parameters. The constraint set $\mathcal{C}(s_k)$ is state-dependent, where s_k acts as a parameter. An action u_k is considered safe, i.e., $u_k \in \mathcal{C}(s_k)$, if executing it from state s_k does not lead to an immediate constraint violation (e.g. collision). The objective is to learn a policy $\pi_\theta : \mathcal{X} \rightarrow \mathcal{Y}$, that maximizes the expected cumulative reward. The goal is, therefore, to solve for a policy π_θ such that for any given state s_k , the chosen action $u_k = \pi_\theta(s_k)$ is both optimal long-term and safe.

Algorithm	Reward Mean \uparrow	Reward Std \downarrow	Cost Mean \downarrow	Cost Std \downarrow	Time Mean (s) \downarrow
PPO_FAB	-1.12	1.11	24.32	37.84	2.75
PPO	13.26	14.05	167.46	87.06	2.47
PPO_LAG	2.24	5.10	54.10	64.50	2.40
TRPO	15.58	10.31	164.14	88.43	2.59
TRPO_LAG	2.37	8.46	89.04	187.67	2.78

Table 2: SafetyGym Goal Results

Algorithm	Reward Mean \uparrow	Reward Std \downarrow	Cost Mean \downarrow	Cost Std \downarrow	Time Mean (s) \downarrow
PPO_FAB	-0.07	0.52	7.70	37.51	3.71
PPO	0.41	3.11	59.86	120.18	3.78
PPO_LAG	0.60	1.58	31.34	58.17	4.09
TRPO	0.20	2.50	106.88	216.19	4.03
TRPO_LAG	-0.77	5.51	28.22	67.21	4.05

Table 3: SafetyGym Push Results

Environment. We evaluate our approach using the SafetyPointGoal2-v0 and SafetyPointPush2-v0 environments from the Safety Gymnasium benchmark suite (Ji et al., 2023). In the Goal task, an agent must navigate to a series of goal locations while avoiding randomly-placed hazards. In the Push task, an agent must navigate to and manipulate a movable box, pushing it toward a target location while avoiding randomly placed hazards. The state vector comprises simulated sensor readings (e.g., accelerometers, gyroscope, and a LiDAR-like sensor for hazard detection), while the action vector controls the agent’s forward/backward movement, as well as its rotational velocities.

Baselines. We compare FAB projections against:

- **PPO** (Schulman et al., 2017). A state-of-the-art unconstrained RL algorithm.
- **TRPO** (Schulman et al., 2015). Another robust unconstrained RL algorithm.
- **Lagrangian PPO, Lagrangian TRPO** (Stooke et al., 2020; Ray et al., 2019). Constrained versions of PPO and TRPO that use Lagrangian relaxation to penalize constraint violations during training.

Implementation details. For each environment, the autoencoder was trained on a dataset of 100,000 state-action pairs, approximately evenly balanced between safe and unsafe examples, which were collected by running a random policy in the environment. The performance of these algorithms was evaluated over 50 independent seeds. All methods were run on a workstation equipped with one NVIDIA RTX 5090 GPU. We choose \mathcal{S} to be a 0.5-radius hypersphere. Hyperparameters are given in Appendix A.2.

Evaluation metrics. We assess performance based on several metrics, averaged over the evaluation episodes: episode rewards, episode costs, and mean inference time. Episode reward is the cumulative rewards obtained by the agent in one episode (which is 1000 steps). Episode cost is the total number of constraint violations (i.e., entering hazard zones).

Results

As presented in 2 and 3, PPO_FAB achieves the lowest costs among the evaluated methods, substantially outperforming standard baselines like PPO (167.46, 59.86) and TRPO (164.14, 106.88). Furthermore, it also exhibits the highest levels of reliability with the lowest standard deviation in both cost and reward. However,

423 this focus on safety also corresponds to a more conservative reward-seeking policy, resulting in lower reward
424 means (-1.12, -0.07) compared to the other algorithms.
425

426 6 CONCLUSION

428 In this work, we introduced a novel-data driven method, Fast Autoencoder-Based (FAB) projections for
429 feasibility improvement, to address the critical challenge of enforcing nonconvex constraints in learning-
430 based systems. Our method leverages a specially trained autoencoder to learn an approximate projection
431 from a convex latent set to a general (e.g., potentially nonconvex or disjoint) feasible set. By structuring the
432 autoencoder’s latent space to correspond to a simple convex shape that is easy to map into, we can perform
433 a fast projection in this latent space at inference and decode the result back into a point that is highly likely
434 to be feasible.

435 On constrained optimization benchmarks, FAB consistently achieved near-perfect feasibility rates (often
436 approaching 100%) across a diverse set of challenging, nonconvex feasible sets. In the SafetyGym bench-
437 mark, it consistently achieved low costs, with the median cost being 0.0. Crucially, it achieved both of
438 these with inference times that were significantly faster than those of other methods. While methods with
439 formal guarantees can ensure exact feasibility, their computational overhead can be prohibitive for real-time
440 applications. Therefore, FAB presents a compelling alternative for latency-sensitive domains where rapid,
441 high-fidelity approximate projections may be sufficient.

442 Limitations of FAB include lack of hard guarantees, as well as its reliance on having a representative dataset
443 of feasible and infeasible points in order for autoencoder training to perform well. Future work includes
444 identifying ways to improve sample efficiency, distributional performance, and interpretability of the data-
445 driven mapping – e.g., through the use of specialized neural network structures such as input-convex neural
446 networks (Amos et al., 2017) or operator learning-based networks (Kovachki et al., 2024), or even simpler
447 parameterized functions – as well as exploring toolkits such as formal verification to better understand and
448 assess the behavior of the learned data-driven mapping after it is trained. Another future direction includes
449 exploring adaptive co-training of the data-driven feasibility improvement mapping with the neural network to
450 which it will be attached (rather than training the mapping fully separately), motivated by prior results on the
451 benefits of end-to-end learning (Cameron et al., 2022). Overall, we believe that the paradigm of data-driven
452 feasibility mapping offers a compelling yet underexplored middle ground between penalty methods and
453 exact feasibility enforcement, and we look forward to future work that investigates alternative approaches
454 within this space beyond the one presented here.

REFERENCES

- Brandon Amos. Tutorial on amortized optimization. *Foundations and Trends® in Machine Learning*, 16(5): 592–732, 2023.
- Brandon Amos, Lei Xu, and J Zico Kolter. Input convex neural networks. In *International conference on machine learning*, pp. 146–155. PMLR, 2017.
- Tao Bai, Jinqi Luo, Jun Zhao, Bihan Wen, and Qian Wang. Recent advances in adversarial training for adversarial robustness. *arXiv preprint arXiv:2102.01356*, 2021.
- Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011. URL https://web.stanford.edu/~boyd/papers/admm_distr_stats.pdf.
- Chris Cameron, Jason Hartford, Taylor Lundy, and Kevin Leyton-Brown. The perils of learning before optimizing. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pp. 3708–3715, 2022.
- Bingqing Chen, Priya L. Donti, Kyri Baker, J. Zico Kolter, and Mario Bergés. Enforcing policy feasibility constraints through differentiable projection for energy optimization. In *Proceedings of the Twelfth ACM International Conference on Future Energy Systems (e-Energy '21)*, pp. 199–210. ACM, 2021. ISBN 978-1-4503-8333-2. doi: 10.1145/3447555.3464874.
- Wenbo Chen, Seonho Park, Mathieu Tanneau, and Pascal Van Hentenryck. Learning optimization proxies for large-scale security-constrained economic dispatch. *Electric Power Systems Research*, 213:108566, 2022.
- Priya L. Donti, David Rolnick, and J Zico Kolter. DC3: A learning method for optimization with hard constraints. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=V1ZHVxJ6dSS>.
- Jim Douglas, Jr. and H. H. Rachford. On the numerical solution of heat conduction problems in two and three space variables. *Transactions of the American Mathematical Society*, 82(2):421–439, 1956.
- Ferdinando Fioretto, Pascal Van Hentenryck, Terrence WK Mak, Cuong Tran, Federico Baldo, and Michele Lombardi. Lagrangian duality for constrained deep learning. In *Machine learning and knowledge discovery in databases. applied data science and demo track: European conference, ECML pKDD 2020, Ghent, Belgium, September 14–18, 2020, proceedings, part v*, pp. 118–135. Springer, 2021.
- Javier Garcia and Fernando Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1):1437–1480, 2015.
- Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- Shangding Gu, Long Yang, Yali Du, Guang Chen, Florian Walter, Jun Wang, and Alois Knoll. A review of safe reinforcement learning: Methods, theories and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- Alexander Hans, Daniel Schneegaß, Anton Maximilian Schäfer, and Steffen Udfluft. Safe exploration for reinforcement learning. In *ESANN*, pp. 143–148, 2008.

- 517 Jiaming Ji, Borong Zhang, Jiayi Zhou, Xuehai Pan, Weidong Huang, Ruiyang Sun, Yiran Geng, Yifan
518 Zhong, Josef Dai, and Yaodong Yang. Safety gymnasium: A unified safe reinforcement learning bench-
519 mark. *Advances in Neural Information Processing Systems*, 36:18964–18993, 2023.
- 520 Nikola B Kovachki, Samuel Lanthaler, and Andrew M Stuart. Operator learning: Algorithms and anal-
521 ysis. In Siddhartha Mishra and Alex Townsend (eds.), *Numerical Analysis Meets Machine Learning*,
522 volume 25 of *Handbook of Numerical Analysis*, pp. 419–467. Elsevier, 2024. doi: [https://doi.org/](https://doi.org/10.1016/bs.hna.2024.05.009)
523 [10.1016/bs.hna.2024.05.009](https://doi.org/10.1016/bs.hna.2024.05.009). URL [https://www.sciencedirect.com/science/article/](https://www.sciencedirect.com/science/article/pii/S1570865924000097)
524 [pii/S1570865924000097](https://www.sciencedirect.com/science/article/pii/S1570865924000097).
- 525 Enming Liang and Minghua Chen. Efficient bisection projection to ensure NN solution feasibility for opti-
526 mization over general set. 2025. URL <https://openreview.net/forum?id=7TXdglI1g0>.
- 527 Enming Liang, Minghua Chen, and Steven H Low. Low complexity homeomorphic projection to ensure
528 neural-network solution feasibility for optimization over (non-) convex set. In *40th International Confer-*
529 *ence on Machine Learning (ICML 2023)*, pp. 20623–20649, 2023.
- 530 Enming Liang, Minghua Chen, and Steven H Low. Homeomorphic projection to ensure neural-network
531 solution feasibility for constrained optimization. *Journal of Machine Learning Research*, 25(329):1–55,
532 2024.
- 533 Youngjae Min and Navid Azizan. Hardnet: Hard-constrained neural networks with universal approximation
534 guarantees. *arXiv preprint arXiv:2410.10807*, 2024.
- 535 Philipp Nazari, Sebastian Damrich, and Fred A Hamprecht. Geometric autoencoders—what you see is what
536 you decode. *arXiv preprint arXiv:2306.17638*, 2023.
- 537 Hoang T Nguyen and Priya L Donti. FSNet: Feasibility-seeking neural network for constrained optimization
538 with guarantees. *arXiv preprint arXiv:2506.00362*, 2025.
- 539 Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer, 2006.
- 540 Xiang Pan, Minghua Chen, Tianyu Zhao, and Steven H Low. DeepOPF: A feasibility-optimized deep neural
541 network approach for AC optimal power flow problems. *IEEE Systems Journal*, 17(1):673–683, 2022.
- 542 Tu-Hoa Pham, Giovanni De Magistris, and Ryuki Tachibana. Optlayer-practical constrained optimization
543 for deep reinforcement learning in the real world. In *2018 IEEE International Conference on Robotics*
544 *and Automation (ICRA)*, pp. 6236–6243. IEEE, 2018.
- 545 Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learn-
546 ing framework for solving forward and inverse problems involving nonlinear partial differential equations.
547 *Journal of Computational physics*, 378:686–707, 2019.
- 548 Alex Ray, Joshua Achiam, and Dario Amodei. Benchmarking safe exploration in deep reinforcement learn-
549 ing. *arXiv preprint arXiv:1910.01708*, 2019.
- 550 John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy
551 optimization. In *International conference on machine learning*, pp. 1889–1897. PMLR, 2015.
- 552 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy opti-
553 mization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- 554 Adam Stooke, Joshua Achiam, and Pieter Abbeel. Responsive safety in reinforcement learning by pid
555 lagrangian methods. In *International Conference on Machine Learning (ICML)*, 2020.
- 556
- 557
- 558
- 559
- 560
- 561
- 562
- 563

564 Jesus Tordesillas, Jonathan P How, and Marco Hutter. Rayen: Imposition of hard convex constraints on
565 neural networks. *arXiv preprint arXiv:2307.08336*, 2023.

566 Pascal Van Hentenryck. Optimization learning. *arXiv preprint arXiv:2501.03443*, 2025.

568 Ahmed S Zamzam and Kyri Baker. Learning optimal solutions for extremely fast AC optimal power flow.
569 In *2020 IEEE International Conference on Communications, Control, and Computing Technologies for*
570 *Smart Grids (SmartGridComm)*, pp. 1–6. IEEE, 2020.

571

572

573

A ADDITIONAL EXPERIMENTAL DETAILS

574

575

A.1 CONSTRAINED OPTIMIZATION HYPERPARAMETERS

576

Phase 1 Training (phase1_training.py)

577

578

Batch size 256

579

Epochs 500

580

Learning rate 0.001

581

Optimizer Adam

582

Validation split 0.2

583

Hidden dimension 64

584

Number of decoders 1

585

Number of samples 60000

586

Phase 2 Training (phase2_training.py)

587

588

Batch size 256

589

Epochs 150

590

Learning rates AE=0.0005, Discriminator=0.001

591

Loss weights $\lambda_{\text{recon}} = 1.0$, $\lambda_{\text{feasibility}} = 1.0$, $\lambda_{\text{latent}} = 1.0$, $\lambda_{\text{hinge}} = 0.1$, $\lambda_{\text{geometric}} = 0.1$

592

Optimizer Adam

593

Discriminator type "absolute"

594

Critic steps 3

595

Validation split 0.2 (implicit from train_test_split)

596

Normalization enabled

597

Hidden dimension 64

598

Number of decoders 1

599

600

A.2 SAFETY GYM HYPERPARAMETERS

601

Phase 1 Training (safety-gym/phase1_training.py)

602

Batch size: 256

603

Epochs: 500

611 **Learning rate:** 0.001
612 **Optimizer:** Adam
613 **Validation split:** 0.2
614 **Reconstruction weight:** 1.0
615 **Hidden dimension:** 64
616 **Number of decoders:** 1
617 **Number of samples:** 100000
618 **State dimension:** 60 (safety_gym)
619 **Action dimension:** 2 (safety_gym)
620 **Latent dimension:** equal to action dimension
621 **Phase 2 Training (safety-gym/phase2_training.py)**
622 **Batch size:** 256
623 **Epochs:** 100
624 **Learning rates:** AE=0.0001, Discriminator=0.0002
625 **Loss weights:** $\lambda_{recon} = 1.0$, $\lambda_{feasibility} = 0.5$, $\lambda_{latent} = 0.5$, $\lambda_{hinge} = 0.5$, $\lambda_{geometric} = 0.1$
626 **Optimizer:** Adam
627 **Discriminator type:** "absolute"
628 **Critic steps:** 1
629 **Validation split:** 0.2
630 **Hidden dimension:** 64
631 **Number of decoders:** 1
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657

B EXPERIMENTAL TABLES

Method	Quadratic			Linear			Dist. Min.			
	↑ Feas (%)	↓ Time (ms)	↓ Gap	↑ Feas (%)	↓ Time (ms)	↓ Gap	↑ Feas (%)	↓ Time (ms)	↓ Gap	
Classical	Projected Gradient	80.5	79.07	1.18	76.3	48.23	0.90	79.2	69.79	4.65
		$\sigma=39.6$	$\sigma=91.65$	$\sigma=1.58$	$\sigma=42.5$	$\sigma=52.86$	$\sigma=0.99$	$\sigma=40.6$	$\sigma=25.09$	$\sigma=5.87$
	Penalty Method	20.4	78.39	1.36	10.4	61.58	1.52	14.6	38.67	6.39
		$\sigma=40.3$	$\sigma=71.27$	$\sigma=1.67$	$\sigma=30.5$	$\sigma=111.21$	$\sigma=1.49$	$\sigma=35.3$	$\sigma=10.02$	$\sigma=7.07$
	Augmented Lagrangian	20.9	85.71	1.33	10.9	56.98	1.56	15.7	42.58	6.43
	$\sigma=40.7$	$\sigma=119.17$	$\sigma=1.59$	$\sigma=31.1$	$\sigma=62.66$	$\sigma=1.56$	$\sigma=36.3$	$\sigma=6.69$	$\sigma=6.70$	
Interior Point	78.3	82.80	1.64	77.3	46.37	1.03	79.9	42.00	6.12	
	$\sigma=41.2$	$\sigma=110.99$	$\sigma=2.02$	$\sigma=41.9$	$\sigma=33.68$	$\sigma=1.03$	$\sigma=40.0$	$\sigma=12.03$	$\sigma=6.41$	
ML-Based	FSNet	98.7	4.42	1.13	98.3	3.40	0.96	99.3	4.03	5.30
		$\sigma=11.5$	$\sigma=2.77$	$\sigma=1.45$	$\sigma=12.8$	$\sigma=2.08$	$\sigma=1.01$	$\sigma=8.5$	$\sigma=1.67$	$\sigma=5.51$
	Homeomorphic Projection	73.9	247	1.74	76.1	240	1.67	76.1	243	8.12
		$\sigma=43.9$	$\sigma=105$	$\sigma=1.90$	$\sigma=42.6$	$\sigma=104$	$\sigma=1.48$	$\sigma=42.6$	$\sigma=105$	$\sigma=7.80$
	FAB: 1 Decoder	96.8	0.55	0.85	83.5	0.43	0.73	88.3	0.41	4.22
		$\sigma=17.6$	$\sigma=0.40$	$\sigma=1.16$	$\sigma=37.1$	$\sigma=0.27$	$\sigma=0.73$	$\sigma=32.2$	$\sigma=0.20$	$\sigma=4.15$
	FAB: 2 Decoders	100.0	0.59	1.31	100.0	0.55	0.94	100.0	0.50	5.64
		$\sigma=0.0$	$\sigma=0.35$	$\sigma=1.36$	$\sigma=0.0$	$\sigma=0.37$	$\sigma=0.89$	$\sigma=0.0$	$\sigma=0.25$	$\sigma=5.73$
	FAB: 3 Decoders	100.0	0.80	1.02	100.0	0.62	0.90	100.0	0.60	5.44
		$\sigma=0.0$	$\sigma=1.03$	$\sigma=1.15$	$\sigma=0.0$	$\sigma=0.36$	$\sigma=0.84$	$\sigma=0.0$	$\sigma=0.32$	$\sigma=4.84$
FAB: Only phase 1	82.5	0.64	0.86	69.9	1.07	0.72	64.1	0.54	3.83	
	$\sigma=38.0$	$\sigma=0.14$	$\sigma=1.21$	$\sigma=45.9$	$\sigma=4.12$	$\sigma=0.69$	$\sigma=48.0$	$\sigma=0.09$	$\sigma=4.07$	
FAB: No hinge	28.4	0.63	1.13	18.8	0.76	0.85	28.1	0.56	5.14	
	$\sigma=45.1$	$\sigma=0.16$	$\sigma=1.23$	$\sigma=39.1$	$\sigma=0.84$	$\sigma=0.80$	$\sigma=44.9$	$\sigma=0.16$	$\sigma=4.96$	
FAB: No D_ξ	75.1	0.64	0.86	54.4	0.67	0.70	13.6	0.55	3.61	
	$\sigma=43.2$	$\sigma=0.22$	$\sigma=1.21$	$\sigma=49.8$	$\sigma=0.31$	$\sigma=0.69$	$\sigma=34.3$	$\sigma=0.19$	$\sigma=4.00$	

Table 4: Mean and std. dev. for feasibility (%), time (ms), and optimality gap for each method: Two Moons.

Method	Quadratic			Linear			Dist. Min.			
	↑ Feas (%)	↓ Time (ms)	↓ Gap	↑ Feas (%)	↓ Time (ms)	↓ Gap	↑ Feas (%)	↓ Time (ms)	↓ Gap	
Classical	Projected Gradient	100.0 $\sigma=0.0$	38.41 $\sigma=39.50$	0.34 $\sigma=0.41$	100.0 $\sigma=0.0$	18.99 $\sigma=10.97$	0.78 $\sigma=0.69$	100.0 $\sigma=0.0$	31.18 $\sigma=27.47$	1.01 $\sigma=1.86$
	Penalty Method	75.8 $\sigma=42.8$	61.37 $\sigma=33.75$	0.87 $\sigma=1.32$	38.0 $\sigma=48.5$	40.03 $\sigma=33.05$	1.17 $\sigma=1.15$	57.3 $\sigma=49.5$	39.41 $\sigma=13.16$	4.84 $\sigma=5.99$
	Augmented Lagrangian	75.0 $\sigma=43.3$	70.70 $\sigma=74.02$	0.89 $\sigma=1.38$	40.4 $\sigma=49.1$	39.49 $\sigma=20.97$	1.19 $\sigma=1.10$	58.1 $\sigma=49.3$	48.25 $\sigma=32.70$	4.99 $\sigma=6.17$
	Interior Point	96.6 $\sigma=18.1$	70.65 $\sigma=75.77$	1.77 $\sigma=1.99$	92.9 $\sigma=25.7$	40.25 $\sigma=21.66$	1.45 $\sigma=1.46$	95.0 $\sigma=21.8$	53.22 $\sigma=62.60$	9.43 $\sigma=8.70$
ML-Based	FSNet	99.9 $\sigma=3.6$	3.13 $\sigma=3.00$	1.03 $\sigma=1.91$	99.7 $\sigma=5.2$	3.76 $\sigma=2.06$	1.05 $\sigma=1.04$	99.7 $\sigma=5.2$	5.41 $\sigma=11.90$	7.24 $\sigma=7.61$
	Homeomorphic Projection	100.0 $\sigma=0.0$	128.653 $\sigma=48.4$	3.70 $\sigma=3.35$	100.0 $\sigma=0.0$	126.350 $\sigma=48.9$	2.35 $\sigma=2.07$	100.0 $\sigma=0.0$	126.455 $\sigma=48.9$	12.49 $\sigma=11.84$
	FAB: 1 Decoder	99.9 $\sigma=2.6$	0.56 $\sigma=0.45$	0.34 $\sigma=0.40$	99.9 $\sigma=2.6$	0.60 $\sigma=1.61$	0.77 $\sigma=0.65$	95.4 $\sigma=20.9$	0.64 $\sigma=0.93$	2.01 $\sigma=1.96$
	FAB: 2 Decoder	100.0 $\sigma=0.0$	0.60 $\sigma=0.32$	0.37 $\sigma=0.43$	99.5 $\sigma=7.3$	0.49 $\sigma=0.26$	0.73 $\sigma=0.59$	94.7 $\sigma=22.3$	0.71 $\sigma=0.83$	2.30 $\sigma=2.28$
	FAB: 3 Decoder	99.9 $\sigma=2.6$	0.65 $\sigma=0.29$	0.34 $\sigma=0.39$	99.9 $\sigma=3.6$	0.55 $\sigma=0.24$	0.75 $\sigma=0.62$	94.5 $\sigma=22.9$	0.93 $\sigma=2.39$	2.02 $\sigma=2.00$
	FAB: Only phase 1	100.0 $\sigma=0.0$	0.61 $\sigma=0.08$	0.34 $\sigma=0.41$	100.0 $\sigma=0.0$	0.54 $\sigma=0.08$	0.76 $\sigma=0.65$	100.0 $\sigma=0.0$	0.53 $\sigma=0.08$	3.05 $\sigma=3.04$
	FAB: No hinge	100.0 $\sigma=0.0$	0.60 $\sigma=0.07$	0.38 $\sigma=0.49$	100.0 $\sigma=0.0$	0.54 $\sigma=0.07$	0.83 $\sigma=0.75$	100.0 $\sigma=0.0$	0.52 $\sigma=0.04$	5.00 $\sigma=4.27$
	FAB: No D_ξ	100.0 $\sigma=0.0$	0.86 $\sigma=2.05$	0.35 $\sigma=0.40$	100.0 $\sigma=0.0$	0.54 $\sigma=0.07$	0.78 $\sigma=0.65$	96.8 $\sigma=17.6$	0.52 $\sigma=0.04$	1.62 $\sigma=1.64$

Table 5: Mean and std. dev. for feasibility (%), time (ms), and optimality gap for each method: Star Shaped constraint family. 5 seeds, 300 problems each per problem per objective type.

Method	Quadratic			Linear			Dist. Min.			
	↑ Feas (%)	↓ Time (ms)	↓ Gap	↑ Feas (%)	↓ Time (ms)	↓ Gap	↑ Feas (%)	↓ Time (ms)	↓ Gap	
Classical	Projected Gradient	100.0	31.24	0.96	100.0	14.83	1.16	100.0	24.36	0.40
		$\sigma=0.0$	$\sigma=13.55$	$\sigma=1.27$	$\sigma=0.0$	$\sigma=14.92$	$\sigma=1.16$	$\sigma=0.0$	$\sigma=17.34$	$\sigma=0.88$
	Penalty Method	35.5	56.68	1.12	42.5	39.15	1.35	49.0	40.05	5.55
		$\sigma=47.8$	$\sigma=44.47$	$\sigma=1.41$	$\sigma=49.4$	$\sigma=63.11$	$\sigma=1.30$	$\sigma=50.0$	$\sigma=31.26$	$\sigma=6.60$
Augmented Lagrangian	37.2	57.59	1.16	42.1	52.09	1.39	49.3	43.93	5.70	
	$\sigma=48.3$	$\sigma=43.84$	$\sigma=1.43$	$\sigma=49.4$	$\sigma=85.00$	$\sigma=1.30$	$\sigma=50.0$	$\sigma=46.15$	$\sigma=6.91$	
Interior Point	93.4	54.53	2.67	89.9	39.76	1.78	92.6	44.17	11.89	
	$\sigma=24.8$	$\sigma=34.56$	$\sigma=2.71$	$\sigma=30.2$	$\sigma=29.55$	$\sigma=1.79$	$\sigma=26.2$	$\sigma=38.55$	$\sigma=11.28$	
ML-Based	FSNet	99.9	1.75	1.29	98.5	3.04	1.72	99.1	2.03	9.72
		$\sigma=2.6$	$\sigma=1.01$	$\sigma=1.72$	$\sigma=12.0$	$\sigma=6.11$	$\sigma=2.01$	$\sigma=9.3$	$\sigma=0.97$	$\sigma=14.52$
	Homeomorphic Projection	93.7	166	4.83	93.4	168	2.56	93.4	173	14.08
		$\sigma=24.4$	$\sigma=86$	$\sigma=4.31$	$\sigma=24.8$	$\sigma=87$	$\sigma=2.32$	$\sigma=24.8$	$\sigma=90$	$\sigma=13.51$
	FAB: 1 Decoder	100.0	0.45	1.75	100.0	0.38	1.48	99.9	0.48	5.72
		$\sigma=0.0$	$\sigma=0.21$	$\sigma=1.67$	$\sigma=0.0$	$\sigma=0.13$	$\sigma=1.36$	$\sigma=2.6$	$\sigma=0.36$	$\sigma=7.72$
	FAB: 2 Decoders	0.2	0.50	1.49	8.3	0.49	1.52	81.4	0.52	1.54
		$\sigma=4.5$	$\sigma=0.17$	$\sigma=1.44$	$\sigma=27.5$	$\sigma=0.30$	$\sigma=1.33$	$\sigma=38.9$	$\sigma=0.30$	$\sigma=1.53$
	FAB: 3 Decoders	0.7	0.86	1.50	3.1	0.61	1.60	82.5	0.71	1.25
		$\sigma=8.1$	$\sigma=1.94$	$\sigma=1.43$	$\sigma=17.4$	$\sigma=0.44$	$\sigma=1.47$	$\sigma=38.0$	$\sigma=1.21$	$\sigma=1.25$
FAB: Only phase 1	0.1	0.77	0.89	2.0	0.54	1.12	25.9	0.51	4.61	
	$\sigma=3.6$	$\sigma=2.51$	$\sigma=0.82$	$\sigma=14.0$	$\sigma=0.11$	$\sigma=0.81$	$\sigma=43.8$	$\sigma=0.04$	$\sigma=3.63$	
FAB: No hinge	1.0	0.74	1.30	1.9	0.58	1.48	68.7	0.51	7.08	
	$\sigma=9.9$	$\sigma=0.57$	$\sigma=1.29$	$\sigma=13.5$	$\sigma=0.21$	$\sigma=1.29$	$\sigma=46.4$	$\sigma=0.03$	$\sigma=7.85$	
FAB: No D_ξ	1.7	0.74	1.02	1.5	0.72	1.23	82.7	0.51	1.58	
	$\sigma=12.8$	$\sigma=0.67$	$\sigma=0.92$	$\sigma=12.3$	$\sigma=2.67$	$\sigma=0.96$	$\sigma=37.9$	$\sigma=0.03$	$\sigma=2.15$	

Table 6: Mean and std. dev. for feasibility (%), time (ms), and optimality gap for each method: Concentric Circles constraint family. 5 seeds, 300 problems each per problem per objective type.

Method	Quadratic			Linear			Dist. Min.			
	↑ Feas (%)	↓ Time (ms)	↓ Gap	↑ Feas (%)	↓ Time (ms)	↓ Gap	↑ Feas (%)	↓ Time (ms)	↓ Gap	
Incomplete Coverage	Cov-10	100.0	1.12	0.82	100.0	0.87	1.53	100.0	0.85	10.25
		$\sigma=0.0$	$\sigma=1.10$	$\sigma=1.02$	$\sigma=0.0$	$\sigma=0.04$	$\sigma=1.39$	$\sigma=0.0$	$\sigma=0.04$	$\sigma=7.33$
	Cov-25	100.0	1.08	0.81	100.0	0.86	1.52	100.0	0.85	10.15
		$\sigma=0.0$	$\sigma=0.09$	$\sigma=1.01$	$\sigma=0.0$	$\sigma=0.01$	$\sigma=1.39$	$\sigma=0.0$	$\sigma=0.03$	$\sigma=7.20$
Cov-50	100.0	1.08	0.80	100.0	0.86	1.47	100.0	0.85	9.15	
	$\sigma=0.0$	$\sigma=0.09$	$\sigma=0.98$	$\sigma=0.0$	$\sigma=0.01$	$\sigma=1.30$	$\sigma=0.0$	$\sigma=0.03$	$\sigma=6.97$	
Cov-75	100.0	1.08	0.81	100.0	0.86	1.52	100.0	0.85	9.46	
	$\sigma=0.0$	$\sigma=0.09$	$\sigma=1.01$	$\sigma=0.0$	$\sigma=0.03$	$\sigma=1.38$	$\sigma=0.0$	$\sigma=0.03$	$\sigma=6.86$	
Decoder capacity	W32-D2	99.5	1.09	0.62	98.5	0.79	1.21	72.5	0.80	6.36
		$\sigma=6.8$	$\sigma=1.48$	$\sigma=0.80$	$\sigma=12.0$	$\sigma=0.03$	$\sigma=1.03$	$\sigma=44.6$	$\sigma=0.03$	$\sigma=5.14$
	W32-D4	98.6	1.11	0.67	98.5	0.86	1.35	77.3	0.88	9.52
		$\sigma=11.8$	$\sigma=0.08$	$\sigma=0.88$	$\sigma=12.3$	$\sigma=0.03$	$\sigma=1.23$	$\sigma=41.9$	$\sigma=0.03$	$\sigma=7.37$
	W32-D6	99.4	1.17	0.72	98.7	0.93	1.20	96.8	0.95	6.54
		$\sigma=7.7$	$\sigma=0.08$	$\sigma=0.81$	$\sigma=11.5$	$\sigma=0.03$	$\sigma=1.06$	$\sigma=17.6$	$\sigma=0.03$	$\sigma=5.17$
	W64-D2	100.0	1.03	0.58	100.0	0.79	1.17	100.0	0.80	6.19
		$\sigma=0.0$	$\sigma=0.07$	$\sigma=0.77$	$\sigma=0.0$	$\sigma=0.03$	$\sigma=0.99$	$\sigma=0.0$	$\sigma=0.02$	$\sigma=5.09$
W64-D6	67.3	1.16	0.76	68.6	0.93	1.42	52.2	0.94	7.10	
	$\sigma=46.9$	$\sigma=0.08$	$\sigma=0.93$	$\sigma=46.4$	$\sigma=0.03$	$\sigma=1.28$	$\sigma=50.0$	$\sigma=0.02$	$\sigma=5.50$	
W128-D2	99.9	1.04	0.57	99.7	0.79	1.22	95.3	0.80	6.13	
	$\sigma=3.7$	$\sigma=0.07$	$\sigma=0.76$	$\sigma=5.2$	$\sigma=0.03$	$\sigma=1.03$	$\sigma=21.2$	$\sigma=0.02$	$\sigma=5.08$	
W128-D4	100.0	1.10	0.82	100.0	0.86	1.51	100.0	0.87	9.70	
	$\sigma=0.0$	$\sigma=0.07$	$\sigma=0.98$	$\sigma=0.0$	$\sigma=0.03$	$\sigma=1.36$	$\sigma=0.0$	$\sigma=0.02$	$\sigma=7.01$	
W128-D6	100.0	1.16	0.82	100.0	0.93	1.55	100.0	0.94	10.26	
	$\sigma=0.0$	$\sigma=0.08$	$\sigma=1.02$	$\sigma=0.0$	$\sigma=0.03$	$\sigma=1.41$	$\sigma=0.0$	$\sigma=0.06$	$\sigma=7.38$	

Table 7: Mean and std. dev. for feasibility (%), time (ms), and optimality gap for each ablation: Blob with Bite.

In tables 7-8, we provide ablation results for incomplete coverage and decoder capacity. The incomplete coverage ablations test the performance of our method when the training data only consisted of part of the constraint set (10%, 25%, 50%, and 75%). The decoder capacity ablations test the performance of multiple decoder depth and width configurations. W indicates decoder width, while D indicates decoder depth.

Method	Quadratic			Linear			Dist. Min.			
	↑ Feas (%)	↓ Time (ms)	↓ Gap	↑ Feas (%)	↓ Time (ms)	↓ Gap	↑ Feas (%)	↓ Time (ms)	↓ Gap	
Incomplete Coverage	Cov-10	100.0	1.06	1.30	100.0	0.87	1.43	100.0	0.88	9.57
		$\sigma=0.0$	$\sigma=0.11$	$\sigma=1.25$	$\sigma=0.0$	$\sigma=0.02$	$\sigma=1.40$	$\sigma=0.0$	$\sigma=0.03$	$\sigma=7.80$
	Cov-25	100.0	1.04	1.21	100.0	0.86	1.00	100.0	0.87	5.31
		$\sigma=0.0$	$\sigma=0.09$	$\sigma=1.36$	$\sigma=0.0$	$\sigma=0.01$	$\sigma=0.99$	$\sigma=0.0$	$\sigma=0.02$	$\sigma=4.25$
Cov-50	100.0	1.05	0.84	100.0	0.86	0.92	100.0	0.87	5.93	
	$\sigma=0.0$	$\sigma=0.09$	$\sigma=0.93$	$\sigma=0.0$	$\sigma=0.01$	$\sigma=0.91$	$\sigma=0.0$	$\sigma=0.03$	$\sigma=5.43$	
Cov-75	100.0	1.05	0.75	100.0	0.86	0.82	100.0	0.87	4.72	
	$\sigma=0.0$	$\sigma=0.09$	$\sigma=0.78$	$\sigma=0.0$	$\sigma=0.01$	$\sigma=0.75$	$\sigma=0.0$	$\sigma=0.03$	$\sigma=4.39$	
Decoder capacity	W32-D2	100.0	1.01	0.41	100.0	0.81	0.76	100.0	0.81	3.66
		$\sigma=0.0$	$\sigma=0.11$	$\sigma=0.48$	$\sigma=0.0$	$\sigma=0.03$	$\sigma=0.66$	$\sigma=0.0$	$\sigma=0.03$	$\sigma=3.48$
	W32-D4	100.0	1.07	0.37	100.0	0.88	0.79	100.0	0.88	3.55
		$\sigma=0.0$	$\sigma=0.09$	$\sigma=0.44$	$\sigma=0.0$	$\sigma=0.05$	$\sigma=0.68$	$\sigma=0.0$	$\sigma=0.04$	$\sigma=3.46$
	W32-D6	100.0	1.13	0.39	100.0	0.95	0.77	100.0	0.95	3.77
		$\sigma=0.0$	$\sigma=0.09$	$\sigma=0.46$	$\sigma=0.0$	$\sigma=0.03$	$\sigma=0.66$	$\sigma=0.0$	$\sigma=0.02$	$\sigma=3.58$
	W64-D2	100.0	1.00	0.39	100.0	0.80	0.76	100.0	0.81	3.74
		$\sigma=0.0$	$\sigma=0.09$	$\sigma=0.46$	$\sigma=0.0$	$\sigma=0.02$	$\sigma=0.65$	$\sigma=0.0$	$\sigma=0.02$	$\sigma=3.54$
	W64-D6	100.0	1.14	0.87	100.0	0.95	1.06	100.0	0.95	7.90
		$\sigma=0.0$	$\sigma=0.09$	$\sigma=1.10$	$\sigma=0.0$	$\sigma=0.03$	$\sigma=1.11$	$\sigma=0.0$	$\sigma=0.02$	$\sigma=7.51$
W128-D2	100.0	0.99	0.38	100.0	0.81	0.77	100.0	0.81	3.62	
	$\sigma=0.0$	$\sigma=0.09$	$\sigma=0.44$	$\sigma=0.0$	$\sigma=0.02$	$\sigma=0.67$	$\sigma=0.0$	$\sigma=0.02$	$\sigma=3.48$	
W128-D4	100.0	1.05	0.38	100.0	0.88	0.83	100.0	0.88	3.51	
	$\sigma=0.0$	$\sigma=0.09$	$\sigma=0.44$	$\sigma=0.0$	$\sigma=0.01$	$\sigma=0.73$	$\sigma=0.0$	$\sigma=0.02$	$\sigma=3.34$	
W128-D6	100.0	1.14	0.37	100.0	0.96	0.74	100.0	0.95	3.40	
	$\sigma=0.0$	$\sigma=0.09$	$\sigma=0.44$	$\sigma=0.0$	$\sigma=0.01$	$\sigma=0.63$	$\sigma=0.0$	$\sigma=0.02$	$\sigma=3.39$	

Table 8: Mean and std. dev. for feasibility (%), time (ms), and optimality gap for each ablation: Star Shaped.

C STATEMENT ON USE OF LARGE LANGUAGE MODELS (LLMs)

LLMs did *not* play a significant role in research ideation and/or writing for this paper. Perplexity was used to facilitate literature discovery as part of the research process, alongside other (non-LLM-based) tools. LLMs were also used to aid in sentence-level rewording for a small number of sentences during the preparation of this manuscript.