

---

# Learning to Solve the Constrained Most Probable Explanation Task in Probabilistic Graphical Models

---

Shivvrat Arya\*

The University of Texas at Dallas

Tahrima Rahman\*

The University of Texas at Dallas

Vibhav Gogate

The University of Texas at Dallas

## Abstract

We propose a *self-supervised learning* approach for solving the following constrained optimization task in log-linear models or Markov networks. Let  $f$  and  $g$  be two log-linear models defined over the sets  $\mathbf{X}$  and  $\mathbf{Y}$  of random variables respectively. Given an assignment  $\mathbf{x}$  to all variables in  $\mathbf{X}$  (evidence) and a real number  $q$ , the constrained most probable explanation (CMPE) task seeks to find an assignment  $\mathbf{y}$  to all variables in  $\mathbf{Y}$  such that  $f(\mathbf{x}, \mathbf{y})$  is maximized and  $g(\mathbf{x}, \mathbf{y}) \leq q$ . In our proposed self-supervised approach, given assignments  $\mathbf{x}$  to  $\mathbf{X}$  (data), we train a deep neural network that learns to output near-optimal solutions to the CMPE problem without requiring access to any pre-computed solutions. The key idea in our approach is to use first principles and approximate inference methods for CMPE to derive novel loss functions that seek to push infeasible solutions towards feasible ones and feasible solutions towards optimal ones. We analyze the properties of our proposed method and experimentally demonstrate its efficacy on several benchmark problems.

## 1 INTRODUCTION

Probabilistic graphical models (PGMs) such as Bayesian and Markov networks (Koller and Friedman, 2009; Darwiche, 2009) compactly represent joint probability distributions over random variables by factorizing the distribution according to a graph structure that encodes conditional independence among the variables. Once learned from data, these models can be used to answer various queries, such as computing the marginal probability distribution over a

subset of variables (MAR) and finding the most likely assignment to all unobserved variables, which is referred to as the most probable explanation (MPE) task.

Recently, Rouhani et al. (2020) proposed an extension to the MPE task in PGMs by introducing constraints. More specifically, given two PGMs  $f$  and  $g$  defined over the set of random variables  $\mathbf{X}$  and a real number  $q$ , the constrained most probable explanation (CMPE) task seeks to find the most likely state  $\mathbf{X} = \mathbf{x}$  w.r.t.  $f$  such that the constraint  $g(\mathbf{x}) \leq q$  is satisfied. Even though both MPE and CMPE are NP-hard in general, CMPE is considerably more difficult to solve in practice than MPE. Notably, CMPE is NP-hard even on PGMs having no edges, such as zero treewidth or independent PGMs, while MPE can be solved in linear time. Rouhani et al. (2020) and later Rahman et al. (2021) showed that several probabilistic inference queries are special cases of CMPE. This includes queries such as finding the decision preserving most probable explanation (Choi et al., 2012), finding the nearest assignment (Rouhani et al., 2018) and robust estimation (Darwiche and Hirth, 2023, 2020).

Our interest in the CMPE task is motivated by its extensive applicability to various *neuro-symbolic inference* tasks. Many of these tasks can be viewed as specific instances of CMPE. Specifically, when  $f(\mathbf{x})$  represents a function encoded by a neural network and  $g(\mathbf{x}) \leq q$  signifies particular symbolic or weighted constraints that the neural network must adhere to, the neuro-symbolic inference task involves determining the most likely prediction with respect to  $f$  while ensuring that the constraint  $g(\mathbf{x}) \leq q$  is satisfied. Another notable application of CMPE involves transferring abstract knowledge and inferences from simulations to real-world contexts. For example, in robotics, numerous simulations can be employed to instruct the robot on various aspects, such as object interactions, robot-world interactions, and underlying physical principles, encapsulating this abstract knowledge within the constraint  $g(\mathbf{x}) \leq q$ . Subsequently, with a neural network  $f$  trained on a limited amount of real-world data, characterized by richer feature sets and objectives,  $g$  can be used to reinforce the predictions made by  $f$ , ensuring that the robot identifies the most likely prediction with respect to  $f$  while satisfying the constraint  $g(\mathbf{x}) \leq q$ . This strategy enhances the reliability of

---

\*These authors contributed equally to this work.

the robot’s predictions and underscores the practical significance of CMPE.

In this paper, we explore novel machine learning (ML) approaches for solving the CMPE task, drawing inspiration from recent developments in *learning to optimize* (Donti et al., 2021; Fioretto et al., 2020; Park and Van Hentenryck, 2022; Zamzam and Baker, 2019). The main idea in these works is to train a deep neural network that takes the parameters, observations, etc. of a constrained optimization problem as *input* and outputs a near-optimal solution to the optimization problem.

In practice, a popular approach for solving optimization problems is to use search-based solvers such as Gurobi and SCIP. However, a drawback of these off-the-shelf solvers is their inability to efficiently solve large problems, especially those with dense global constraints, such as the CMPE problem. In contrast, neural networks are efficient because once trained, the time complexity of solving an optimization problem using them scales linearly with the network’s size. This attractive property has also driven their application in solving probabilistic inference tasks such as MAR and MPE inference (Gilmer et al., 2017; Kuck et al., 2020; Zhang et al., 2020; Satorras and Welling, 2021). However, all of these works require access to exact inference techniques in order to train the neural network. As a result, they are feasible only for small graphical models on which exact inference is tractable. Recently, Cui et al. (2022) proposed to solve the MPE task by training a variational distribution that is parameterized by a neural network in a self-supervised manner (without requiring access to exact inference methods). To the best of our knowledge, there is no prior work on using neural networks for solving the CMPE problem.

In this paper, we propose a new self-supervised approach for training neural networks which takes observations or evidence as input and outputs a near optimal solution to the CMPE task. Existing self-supervised approaches (Fioretto et al., 2020; Park and Van Hentenryck, 2022) in the *learning to optimize* literature either relax the constrained objective function using Lagrangian relaxation and then use the Lagrangian dual as a loss function or use the Augmented Lagrangian method. We show that these methods can be easily adapted to solve the CMPE task. Unfortunately, an issue with them is that an optimal solution to the Lagrangian dual is not guaranteed to be an optimal solution to the CMPE task (because of the non-convexity of CMPE, there is a duality gap). To address this issue, we propose a new loss function based on first principles and show that an optimal solution to the loss function is also an optimal solution to the CMPE task. Moreover, our new loss function has several desirable properties, which include: (a) during training, when the constraint is violated, it focuses on decreasing the strength of the violation, and (b) when constraints are not violated, it focuses on increasing the value of the objective function associated with the CMPE task.

We conducted a comprehensive empirical evaluation, comparing several supervised and self-supervised approaches to our proposed method. To the best of our knowledge, these are the first empirical results on using machine learning, either supervised or self-supervised, to solve the CMPE task in PGMs. On a number of benchmark models, our experiments show that neural networks trained using our proposed loss function are more efficient and accurate compared to models trained to minimize competing supervised and self-supervised loss functions from the literature.

## 2 Notation and Background

We denote random variables by upper-case letters (e.g.,  $X$ ,  $Y$ ,  $Z$ , etc.), their corresponding assignments by lower-case letters (e.g.,  $x$ ,  $y$ ,  $z$ , etc.), sets of random variables by bold upper-case letters (e.g.,  $\mathbf{X}$ ,  $\mathbf{Y}$ ,  $\mathbf{Z}$ , etc.) and assignments to them by bold lower-case letters (e.g.,  $\mathbf{x}$ ,  $\mathbf{y}$ ,  $\mathbf{z}$ , etc.).  $\mathbf{z}_{\mathbf{X}}$  denotes the projection of the complete assignment  $\mathbf{z}$  on to the subset  $\mathbf{X}$  of  $\mathbf{Z}$ . For simplicity of exposition, we assume that discrete and continuous random variables take values from the set  $\{0,1\}$  and  $[0,1]$  respectively.

We use the multilinear polynomial representation (Sherali and Adams, 2009; Sherali and Tuncbilek, 1992; Horst and Tuy, 1996) to concisely describe our proposed method as well as for specifying discrete, continuous, and mixed constrained optimization problems. Let  $\mathbf{Z} = \{Z_1, \dots, Z_n\}$  be a set of random variables. Let  $[n] = \{1, \dots, n\}$  and  $i \in [n]$  be an index over the variables of  $\mathbf{Z}$ . Let  $2^{[n]}$  denote the set of subsets of indices of  $[n]$ ; thus each element of  $2^{[n]}$  denotes a (unique) subset of  $\mathbf{Z}$ . Let  $\mathcal{I} \subseteq 2^{[n]}$  and let  $w_I \in \mathbb{R}$  where  $I \in \mathcal{I}$  be a real number (weight) associated with each element  $I$  of  $\mathcal{I}$ . Then, a multilinear polynomial is given by

$$f(\mathbf{z}) = f(z_1, \dots, z_n) = \sum_{I \in \mathcal{I}} w_I \prod_{i \in I} z_i \quad (1)$$

where  $\mathbf{z} = (z_1, \dots, z_n)$  is an assignment to all variables in  $\mathbf{Z}$ . We will call  $f(\mathbf{z})$  the weight of  $\mathbf{z}$ .

It is known that weighting functions, namely the sum of log of conditional probability tables and log-potentials associated with Bayesian and Markov networks respectively can be expressed as multilinear polynomials (see for example (Koller and Friedman, 2009)).

**Example 1.** *Figure 1 shows a multilinear representation for a Markov network. The weight of the assignment ( $X_1 = 0, X_2 = 1, Y_1 = 0, Y_2 = 1$ ) is 14 and 18 w.r.t.  $\mathcal{M}_1$  and  $\mathcal{M}_2$  respectively.*

### 2.1 Constrained Most Probable Explanation

We are interested in solving the following constrained most probable explanation (CMPE) task. Let  $\mathbf{X}$  and  $\mathbf{Y}$  be two subsets of  $\mathbf{Z}$  such that  $\mathbf{Z} = \mathbf{X} \cup \mathbf{Y}$  and  $\mathbf{X} \cap \mathbf{Y} = \emptyset$ . We




Figure 1: Two Markov networks  $\mathcal{M}_1$  and  $\mathcal{M}_2$  having the same chain-like structure and defined over the same set  $\{X_1, X_2, Y_1, Y_2\}$  of variables.  $\mathcal{M}_1$  is defined by the set of log-potentials  $\{h_1, h_2, h_3\}$  and  $\mathcal{M}_2$  is defined by the set of log-potentials  $\{t_1, t_2, t_3\}$ . Each log-potential can be expressed as a local multilinear polynomial function. The global multilinear function representing  $\mathcal{M}_1$  and  $\mathcal{M}_2$  are  $h(x_1, x_2, y_1, y_2) = 18 - 3x_1 + x_2 - 7y_1 - y_2 + 5x_1y_1 - 4x_2y_2 - y_1y_2$  and  $t(x_1, x_2, y_1, y_2) = 28 + 2x_1 - 7x_2 - 4y_1 - 2y_2 + 2x_1y_1 - x_2y_2$  respectively which are obtained by adding the local functions associated with the respective models and then simplifying, i.e.,  $h(x_1, x_2, y_1, y_2) = h_1(x_1, y_1) + h_2(y_1, y_2) + h_3(x_2, y_2)$ .  $t(x_1, x_2, y_1, y_2)$  is obtained similarly.

will refer to  $\mathbf{Y}$  as *decision variables* and  $\mathbf{X}$  as *evidence variables*. Given assignments  $\mathbf{x}$  and  $\mathbf{y}$ , let  $(\mathbf{x}, \mathbf{y})$  denote their composition. Let  $h$  and  $t$  denote two multilinear polynomials over  $\mathbf{Z}$  obtained from two Markov networks  $\mathcal{M}_1$  and  $\mathcal{M}_2$  respectively that represent two (possibly different) joint probability distributions over  $\mathbf{Z}$ . Then given a real number  $q$  and an assignment  $\mathbf{x}$  to all variables in  $\mathbf{X}$ , the CMPE task is to find an assignment  $\mathbf{y}^*$  to all the variables in  $\mathbf{Y}$  such that  $h(\mathbf{x}, \mathbf{y}^*)$  is maximized (namely the probability of the assignment w.r.t.  $\mathcal{M}_1$  is maximized) and  $t(\mathbf{x}, \mathbf{y}^*) \leq q$  (namely the probability of the assignment w.r.t.  $\mathcal{M}_2$  is bounded by a constant). Formally,

$$\underset{\mathbf{y}}{\text{maximize}} \quad h(\mathbf{x}, \mathbf{y}) \quad \text{s.t.} \quad t(\mathbf{x}, \mathbf{y}) \leq q \quad (2)$$

For brevity, we will abuse notation and use  $h_{\mathbf{x}}(\mathbf{y})$  and  $t_{\mathbf{x}}(\mathbf{y})$  to denote  $h(\mathbf{x}, \mathbf{y})$  and  $t(\mathbf{x}, \mathbf{y})$  respectively. The most probable explanation (MPE) task in probabilistic graphical models (Koller and Friedman, 2009) is a special case of CMPE; MPE is just CMPE without the constraint  $t_{\mathbf{x}}(\mathbf{y}) \leq q$ . The goal in MPE is to find an assignment  $\mathbf{y}^*$  to  $\mathbf{Y}$  such that the weight  $h_{\mathbf{x}}(\mathbf{y}^*)$  of the assignment is maximized given evidence  $\mathbf{x}$ . Similar to MPE, CMPE is NP-hard in general, with the caveat that CMPE is much harder than MPE. Specifically, CMPE is NP-hard even on independent graphical models (having zero treewidth), where MPE can be solved

in linear time by independently maximizing each univariate function (Rouhani et al., 2020).

**Example 2.** Given  $X_1 = 1$ ,  $X_2 = 1$  and  $q = 20$ , the CMPE solution of the example problem in figure 1 is  $(y_1^*, y_2^*) = (0, 1)$  with a value  $h(1, 1, 0, 1) = 11$ , whereas the MPE solution is  $(y_1^*, y_2^*) = (0, 0)$  with value  $h(1, 1, 0, 0) = 16$ .

Since we are interested in machine learning approaches to solve the CMPE task and such approaches employ loss functions, it is convenient to express CMPE as a minimization task with a " $\leq 0$ " constraint. This can be accomplished by negating  $h$  and subtracting  $q$  from  $t$ . Formally, let  $f_{\mathbf{x}}(\mathbf{y}) = -h_{\mathbf{x}}(\mathbf{y})$  and  $g_{\mathbf{x}}(\mathbf{y}) = t_{\mathbf{x}}(\mathbf{y}) - q$ . Then Eq. (2) is equivalent to the following minimization problem:

$$\underset{\mathbf{y}}{\text{minimize}} \quad f_{\mathbf{x}}(\mathbf{y}) \quad \text{s.t.} \quad g_{\mathbf{x}}(\mathbf{y}) \leq 0 \quad (3)$$

Let  $\mathbf{y}^*$  be the optimal solution to the problem given in Eq. (3) and let  $p_{\mathbf{x}}^* = f_{\mathbf{x}}(\mathbf{y}^*)$ . Also, without loss of generality, we assume that  $f_{\mathbf{x}}$  is strictly positive, i.e.,  $\forall \mathbf{y}, f_{\mathbf{x}}(\mathbf{y}) > 0$ .

If all variables in  $\mathbf{Y}$  are binary (or discrete in general), Eq. (3) can be formulated as an (equivalent) integer linear programming (ILP) problem by introducing auxiliary integer variables for each multilinear term (e.g.,  $y_{1,2} = y_1y_2$ ,  $y_{2,3} = y_2y_3$ , etc.) and adding appropriate constraints to model the equivalence between the auxiliary variables and multilinear terms (see for example (Koller and Friedman, 2009), Chapter 13). Therefore, in practice, (3) can be solved optimally using mixed integer linear programming (MILP) solvers such as Gurobi (Gurobi Optimization, 2021) and SCIP (Achterberg et al., 2008; Achterberg, 2009).

Unfortunately, due to a presence of a dense global constraint, namely  $g_{\mathbf{x}}(\mathbf{y}) \leq 0$  in Eq. (3), the MILP solvers often perform poorly. Instead, in practice, application designers often use efficient, specialized algorithms that exploit problem structure for lower bounding  $p_{\mathbf{x}}^*$ , and then using these lower bounds in an *anytime* branch-and-bound algorithm to obtain an upper bound on  $p_{\mathbf{x}}^*$ .

## 2.2 Specialized Lower Bounding Algorithms

Recently, Rahman et al. (2021) proposed two new approaches for computing upper bounds on the optimal value of the maximization problem given in Eq. (2). These methods can be easily adapted to obtain a lower bound on  $p_{\mathbf{x}}^*$ ; because an upper bound on the maximization problem is a lower bound on the corresponding minimization problem. We present the adaptations of Rahman et al.'s approach next.

The first approach is based on the Lagrangian relaxation method that introduces a *Lagrange multiplier*  $\mu \geq 0$  to transform the constrained minimization problem to the following unconstrained problem: minimize  $f_{\mathbf{x}}(\mathbf{y}) + \mu g_{\mathbf{x}}(\mathbf{y})$ . Let  $d_{\mu}^*$  denote the optimal value of the unconstrained problem. Then, it is easy to show that  $d_{\mu}^* \leq p_{\mathbf{x}}^*$ . The largest upper

bound is obtained by finding a value of  $\mu$  that maximizes  $d_\mu^*$ . More formally,

$$\max_{\mu \geq 0} d_\mu^* = \max_{\mu \geq 0} \min_{\mathbf{y}} f_{\mathbf{x}}(\mathbf{y}) + \mu g_{\mathbf{x}}(\mathbf{y}) \leq p_{\mathbf{x}}^* \quad (4)$$

Rahman et al. (2021) proposed to solve the inner minimization problem using exact techniques from the graphical models literature such as variable/bucket elimination (Dechter, 1999), branch and bound search and best-first search (Marinescu and Dechter, 2012, 2009; Wu et al., 2020). When exact inference is not feasible, Rahman et al. (2021) proposed to solve the inner problem using approximate inference techniques such as mini-bucket elimination, dual-decomposition and join-graph based bounding algorithms (Choi and Darwiche, 2011; Dechter and Rish, 2003; Wainwright et al., 2005; Globerson and Jaakkola, 2007; Komodakis et al., 2007; Ihler et al., 2012). The outer maximization problem is solved using sub-gradient ascent.

The second approach by (Rahman et al., 2021) uses the Lagrangian decomposition method to transform the problem into a multi-choice knapsack problem (MCKP) and then utilizes off-the-shelf MCKP solvers. In our experiments, we use the Lagrange relaxation approach given in Eq. (4).

If the set  $\mathbf{Y}$  contains continuous variables, then it is not possible to reduce it to an equivalent MILP/LP (Horst and Tuy, 1996; Sherali and Tuncbilek, 1992). However, by leveraging linearization methods (Sherali and Tuncbilek, 1992; Sherali and Adams, 2009) and solving the resulting problem using linear programming (LP) solvers, we can still obtain good lower bounds on  $p_{\mathbf{x}}^*$ .

### 3 Solving CMPE using Methods from the Learning to Optimize Literature

In this section, we show how techniques developed in the learning to optimize literature (Donti et al., 2021; Fioretto et al., 2020; Park and Van Hentenryck, 2022; Zamzam and Baker, 2019) which seeks to develop machine learning approaches for solving constrained optimization problems can be leveraged to solve the CMPE task. The main idea is to train a deep neural network  $\mathcal{F}_\Theta : \mathbf{X} \rightarrow \mathbf{Y}$  parameterized by the set  $\Theta \in \mathbb{R}^M$  such that at test time given evidence  $\mathbf{x}$ , the network is able to predict an (near) optimal solution  $\hat{\mathbf{y}}$  to the CMPE problem. Note that as far as we are aware, no prior work exists on solving CMPE using deep neural networks.

#### 3.1 Supervised Methods

In order to train the parameters of  $\mathcal{F}_\Theta$  in a supervised manner, we need to acquire labeled data in the form  $\mathcal{D} = \{\langle \mathbf{x}_i, \mathbf{y}_i \rangle\}_{i=1}^N$  where each label  $\mathbf{y}_i$  is an optimal solution to the problem given in Eq. (3) given  $\mathbf{x}_i$ . In practice, we can generate the assignments  $\{\mathbf{x}_i\}_{i=1}^N$  by sampling them from the graphical model corresponding to  $f$  and the labels

$\{\mathbf{y}_i\}_{i=1}^N$  by solving the minimization problem given in Eq. (3) using off-the-shelf solvers such as Gurobi and SCIP.

Let  $\hat{\mathbf{y}}_i = \mathcal{F}_\Theta(\mathbf{x}_i)$  denote the labels predicted by the neural network for  $\mathbf{x}_i$ . Following Zamzam and Baker (2019), we propose to train  $\mathcal{F}_\Theta$  using the following two loss functions

$$\text{Mean-Squared Error (MSE)} : \frac{1}{N} \sum_i (\mathbf{y}_i - \hat{\mathbf{y}}_i)^2 \quad (5)$$

$$\text{Mean-Absolute-Error (MAE)} : \frac{1}{N} \sum_i |\mathbf{y}_i - \hat{\mathbf{y}}_i| \quad (6)$$

Experimentally (see the supplementary material), we found that neural networks trained using the MAE and MSE loss functions often output infeasible assignments. To address this issue, following prior work (Nellikath and Chatzivasileiadis, 2021), we propose to add  $\lambda_{\mathbf{x}} \max\{0, g_{\mathbf{x}}(\hat{\mathbf{y}})\}$  to the loss function where  $\lambda_{\mathbf{x}}$  is a penalty coefficient.

In prior work (Fioretto et al., 2020), it was observed that the quality of the solutions greatly depends on the value chosen for  $\lambda_{\mathbf{x}}$ . Moreover, it is not straightforward to choose it optimally because it varies for each  $\mathbf{x}$ . To circumvent this issue, we propose to update  $\lambda_{\mathbf{x}}$  via a *Lagrangian dual* method (Nocedal and Wright, 2006). More specifically, we propose to use the following subgradient method to optimize the value of  $\lambda_{\mathbf{x}}$ . While training a neural network, let  $\lambda_{\mathbf{x}_i}^k$  and  $\hat{\mathbf{y}}_i^k$  denote the values of the penalty co-efficient and the predicted assignment respectively at the  $k$ -th epoch and for the  $i$ -th example in  $\mathcal{D}$  (if the  $i$ -th example is part of the current mini-batch), then, we update  $\lambda_{\mathbf{x}_i}^{k+1}$  using

$$\lambda_{\mathbf{x}_i}^{k+1} = \lambda_{\mathbf{x}_i}^k + \rho \max\{0, g_{\mathbf{x}_i}(\hat{\mathbf{y}}_i^k)\} \quad (7)$$

where  $\rho$  is the Lagrangian step size. In our experiments, we evaluated both the naive and the penalty based supervised loss approaches (for CMPE) and found that the penalty method with MSE loss yields the best results. Therefore, in our experiments, we use it as a strong supervised baseline.

#### 3.2 Self-Supervised Methods

Supervised methods require pre-computed solutions for numerous NP-hard/multilinear problem instances, which are computationally expensive to derive. Therefore, we propose to train the neural network in a self-supervised manner that does not depend on the pre-computed results. Utilizing findings from (Kotary et al., 2021) and (Park and Van Hentenryck, 2022), we introduce two self-supervised approaches: one is grounded in the *penalty method*, and the other builds upon the *augmented Lagrangian method*.

**Penalty Method** (Donti et al., 2021; Kotary et al., 2021; Fioretto et al., 2020). In the penalty method, we solve the constrained minimization problem by iteratively transforming it into a sequence of unconstrained problems. Each unconstrained problem at iteration  $k$  is constructed by adding

a term, which consists of a penalty parameter  $\lambda_{\mathbf{x}}^k$  multiplied by a function  $\max\{0, g_{\mathbf{x}}(\mathbf{y})\}^2$  that quantifies the constraint violations, to the objective function. Formally, the optimization problem at the  $k$ -th step is given by:

$$\min_{\mathbf{y}} f_{\mathbf{x}}(\mathbf{y}) + \frac{\lambda_{\mathbf{x}}^k}{2} \max\{0, g_{\mathbf{x}}(\mathbf{y})\}^2 \quad (8)$$

Here,  $\lambda_{\mathbf{x}}^k$  is progressively increased either until the constraint is satisfied or a predefined maximum  $\lambda_{\max}$  is reached.  $\lambda_{\mathbf{x}}^k$  can be updated after a few epochs using simple strategies such as multiplication by a fixed factor (e.g., 2, 10, etc.).

The penalty method can be adapted to learn a neural network in a self-supervised manner as follows. At each epoch  $k$ , we sample an assignment  $\mathbf{x}$  (or multiple samples for a mini-batch) from the graphical model corresponding to  $f$ , predict  $\hat{\mathbf{y}}$  using the neural network and then use the following loss function to update its parameters:

$$\mathcal{L}_{\mathbf{x}}^{\text{pen}}(\hat{\mathbf{y}}) = f_{\mathbf{x}}(\hat{\mathbf{y}}) + \frac{\lambda_{\mathbf{x}}^k}{2} \max\{0, g_{\mathbf{x}}(\hat{\mathbf{y}})\}^2 \quad (9)$$

Determining the optimal  $\lambda_{\mathbf{x}}^k$  is crucial. In prior work, Kotary et al. (2021) and Fioretto et al. (2020) proposed to update it via a subgradient method, similar to the update rule given by Eq. (7). More formally, we can update  $\lambda_{\mathbf{x}}^k$  using:

$$\lambda_{\mathbf{x}}^{k+1} = \lambda_{\mathbf{x}}^k + \rho \max\{0, g_{\mathbf{x}}(\hat{\mathbf{y}}^k)\} \quad (10)$$

where  $\rho$  is the Lagrangian step size.

**Augmented Lagrangian Method (ALM).** In this method, we augment the objective used in the penalty method with a Lagrangian term. More formally, the optimization problem at the  $k$ -th step is given by (compare with Eq. (8)):

$$\min_{\mathbf{y}} f_{\mathbf{x}}(\mathbf{y}) + \frac{\lambda_{\mathbf{x}}^k}{2} \max\{0, g_{\mathbf{x}}(\mathbf{y})\}^2 + \mu_{\mathbf{x}}^k g_{\mathbf{x}}(\mathbf{y}) \quad (11)$$

Here,  $\lambda_{\mathbf{x}}^k$  may be progressively increased similar to the penalty method while  $\mu_{\mathbf{x}}^k$  is updated using

$$\mu_{\mathbf{x}}^{k+1} = \max\{0, \mu_{\mathbf{x}}^k + \lambda_{\mathbf{x}}^k g_{\mathbf{x}}(\mathbf{y}^k)\} \quad (12)$$

Recently, Park and Van Hentenryck (2022) proposed a self-supervised primal-dual learning method that leverages two distinct networks to emulate the functionality of ALM: the first (primal) network takes as input  $\mathbf{x}$  and outputs  $\mathbf{y}$  while the second network focuses on learning the dual aspects; specifically it takes  $\mathbf{x}$  as input and outputs  $\mu_{\mathbf{x}}^k$ . The training process uses a sequential approach, where one network is trained while the other remains frozen to furnish the requisite values for the loss computation.

The primal network uses the following loss function:

$$\mathcal{L}_{\mathbf{x}}^{A,p}(\hat{\mathbf{y}}|\mu, \lambda) = f_{\mathbf{x}}(\hat{\mathbf{y}}) + \frac{\lambda}{2} \max\{0, g_{\mathbf{x}}(\hat{\mathbf{y}})\}^2 + \mu g_{\mathbf{x}}(\hat{\mathbf{y}})$$

While the dual network uses the following loss function

$$\mathcal{L}_{\mathbf{x}}^{A,d}(\hat{\mu}|\mathbf{y}, \lambda, \mu^k) = \|\hat{\mu} - \max\{0, \mu^k + \lambda g_{\mathbf{x}}(\mathbf{y})\}\|$$

where  $\hat{\mu}$  is the predicted value of the Lagrangian multiplier.

### 3.2.1 Drawbacks of the Penalty and ALM Methods

A limitation of the penalty-based self-supervised method is that it does not guarantee a global minimum unless specific conditions are met. In particular, the optimal solution w.r.t. the loss function (see Eq. (9)) may be far away from the optimal solution  $\mathbf{y}^*$  of the problem given in Eq. (3), unless the penalty co-efficient  $\lambda_{\mathbf{x}}^k \rightarrow \infty$ . Moreover, when  $\lambda_{\mathbf{x}}^k$  is large for all  $\mathbf{x}$ , the gradients will be uninformative. In the case of ALM method (cf. (Nocedal and Wright, 2006)), for global minimization, we require that either  $\lambda_{\mathbf{x}}^k \rightarrow \infty$  or  $\forall \mathbf{x}$  with  $g_{\mathbf{x}}(\mathbf{y}) > 0$ ,  $\mu_{\mathbf{x}}^k$  should be such that  $\min_{\mathbf{y}} f_{\mathbf{x}}(\mathbf{y}) + \mu_{\mathbf{x}}^k g_{\mathbf{x}}(\mathbf{y}) > p_{\mathbf{x}}^*$ . Additionally, ALM introduces a dual network, increasing the computational complexity and potentially leading to negative information transfer when the dual network’s outputs are inaccurate. These outputs are subsequently utilized in the loss to train the primal network for the following iteration, thereby exerting a negative effect. To address these limitations, next, we introduce a self-supervised method that achieves global minimization without the need for a dual network or infinite penalty coefficients.

## 4 A NOVEL SELF-SUPERVISED CMPE SOLVER

An appropriately designed loss function should have the following characteristics. For feasible solutions, namely when  $g_{\mathbf{x}}(\mathbf{y}) \leq 0$ , the loss function should be proportional to  $f_{\mathbf{x}}(\mathbf{y})$ . While for infeasible assignments, it should equal infinity. This loss function will ensure that once a feasible solution is found, the neural network will only explore the space of feasible solutions. Unfortunately, infinity does not provide any gradient information, and the neural network will get stuck in the infeasible region if the neural network generates an infeasible assignment during training.

An alternative approach is to use  $g$  as a loss function when the constraint is not satisfied (i.e.,  $g_{\mathbf{x}}(\mathbf{y}) > 0$ ) in order to push the infeasible solutions towards feasible ones (Liu and Cherian, 2023). Unfortunately, this approach will often yield feasible solutions that lie at the boundary  $g_{\mathbf{x}}(\mathbf{y}) = 0$ . For instance, for a boundary assignment  $\mathbf{y}_b$  where  $g_{\mathbf{x}}(\mathbf{y}_b) = 0$  but  $f_{\mathbf{x}}(\mathbf{y}_b) > 0$  (or decreasing), the sub-gradient will be zero, and the neural network will treat the boundary assignment as an optimal one.

To circumvent this issue, we propose a loss function which has the following two properties: (1) It is proportional to  $g$  in the infeasible region with  $f$  acting as a control in the boundary region (when  $g$  is zero); and (2) It is proportional to  $f$  in the feasible region. Formally,

$$\mathcal{L}_{\mathbf{x}}(\hat{\mathbf{y}}) = \begin{cases} f_{\mathbf{x}}(\hat{\mathbf{y}}) & \text{if } g_{\mathbf{x}}(\hat{\mathbf{y}}) \leq 0 \\ \alpha_{\mathbf{x}}(f_{\mathbf{x}}(\hat{\mathbf{y}}) + g_{\mathbf{x}}(\hat{\mathbf{y}})) & \text{if } g_{\mathbf{x}}(\hat{\mathbf{y}}) > 0 \end{cases} \quad (13)$$

where  $\alpha_{\mathbf{x}}$  is a function of the evidence  $\mathbf{x}$ . Our goal is to find

a bound for  $\alpha_{\mathbf{x}}$  such that the following desirable property is satisfied and the bound can be computed in polynomial time for each  $\mathbf{x}$  by leveraging bounding methods for CMPE.

**Property (Consistent Loss):** The loss for all infeasible assignments is higher than the optimal value  $p_{\mathbf{x}}^*$ . To satisfy this property, we have to ensure that:

$$\forall \hat{\mathbf{y}} \text{ s.t. } g_{\mathbf{x}}(\hat{\mathbf{y}}) > 0, \quad \alpha_{\mathbf{x}} (f_{\mathbf{x}}(\hat{\mathbf{y}}) + g_{\mathbf{x}}(\hat{\mathbf{y}})) > p_{\mathbf{x}}^*$$

which implies that the following condition holds.

$$\alpha_{\mathbf{x}} \left( \min_{\hat{\mathbf{y}}} f_{\mathbf{x}}(\hat{\mathbf{y}}) + g_{\mathbf{x}}(\hat{\mathbf{y}}) \text{ s.t. } g_{\mathbf{x}}(\hat{\mathbf{y}}) > 0 \right) > p_{\mathbf{x}}^*$$

Let  $q_{\mathbf{x}}^*$  denote the optimal value of  $\min_{\hat{\mathbf{y}}} f_{\mathbf{x}}(\hat{\mathbf{y}}) + g_{\mathbf{x}}(\hat{\mathbf{y}}) \text{ s.t. } g_{\mathbf{x}}(\hat{\mathbf{y}}) > 0$ . Then,  $\alpha_{\mathbf{x}} > \frac{p_{\mathbf{x}}^*}{q_{\mathbf{x}}^*}$ .

**Proposition 4.1.** *If  $\mathcal{L}_{\mathbf{x}}(\hat{\mathbf{y}})$  is consistent, i.e.,  $\alpha_{\mathbf{x}} > \frac{p_{\mathbf{x}}^*}{q_{\mathbf{x}}^*}$  then  $\min_{\hat{\mathbf{y}}} \mathcal{L}_{\mathbf{x}}(\hat{\mathbf{y}}) = p_{\mathbf{x}}^*$ , namely  $\mathcal{L}_{\mathbf{x}}(\hat{\mathbf{y}})$  is an optimal loss function.*

*Proof.* From equation (13), we have

$$\begin{aligned} \min_{\hat{\mathbf{y}}} \mathcal{L}_{\mathbf{x}}(\hat{\mathbf{y}}) &= \min \left\{ \min_{\hat{\mathbf{y}}} f_{\mathbf{x}}(\hat{\mathbf{y}}) \text{ s.t. } g_{\mathbf{x}}(\hat{\mathbf{y}}) \leq 0, \right. \\ &\quad \left. \alpha_{\mathbf{x}} \left( \min_{\hat{\mathbf{y}}} f_{\mathbf{x}}(\hat{\mathbf{y}}) + g_{\mathbf{x}}(\hat{\mathbf{y}}) \text{ s.t. } g_{\mathbf{x}}(\hat{\mathbf{y}}) > 0 \right) \right\} \\ &= \min \{ p_{\mathbf{x}}^*, \alpha_{\mathbf{x}} q_{\mathbf{x}}^* \} \end{aligned} \quad (14)$$

Because  $\mathcal{L}_{\mathbf{x}}(\hat{\mathbf{y}})$  is consistent, namely,  $\alpha_{\mathbf{x}} > \frac{p_{\mathbf{x}}^*}{q_{\mathbf{x}}^*}$ , we have

$$\min \{ p_{\mathbf{x}}^*, \alpha_{\mathbf{x}} q_{\mathbf{x}}^* \} = p_{\mathbf{x}}^* \quad (15)$$

From equations (14) and (15), the proof follows.  $\square$

We assume that  $f_{\mathbf{x}}(\mathbf{y})$  and  $g_{\mathbf{x}}(\mathbf{y})$  are bounded functions, namely for any assignment  $(\mathbf{x}, \mathbf{y})$ ,  $l_f \leq f_{\mathbf{x}}(\mathbf{y}) \leq u_f$  and  $l_g \leq g_{\mathbf{x}}(\mathbf{y}) \leq u_g$  where  $-\infty < s < \infty$  and  $s \in \{l_f, u_f, l_g, u_g\}$ . Also, for simplicity, we assume that  $f_{\mathbf{x}}(\mathbf{y})$  is a strictly positive function, namely  $l_f > 0$ .

Thus, based on the assumptions given above, we have

$$\frac{p_{\mathbf{x}}^*}{q_{\mathbf{x}}^*} \leq \frac{u_f}{l_f} \text{ and } 0 < \alpha_{\mathbf{x}} \leq \frac{u_f}{l_f}$$

The above assumptions will ensure that the gradients are bounded, because  $\alpha_{\mathbf{x}}$ ,  $f$  and  $g$  are bounded, and both  $p_{\mathbf{x}}^*$  and  $q_{\mathbf{x}}^*$  are greater than zero.

Next, we show how to compute an upper bound on  $\alpha_{\mathbf{x}}$  using  $\alpha_{\mathbf{x}} > \frac{p_{\mathbf{x}}^*}{q_{\mathbf{x}}^*}$ , thus ensuring that we have an optimal loss function. The terms in the numerator ( $p_{\mathbf{x}}^*$ ) and denominator ( $q_{\mathbf{x}}^*$ ) require solving two instances of the CMPE task. Since solving CMPE exactly is impractical and moreover, since we are interested in self-supervised methods where we do

not assume access to such a solver, we propose to lower bound  $q_{\mathbf{x}}^*$  and upper bound  $p_{\mathbf{x}}^*$ .

For a given instance  $\mathbf{x}$ , a lower bound on  $q_{\mathbf{x}}^*$  can be obtained using the Lagrangian relaxation method described in section 2.2 (see Eq. (4)) for discrete variables and the Reformulation-Linearization method described in Sherali and Tuncbilek (1992) for continuous variables. On the other hand, any feasible solution can serve as an upper bound for  $p_{\mathbf{x}}^*$ . A simple yet efficient approach is to begin with a loose upper bound by upper bounding the MPE task:  $\max_{\mathbf{y}} f_{\mathbf{x}}(\mathbf{y})$ , using fast algorithms such as mini-bucket elimination (Dechter and Rish, 2003) or fast linear programming based approximations (Ihler et al., 2012; Globerson and Jaakkola, 2007) and then keep track of feasible solutions during batch-style gradient descent.

In summary, we proposed a new loss function which uses the quantity  $\alpha_{\mathbf{x}}$ . When the neural network predicts a feasible  $\hat{\mathbf{y}}$ , the loss equals  $f$ , whereas when it predicts an infeasible  $\hat{\mathbf{y}}$ , the loss is such that the infeasible solution can quickly be pushed towards a feasible solution (because it uses gradients from  $g$ ). A key advantage of our proposed loss function is that  $\alpha_{\mathbf{x}}$  is not treated as an optimization variable, and a bound on it can be pre-computed for each example  $\mathbf{x}$ .

#### 4.1 Making The Loss Function Smooth and Continuous

The loss function defined in Eq. (13) is continuous and differential everywhere except at  $g_{\mathbf{x}}(\hat{\mathbf{y}}) = 0$ . There is a *jump discontinuity* at  $g_{\mathbf{x}}(\hat{\mathbf{y}}) = 0$  since  $\lim_{g_{\mathbf{x}}(\hat{\mathbf{y}}) \rightarrow 0^-} f_{\mathbf{x}}(\hat{\mathbf{y}}) \neq \lim_{g_{\mathbf{x}}(\hat{\mathbf{y}}) \rightarrow 0^+} \alpha_{\mathbf{x}} (f_{\mathbf{x}}(\hat{\mathbf{y}}) + g_{\mathbf{x}}(\hat{\mathbf{y}}))$ . To address this issue, we propose the following continuous approximation

$$\begin{aligned} \tilde{\mathcal{L}}_{\mathbf{x}}(\hat{\mathbf{y}}) &= \left( (1 - \sigma(\beta g_{\mathbf{x}}(\hat{\mathbf{y}}))) \cdot [f_{\mathbf{x}}(\hat{\mathbf{y}})] \right) + \\ &\quad \left( \sigma(\beta g_{\mathbf{x}}(\hat{\mathbf{y}})) \cdot [\alpha_{\mathbf{x}} (f_{\mathbf{x}}(\hat{\mathbf{y}}) + \max\{0, g_{\mathbf{x}}(\hat{\mathbf{y}})\})] \right) \end{aligned} \quad (16)$$

where  $\sigma(\cdot)$  is the sigmoid function and  $\beta \geq 0$  is a hyper-parameter that controls the steepness of the sigmoid. At a high level, the above continuous approximation uses a sigmoid function to approximate a Heaviside step function.

## 5 EXPERIMENTAL EVALUATION

In this section, we thoroughly evaluate the effectiveness of our proposed neural networks based solvers for CMPE. We evaluate the competing methods on several test problems using three criteria: optimality gap (relative difference between the optimal solution and the one found by the method), constraint violations (percentage of time the method outputs an infeasible solution), and training and inference times.




Figure 2: Optimality Gap (avg %) and Average Violations for Self-Supervised methods

Table 1: Average gap and constraint violations over test samples for models from the UAI competition.  $\pm$  denotes standard deviation. **Bold** values indicate the methods with the highest performance. Underlined values denote significant violations, particularly those exceeding a threshold of 0.15. For these methods, the gap values are not considered in our analysis.

Methods		Segment12	Segment14	Segment15	Grids17	Grids18
ILP Obj.		463.454	471.205	514.287	2879.469	4160.196
$SL_{pen}$	Gap	<b>0.053 <math>\pm</math> 0.043</b>	0.053 $\pm$ 0.043	0.053 $\pm$ 0.041	0.092 $\pm$ 0.070	0.082 $\pm$ 0.065
	Violations	<u>0.238 <math>\pm</math> 0.426</u>	<u>0.248 <math>\pm</math> 0.432</u>	<u>0.153 <math>\pm</math> 0.361</u>	0.054 $\pm$ 0.226	0.053 $\pm$ 0.224
$SSL_{pen}$	Gap	<b>0.051 <math>\pm</math> 0.042</b>	0.065 $\pm$ 0.048	<b>0.056 <math>\pm</math> 0.044</b>	0.089 $\pm$ 0.055	0.104 $\pm$ 0.062
	Violations	0.149 $\pm$ 0.357	0.127 $\pm$ 0.332	0.086 $\pm$ 0.281	0.004 $\pm$ 0.059	0.005 $\pm$ 0.071
PDL	Gap	0.063 $\pm$ 0.050	0.055 $\pm$ 0.042	0.063 $\pm$ 0.049	0.102 $\pm$ 0.059	0.092 $\pm$ 0.061
	Violations	<b>0.073 <math>\pm</math> 0.261</b>	0.120 $\pm$ 0.326	0.016 $\pm$ 0.126	<b>0.000 <math>\pm</math> 0.000</b>	<b>0.001 <math>\pm</math> 0.022</b>
SS-CMPE	Gap	0.055 $\pm$ 0.045	<b>0.051 <math>\pm</math> 0.040</b>	0.068 $\pm$ 0.051	<b>0.067 <math>\pm</math> 0.049</b>	<b>0.069 <math>\pm</math> 0.051</b>
	Violations	0.093 $\pm$ 0.291	<b>0.107 <math>\pm</math> 0.415</b>	<b>0.002 <math>\pm</math> 0.039</b>	0.001 $\pm$ 0.032	<b>0.001 <math>\pm</math> 0.022</b>

### 5.1 The Loss Functions: Competing Methods

We trained several neural networks to minimize both supervised and self-supervised loss functions. We evaluated both MSE and MAE supervised losses with and without penalty coefficients (see section 3). In the main paper, we show results on the best performing supervised loss, which is MSE with penalty, denoted by  $SL_{pen}$  (results for other supervised loss functions are provided in the supplement).

For self-supervised loss, we experimented with the following three approaches: (1) penalty-based method, (2) ALM, which uses a primal-dual loss (PDL), and the approach described in section 4. We will refer to these three schemes as  $SSL_{pen}$ , PDL, and SS-CMPE, respectively. We used the experimental setup described by Park and Van Hentenryck (2022) for tuning the hyperparameters of PDL and  $SSL_{pen}$ . For the SS-CMPE method, we employed a grid search approach to determine the optimal values for  $\beta$ . The range of values considered for  $\beta$  was  $\{0.1, 1.0, 2.0, 5.0, 10.0, 20.0\}$ .

Note that all methods used the same neural network architecture (described in the supplement) except PDL, which uses two neural networks. We obtained the ground-truth for the supervised training by solving the original ILP problem using SCIP (Achterberg, 2009) and Gurobi (Gurobi Optimization, 2021). We report the objective values of the ILP

solutions in each table (see Tables 1, 2, and 3).

### 5.2 Datasets and Benchmarks

We evaluate the competing algorithms ( $SL_{pen}$ ,  $SSL_{pen}$ , PDL, and SS-CMPE) on a number of log-linear Markov networks ranging from simple models (low treewidth) to high treewidth models. The simple models comprise of learned tractable probabilistic circuits (Choi et al., 2020) without latent variables, specifically, cutset networks (Rahman et al., 2014) from benchmark datasets used in the literature. The complex, high treewidth models are sourced from past UAI inference competitions (Elidan and Globerson, 2010). Finally, we evaluated all methods on the task of generating adversarial examples for neural network classifiers.

### 5.3 High Tree-Width Markov Networks and Tractable Probabilistic Circuits

Our initial series of experiments focus on high treewidth Grids and Image Segmentation Markov networks from the UAI inference competitions (Gogate, 2014, 2016). In this investigation, we generated CMPE problems by utilizing the model employed in the UAI competitions, denoted as  $M_1$ . Subsequently,  $M_2$  was created by adjusting the parameters of  $M_1$  while incorporating a noise parameter  $\epsilon$  drawn from

Table 2: Average gap and constraint violations over test samples from tractable probabilistic models.  $\pm$  denotes standard deviation. **Bold** values indicate the methods with the highest performance. Underlined values denote significant violations, particularly those exceeding a threshold of 0.15. For these methods, the gap values are not considered in our analysis.

Methods		AD	BBC	DNA	20 NewsGroup	WebKB1
ILP Obj.		2519.128	871.567	221.119	921.702	824.493
$SL_{pen}$	Gap	$0.156 \pm 0.057$	$0.036 \pm 0.027$	$0.143 \pm 0.113$	<b><math>0.041 \pm 0.031</math></b>	<b><math>0.044 \pm 0.035</math></b>
	Violations	$0.135 \pm 0.341$	<u><math>0.237 \pm 0.425</math></u>	<u><math>0.151 \pm 0.358</math></u>	$0.084 \pm 0.277$	$0.070 \pm 0.254$
$SSL_{pen}$	Gap	$0.159 \pm 0.055$	$0.045 \pm 0.033$	$0.142 \pm 0.116$	$0.045 \pm 0.036$	$0.058 \pm 0.043$
	Violations	$0.008 \pm 0.089$	$0.056 \pm 0.230$	$0.014 \pm 0.118$	$0.005 \pm 0.071$	$0.025 \pm 0.158$
PDL	Gap	$0.154 \pm 0.055$	$0.051 \pm 0.036$	$0.144 \pm 0.117$	$0.046 \pm 0.035$	$0.059 \pm 0.043$
	Violations	<b><math>0.000 \pm 0.000</math></b>	<b><math>0.025 \pm 0.156</math></b>	<b><math>0.006 \pm 0.077</math></b>	<b><math>0.004 \pm 0.059</math></b>	<b><math>0.012 \pm 0.109</math></b>
SS-CMPE	Gap	<b><math>0.134 \pm 0.054</math></b>	<b><math>0.043 \pm 0.033</math></b>	<b><math>0.138 \pm 0.112</math></b>	$0.045 \pm 0.035$	$0.057 \pm 0.043$
	Violations	$0.006 \pm 0.077$	$0.056 \pm 0.230$	$0.007 \pm 0.083$	$0.005 \pm 0.071$	$0.016 \pm 0.126$

a normal distribution with mean 0 and variance  $\sigma^2 = 0.1$ . To select  $q$ , we randomly generated 100 samples, sorted them based on their weight, and then selected the weight of the 10th, 30th, 60th, 80th, and 90th sample as a value for  $q$ . We assess the impact of changing the value of  $q$  in the supplement. Experiments in the main body of the paper use  $q$  equal to the weight of the 80th random sample. For each network, we randomly chose 60% of variables as evidence ( $\mathbf{X}$ ) and the remaining as query variables ( $\mathbf{Y}$ ). For both the UAI models and tractable probabilistic circuits, we generated 10K samples from  $M_1$ , and used 9K for training and 1K for testing. We used 5-fold cross validation for selecting the hyperparameters.

The results for the UAI datasets are shown in Table 1. We see that for the majority of the datasets, our method produces solutions with superior gap values compared to the other methods. Even in situations where our methods do not achieve better gap values, they exhibit fewer violations. This demonstrates the effectiveness and robustness of our methods in generating solutions that strike a balance between optimizing the objective function and maintaining constraint adherence. For certain datasets, our methods exhibit significantly lower constraint violations, even up to 10 times less than supervised methods.

In the next phase of our study, we employed MPE (Most Probable Explanation) tractable models, which were learned on five high-dimensional datasets (see Lowd and Davis (2010) for details in the datasets): DNA, NewsGroup (c20ng), WebKB1 (cwebkb), AD, and BBC. These learned models served as  $M_1$ . We then applied Gaussian noise as described earlier to generate  $M_2$  based on  $M_1$ . A similar trend can be observed for tractable probabilistic models in Table 2, where our method consistently outperforms the other self-supervised methods across all datasets. Not only does our approach exhibit superior performance in terms of gap values, but it also demonstrates comparable constraint violations. When comparing with the supervised method, our proposed algorithm exhibits significantly fewer con-

straint violations while maintaining a better or comparable gap value. This emphasizes the strength of our method in effectively balancing the optimization objectives and constraint adherence, thereby offering improved overall performance compared to both the self-supervised and supervised approaches in the context of tractable probabilistic models. In Figure 2, we present the average optimality gap and average violations for different dataset groups. It is important to note that results closer to the origin indicate better performance.

#### 5.4 Adversarial Modification on the MNIST Dataset

Table 3: Performance comparison of supervised and self-supervised methods. The table presents the average objective value, gap, and constraint violations over the test examples, along with the training and inference time required for each method for adversarial example generation. **Bold** values signify the methods that achieved the best scores.

Methods	ILP Obj.	Gap	Violation	Time in seconds	
				Train	Inf.
ILP	30.794	0.000	0.000	NA	5.730
$SL_{pen}$	63.670	1.069	0.071	57534.4	0.003
$SSL_{pen}$	76.316	1.480	0.052	<b>469.540</b>	0.003
PDL	66.400	1.158	0.055	839.025	0.003
SS-CMPE	<b>62.400</b>	<b>1.028</b>	<b>0.021</b>	520.149	0.003

We also evaluated our approach on the task of adversarial example generation for discriminative classifiers, specifically neural networks. Adversarial examples play a crucial role in assessing the robustness of models and facilitating the training of more resilient models. The task of Adversarial Example Generation involves producing new images by making minimal modifications to input images that are mis-classified by the model. Rahman et al. (2021) showed that this problem can be reduced to CMPE. Formally, let  $\mathcal{G}$  be a differentiable, continuous function defined over a set of inputs  $\mathcal{X}$ . Given an assignment  $\mathcal{X} = x$ , we introduce the decision variable  $\mathcal{D}$ , which takes the value  $d$  when  $\mathcal{G} > 0$






Figure 3: Qualitative results on the adversarially generated MNIST digits. Each row represents an original image followed by a corresponding image generated adversarially by 8 different methods: ILP, MSE, SL+Penalty, MAE, MAE+Penalty,  $SSL_{pen}$ , PDL, and SS-CMPE.

and  $\bar{d}$  otherwise. In the context of adversarial attacks, given an image  $x$ , our objective is to generate a new image  $x'$  such that the distance between  $x$  and  $x'$  is minimized and the decision is flipped (namely  $\mathcal{G} < 0$ ). We used a log-linear model  $\mathcal{F}$  to represent the sum absolute distance between the pixels. Then the task of adversarial example generation can be formulated as the following CMPE problem: maximize  $\sum_{f \in \mathcal{F}} f(x'|x)$  s.t.  $\mathcal{G}(x'|x) \leq 0$ .

We evaluated the algorithms using the MNIST handwritten digit dataset (LeCun and Cortes, 2010). We trained a multi-layered neural network having >95% test accuracy and used it as our  $\mathcal{G}$  function. To generate adversarial examples corresponding to a given test example, we trained an autoencoder  $A : X \rightarrow X'$  using four loss functions corresponding to  $SL_{pen}$ ,  $SSL_{pen}$ , PDL and SS-CMPE. We used the default train-test split (10K examples for testing and 60K for training).

Table 3 shows quantitative results comparing our proposed SS-CMPE method with other competing methods. We can clearly see that SS-CMPE is superior to competing self-supervised ( $SSL_{pen}$  and PDL) and supervised methods ( $SL_{pen}$ ) in terms of both constraint violations and optimality gap. The second best method in terms of optimality gap is  $SL_{pen}$ . However, its constraint violations are much higher, and its training time is significantly larger because it needs access to labeled data, which in turn requires using computationally expensive ILP solvers. The training time of SS-CMPE is much smaller than PDL (because the latter uses two networks) and is only slightly larger than  $SSL_{pen}$ .

Figure 3 shows qualitative results on adversarial modification to the MNIST digits  $\{1, 2, 6, 7, 8\}$  by all the eight methods. The CMPE task minimally changes an input image such that the corresponding class is flipped according to a discriminative classifier. MSE and our proposed method SS-CMPE are very competitive and were able to generate visually indistinguishable, high-quality modifications whereas the other methods struggled to do so.

**Summary:** Our experiments show that SS-CMPE consistently outperforms competing self-supervised methods, PDL and  $SSL_{pen}$ , in terms of optimality gap and is comparable to PDL in terms of constraint violations. The training time of SS-CMPE is smaller than PDL (by half as much) and is slightly larger than  $SSL_{pen}$ . However, it is considerably better than  $SSL_{pen}$  in terms of constraint violations. SS-CMPE also employs fewer hyperparameters as compared to PDL.

## 6 CONCLUSION AND FUTURE WORK

In this paper, we proposed a new self-supervised learning algorithm for solving the constrained most probable explanation task which at a high level is the task of optimizing a multilinear polynomial subject to a multilinear constraint. Our main contribution is a new loss function for self-supervised learning which is derived from first principles, has the same set of global optima as the CMPE task, and operates exclusively on the primal variables. It also uses only one hyperparameter in the continuous case and two hyperparameters in the discrete case. Experimentally, we evaluated our new self-supervised method with penalty-based and Lagrangian duality-based methods proposed in literature and found that our method is often superior in terms of optimality gap and training time (also requires less hyperparameter tuning) to the Lagrangian duality-based methods and also superior in terms of optimality gap and the number of constraint violations to the penalty-based methods.

Our proposed method has several limitations and we will address them in future work. First, it requires a bound for  $\alpha_x$ . This bound is easy to obtain for graphical models/multilinear objectives but may not be straightforward to obtain for arbitrary non-convex functions. Second, the ideal objective in the infeasible region should be proportional to  $g_x(y)$  but our method uses  $\alpha_x(f_x(y) + g_x(y))$ .

## ACKNOWLEDGMENTS

This work was supported in part by the DARPA Perceptually-Enabled Task Guidance (PTG) Program under contract number HR00112220005, by the DARPA Assured Neuro Symbolic Learning and Reasoning (ANSR) under contract number HR001122S0039 and by the National Science Foundation grant IIS-1652835.

## Bibliography

- Achterberg, T. (2009). Scip: solving constraint integer programs. *Mathematical Programming Computation*, 1:1–41.
- Achterberg, T., Berthold, T., Koch, T., and Wolter, K. (2008). Constraint integer programming: A new approach to integrate cp and mip. In *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems: 5th International Conference, CPAIOR 2008 Paris, France, May 20-23, 2008 Proceedings 5*, pages 6–20. Springer.
- Choi, A. and Darwiche, A. (2011). Relax, compensate and then recover. In Onada, T., Bekki, D., and McCready, E., editors, *New Frontiers in Artificial Intelligence*, pages 167–180, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Choi, A., Xue, Y., and Darwiche, A. (2012). Same-decision probability: A confidence measure for threshold-based decisions. *International Journal of Approximate Reasoning*, 53(9):1415–1428.
- Choi, Y., Vergari, A., and Van den Broeck, G. (2020). Probabilistic circuits: A unifying framework for tractable probabilistic models. *UCLA*. URL: <http://starai.cs.ucla.edu/papers/ProbCirc20.pdf>.
- Cui, Z., Wang, H., Gao, T., Talamadupula, K., and Ji, Q. (2022). Variational message passing neural network for maximum-a-posteriori (map) inference. In *Uncertainty in Artificial Intelligence*, pages 464–474. PMLR.
- Darwiche, A. (2009). *Modeling and Reasoning with Bayesian Networks*. Cambridge University Press.
- Darwiche, A. and Hirth, A. (2020). On the reasons behind decisions. In *Twenty Fourth European Conference on Artificial Intelligence*, volume 325 of *Frontiers in Artificial Intelligence and Applications*, pages 712–720. IOS Press.
- Darwiche, A. and Hirth, A. (2023). On the (complete) reasons behind decisions. *Journal of Logic, Language and Information*, 32(1):63–88.
- Dechter, R. (1999). Bucket elimination: A unifying framework for reasoning. *Artificial Intelligence*, 113:41–85.
- Dechter, R. and Rish, I. (2003). Mini-buckets: A general scheme for bounded inference. *Journal of the ACM (JACM)*, 50(2):107–153.
- Donti, P. L., Rolnick, D., and Kolter, J. Z. (2021). Dc3: A learning method for optimization with hard constraints. *arXiv preprint arXiv:2104.12225*.
- Elidan, G. and Globerson, A. (2010). *The 2010 UAI Approximate Inference Challenge*. Published: Available online at: <http://www.cs.huji.ac.il/project/UAI10/index.php>.
- Fioretto, F., Mak, T. W., and Van Hentenryck, P. (2020). Predicting ac optimal power flows: Combining deep learning and lagrangian dual methods. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 630–637.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. (2017). Neural message passing for quantum chemistry. In *International conference on machine learning*, pages 1263–1272. PMLR.
- Globerson, A. and Jaakkola, T. (2007). Fixing max-product: Convergent message passing algorithms for map lp-relaxations. *Advances in neural information processing systems*, 20:553–560.
- Gogate, V. (2014). Results of the 2014 UAI competition. <https://personal.utdallas.edu/~vibhav.gogate/uai14-competition/index.html>.
- Gogate, V. (2016). Results of the 2016 UAI competition. <https://personal.utdallas.edu/~vibhav.gogate/uai16-competition/index.html>.
- Gurobi Optimization, L. (2021). Gurobi optimizer reference manual.
- Horst, R. and Tuy, H. (1996). *Global Optimization: Deterministic Approaches*. Springer Berlin Heidelberg.
- Ihler, A. T., Flerova, N., Dechter, R., and Otten, L. (2012). Join-graph based cost-shifting schemes. *arXiv preprint arXiv:1210.4878*.
- Kingma, D. P. and Ba, J. (2014). Adam: A Method for Stochastic Optimization. *arXiv*.
- Koller, D. and Friedman, N. (2009). *Probabilistic graphical models: principles and techniques*. MIT press.
- Komodakis, N., Paragios, N., and Tziritas, G. (2007). Mrf optimization via dual decomposition: Message-passing revisited. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE.
- Kotary, J., Fioretto, F., and Hentenryck, P. V. (2021). Learning hard optimization problems: A data generation perspective. *arXiv preprint arXiv: Arxiv-2106.02601*.
- Kuck, J., Chakraborty, S., Tang, H., Luo, R., Song, J., Sabharwal, A., and Ermon, S. (2020). Belief propagation neural networks. *Advances in Neural Information Processing Systems*, 33:667–678.
- LeCun, Y. and Cortes, C. (2010). MNIST handwritten digit database.

- Liu, T. and Cherian, A. (2023). Learning a constrained optimizer: A primal method. In *AAAI 2023 Bridge on Constraint Programming and Machine Learning*.
- Lowd, D. and Davis, J. (2010). Learning Markov Network Structure with Decision Trees. In *2010 IEEE International Conference on Data Mining*, pages 334–343. IEEE.
- Marinescu, R. and Dechter, R. (2009). Memory intensive AND/OR search for combinatorial optimization in graphical models. *AI Journal*, 173(16-17):1492–1524.
- Marinescu, R. and Dechter, R. (2012). Best-First AND/OR Search for Most Probable Explanations. *CoRR*, abs/1206.5268.
- Nellikath, R. and Chatzivasileiadis, S. (2021). Physics-informed neural networks for ac optimal power flow. *arXiv preprint arXiv: Arxiv-2110.02672*.
- Nocedal, J. and Wright, S. J. (2006). *Numerical Optimization*. Springer, New York, NY, USA, 2e edition.
- Park, S. and Van Hentenryck, P. (2022). Self-supervised primal-dual learning for constrained optimization. *arXiv preprint arXiv:2208.09046*.
- Rahman, T., Kothalkar, P., and Gogate, V. (2014). Cutset networks: A simple, tractable, and scalable approach for improving the accuracy of chow-liu trees. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2014, Nancy, France, September 15-19, 2014. Proceedings, Part II 14*, pages 630–645. Springer.
- Rahman, T., Rouhani, S., and Gogate, V. (2021). Novel upper bounds for the constrained most probable explanation task. *Advances in Neural Information Processing Systems*, 34:9613–9624.
- Rouhani, S., Rahman, T., and Gogate, V. (2018). Algorithms for the nearest assignment problem. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 5096–5102. ijcai.org.
- Rouhani, S., Rahman, T., and Gogate, V. (2020). A novel approach for constrained optimization in graphical models. *Advances in Neural Information Processing Systems*, 33:11949–11960.
- Satorras, V. G. and Welling, M. (2021). Neural enhanced belief propagation on factor graphs. In *International Conference on Artificial Intelligence and Statistics*, pages 685–693. PMLR.
- Sherali, H. D. and Adams, W. P. (2009). A reformulation-linearization technique (rlt) for semi-infinite and convex programs under mixed 0-1 and general discrete restrictions. *Discrete Applied Mathematics*, 157(6):1319–1333. Reformulation Techniques and Mathematical Programming.
- Sherali, H. D. and Tuncbilek, C. H. (1992). A global optimization algorithm for polynomial programming problems using a reformulation-linearization technique. *Journal of Global Optimization*, 2:101–112.
- Ucla-Starai (2023). Density-Estimation-Datasets. [Online; accessed 17. May 2023].
- Van Haaren, J. and Davis, J. (2012). Markov Network Structure Learning: A Randomized Feature Generation Approach. *AAAI*, 26(1):1148–1154.
- Wainwright, M. J., Jaakkola, T. S., and Willsky, A. S. (2005). Map estimation via agreement on trees: message-passing and linear programming. *IEEE transactions on information theory*, 51(11):3697–3717.
- Wu, B., Shen, L., Zhang, T., and Ghanem, B. (2020). Map inference via l2 sphere linear program reformulation. *International Journal of Computer Vision*, 128(7):1913–1936.
- Zamzam, A. and Baker, K. (2019). Learning optimal solutions for extremely fast ac optimal power flow. *arXiv preprint arXiv: Arxiv-1910.01213*.
- Zhang, Z., Wu, F., and Lee, W. S. (2020). Factor graph neural networks. *Advances in Neural Information Processing Systems*, 33:8577–8587.

## Checklist

1. For all models and algorithms presented, check if you include:
  - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
  - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
  - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes]
2. For any theoretical claim, check if you include:
  - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
  - (b) Complete proofs of all theoretical results. [Yes]
  - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
  - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
  - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
  - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
  - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
  - (a) Citations of the creator If your work uses existing assets. [Yes]
  - (b) The license information of the assets, if applicable. [Not Applicable]
  - (c) New assets either in the supplemental material or as a URL, if applicable. [Yes]
  - (d) Information about consent from data providers/curators. [Not Applicable]
  - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
  - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
  - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
  - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

## A DERIVING UPPER BOUNDS FOR $p_x^*$ AND LOWER BOUNDS FOR $q_x^*$

To compute the value of  $\alpha_x$ , we utilize the following equation:

$$\alpha_x > \frac{p_x^*}{q_x^*} \quad (17)$$

We choose to set a lower bound for the denominator  $q_x^*$  and an upper bound for the numerator  $p_x^*$  due to the computational challenges of finding exact solutions for the CMPE task.

To estimate an upper bound for the optimal value ( $p_x^*$ ) of the constrained optimization problem

$$\min_{\hat{y}} f_x(\hat{y}) \text{ s.t. } g_x(\hat{y}) \leq 0, \quad (18)$$

we begin by seeking a loose upper bound through solving the unconstrained task  $\max_y f_x(\hat{y})$  by utilizing mini-bucket elimination (Dechter and Rish, 2003). Subsequently, feasible solutions are tracked during batch-style gradient descent to refine the initial upper bound (note that the weight of any feasible solution is an upper bound on  $p_x^*$ ). For each iteration, the feasible solution with the optimal objective value for each example is stored and subsequently utilized.

To derive a lower bound for  $q_x^*$ , which represents the optimal solution for the following constrained optimization problem,

$$\min_{\hat{y}} f_x(\hat{y}) + g_x(\hat{y}) \text{ s.t. } g_x(\hat{y}) > 0, \quad (19)$$

we can employ the methodologies delineated in Rahman et al. (2021). These techniques provide a mechanism for either upper bounding or lower bounding the CMPE task, contingent on whether it is formulated as a maximization or minimization problem, respectively.

To establish a lower bound for  $q_x^*$ , the constrained optimization task is initially transformed into an unconstrained formulation via Lagrange Relaxation. This results in the following optimization problem:

$$\max_{\mu \geq 0} \min_{\hat{y}} f_x(\hat{y}) + (1 - \mu) \times g_x(\hat{y}) \quad (20)$$

Here,  $\mu$  denotes the Lagrangian multiplier. By addressing this dual optimization problem, we enhance the precision of the lower bound for  $q_x^*$ . For the inner minimization task, the mini-bucket elimination method is employed. The outer maximization is solved through the utilization of sub-gradient descent.

## B EXTENSIONS

**Adding a Penalty for Constraint Violations.** A penalty of the form  $\max\{0, g_x(\hat{y})\}^2$  can be easily added to the loss function as described by the following equation

$$\mathcal{L}_x(\hat{y}) = \begin{cases} f_x(\hat{y}) & \text{if } g_x(\hat{y}) \leq 0 \\ \alpha_x(f_x(\hat{y}) + g_x(\hat{y})) + \rho \max\{0, g_x(\hat{y})\}^2 & \text{if } g_x(\hat{y}) > 0 \end{cases}$$

where  $\rho \geq 0$  is a hyperparameter.

## C EXPERIMENTAL SETUP AND DETAILS

### C.A Dataset and Model Description

Table ST4 provides a comprehensive overview of the characteristics of each binary dataset, including the number of variables and functions present in each dataset. These datasets were specifically chosen to provide diverse and representative examples for evaluating the performance and scalability of our algorithms.

We used the following two classes of Markov networks from the UAI competitions (Elidan and Globerson, 2010; Gogate, 2014, 2016): Ising models (Grids) and Image Segmentation networks. Specifically, we used the Grids\_17 and Grids\_18 networks and Segmentation\_12, Segmentation\_14 and Segmentation\_15 networks.

We learned MPE tractable cutset networks without latent variables using the scheme of Rahman et al. (2014) on five high-dimensional datasets: DNA (Van Haaren and Davis, 2012; UCLA-Starai, 2023), NewsGroup (c20ng) (Lowd and Davis, 2010; UCLA-Starai, 2023), WebKB1 (cwebkb) (Lowd and Davis, 2010; UCLA-Starai, 2023), AD (Van Haaren and Davis, 2012; UCLA-Starai, 2023), and BBC (Van Haaren and Davis, 2012; UCLA-Starai, 2023). These datasets are widely used in the probabilistic circuits literature (Lowd and Davis, 2010). Note that CMPE is intractable on these models even though MPE is tractable.

Table ST4: Dataset and Model Descriptions

Dataset	Number of Variables	Number of Functions
Tractable Probabilistic Circuits		
AD	1556	1556
BBC	1058	1058
20NewsGroup	910	910
WebKB	839	839
DNA	180	180
High Tree-Width Markov Networks		
Grids17	400	1160
Grids18	400	1160
Segmentation12	229	851
Segmentation14	226	845
Segmentation15	232	863

### C.B Data Generation

Recall that the CMPE problem uses two Markov networks  $M_1$  and  $M_2$ , and a value of  $q$ . We used the original Markov networks (chosen from the UAI competitions or learned from data) as  $M_1$  and generated  $M_2$  by adding a value  $v$ , which was randomly sampled from a normal distribution with mean 0 and variance 0.1, to each entry in each potential in  $M_1$ . We used the following strategy to generate  $q$ . We generated 100 random samples from  $M_2$ , sorted them according to their weight w.r.t.  $M_2$ , and then chose the 10th, 30th, 60th, and 90th sample as a value for  $q$ . At a high level, as we go from the 10th sample to the 90th sample, namely as  $q$  increases, the constraint (weight w.r.t.  $M_2$  is less than or equal to  $q$ ) becomes less restrictive. In other words, as we increase  $q$ , the set of feasible solutions increases (or stays the same). For each value of  $q$ , we use 60% of the variables as evidence variables  $\mathbf{X}$  and the remaining as  $\mathbf{Y}$ .

For each CMPE problem, we generated 10000 samples, used the first 9000 samples for training and the remaining 1000 samples for testing. For the supervised methods, we generated the optimal assignment to  $\mathbf{Y}$  using an integer linear programming solver called SCIP (Achterberg, 2009).

For our proposed scheme, which we call SS-CMPE, we used approach described in Section A to find the upper bound of  $p_{\mathbf{x}}^*$  and the lower bound of  $q_{\mathbf{x}}^*$

Note that CMPE is a much harder task than MPE. Our scheme can be easily adapted to MPE, all we have to do is use  $f$  to yield a supervised scheme.

### C.C Architecture Design and Training Procedure

In our experimental evaluations, we employed a Multi-Layer Perceptron (MLP) with a Rectified Linear Unit (ReLU) activation function for all hidden layers. The final layer of the MLP utilized a sigmoid activation function, as it was necessary to obtain outputs within the range of  $[0, 1]$  for all our experiments. Each fully connected neural network in our study consisted of three hidden layers with respective sizes of  $[128, 256, \text{and } 512]$ . We maintained this consistent architecture across all our supervised (Zamzam and Baker, 2019; Nellikkath and Chatzivasileiadis, 2021) and self-supervised (Park and Van Hentenryck, 2022; Donti et al., 2021) methods. It is important to highlight that in the adversarial modification

experiments, the neural network possessed an equal number of inputs and outputs, specifically set to  $28 \times 28$  (size of an image in MNIST). However, in the remaining two experiments concerning probabilistic graphical models, the input size was the number of evidence variables ( $|\mathbf{X}|$ ), while the output size was  $|\mathbf{Y}|$ .

For PDL (Park and Van Hentenryck, 2022), the dual network had one hidden layer with 128 nodes. The number of outputs of the dual network corresponds to the number of constraints in the optimization problem. It is worth emphasizing that our method is not constrained to the usage of Multi-Layer Perceptrons (MLPs) exclusively, and we have the flexibility to explore various neural network architectures. This flexibility allows us to consider and utilize alternative architectures that may better suit the requirements and objectives of other optimization tasks.

Regarding the training process, all methods underwent 300 epochs using the Adam optimizer (Kingma and Ba, 2014) with a learning rate of  $1e^{-3}$ . We employed a Learning Rate Scheduler to dynamically adapt the learning rate as the loss reaches a plateau. The training and testing processes for all models were conducted on a single NVIDIA A40 GPU.

### C.D Hyper-parameters

The number of instances in the minibatch was set to 128 for all the experiments. We decay the learning rate in all the experiments by 0.9 when the loss becomes a plateau. Given the empirical observations that learning rate decay often leads to early convergence in most cases and does not yield beneficial results for the supervised baselines, we have made the decision not to apply learning rate decay to these methods. This choice is based on the understanding that the baselines perform optimally without this particular form of learning rate adjustment. For detailed information regarding the hyper-parameters utilized in the benchmarking methods, we refer readers to the corresponding papers associated with each method. As stated in the main text, the optimal hyperparameters were determined using a grid search approach. For the  $SS-CMPE$  method, for each dataset, the hyperparameters were selected from the following available options -

- $\beta$  -  $\{0.1, 1.0, 2.0, 5.0, 10.0, 20.0\}$
- $\rho$  -  $\{0.01, 0.1, 1, 10, 100\}$

In the optimization problem of  $SS-CMPE$ , the parameter  $\rho$  is employed to penalize the violation of constraints. The methodology for this approach, denoted as  $SS-CMPE_{pen}$ , is explained in detail in Section B. The corresponding experiments are presented in tables ST8 through ST17.

### C.E The Loss Function: Competing Methods

We evaluated eight different loss functions, including our method to train a deep neural network to solve our CMPE tasks. The loss functions used for supervised training are 1) Mean-Squared-Error (MSE), and 2) Mean-Absolute-Error (MAE). Both of these losses were then extended to incorporate penalty terms as suggested by (Nellikath and Chatzivasileiadis, 2021). We denote them as SL+Penalty and MAE+Penalty. We evaluated the self-supervised loss proposed by (Donti et al., 2021) ( $SSL_{pen}$ ) and by (Park and Van Hentenryck, 2022) (PDL). Finally, we extend our self-supervised CMPE loss function to incorporate the penalty term  $\max\{0, g_x(\hat{y})\}^2$  (see section B). We denote it as  $SS-CMPE_{pen}$ .

## D EXAMINING THE INFLUENCE OF $q$ : EVALUATING THE PERFORMANCE OF OUR PROPOSED METHOD FOR CHALLENGING PROBLEM INSTANCES

To determine the value of  $q$ , a total of 100 random samples were generated, and their weights were calculated. Subsequently, the samples were sorted based on their weight in ascending order. The weight values corresponding to the 10th, 30th, 60th, and 90th samples were then chosen as the values for  $q$ . For each value of  $q$ , we compare the average gap and violations obtained by our method ( $SS-CMPE$  and  $SS-CMPE_{pen}$ ) against six other supervised and self-supervised methods. Tables ST8 through ST17 show the scores obtained by each of the eight methods along with their standard deviations on the generated test problems. This study investigates the performance of each method in finding near-optimal feasible solutions for difficult problems which is directly controlled by the percentile rank of  $q$ ; problems with a  $q$  value in the 10th and 30th percentile are considered harder problems to solve as the size of the feasible region is considerably smaller than the size of the infeasible region. As a result all methods have higher violations on these problems than the problems with a  $q$  value in the 60th and 90th percentile.

Table ST5 presents a summary of the performances of  $SS-CMPE$  and other *supervised* methods based on their average gap and the average number of violations on the test data for different values of  $q$ . We compute the minimum gap and violations achieved among the four supervised methods MSE, SL+Penalty, MAE, and MAE+Penalty and label them as the `best supervised` method. We choose the minimum gap and violations among  $SS-CMPE$  and  $SS-CMPE_{pen}$  and label them under the unified term `best SS-CMPE`. We observe that `best SS-CMPE` consistently has significantly lower violations than the best supervised method in all the problem instances, and its gap is often comparable to the gap achieved by the best supervised method, winning when compared to the average gap.

Table ST6 presents a similar summary of the performances of  $SS-CMPE$  and other *self-supervised* methods. We compute the minimum gap and violations achieved among the 2 self-supervised methods  $SSL_{pen}$  and PDL and label them as the `best SSL` method. As before, we choose the minimum gap and violations among  $SS-CMPE$  and  $SS-CMPE_{pen}$  and label them under the unified term `best SS-CMPE`. Although self-supervised methods have larger gaps compared to supervised methods but lesser violations, we observe that  $SS-CMPE$  continues to consistently achieve significantly lower violations than the best performing self-supervised method in all the problem instances, and its gap is often comparable to the gap achieved by the best supervised method, winning when compared to the average gap.

Finally, in table ST7, we present a quick summary of the performances of our best  $SS-CMPE$  method vs all `other` methods. We choose the minimum gap and violations achieved among the 6 `other` supervised and self-supervised methods and the minimum gap and violations among  $SS-CMPE$  and  $SS-CMPE_{pen}$ . In all problems and  $q$  values, the best  $SS-CMPE$  has significantly lower violations compared to other methods while having a very competitive gap.

#### D.A The Feasible-Only Optimality Gaps: Comparing Self-Supervised Approaches

From the results presented in tables ST5 through ST17, we observe that self-supervised approaches produce more feasible solutions compared to supervised approaches. In this section, we present the results of a controlled study that shows how each of the self-supervised approaches perform in terms of finding optimal solutions in the feasible region.

We selected a subset of problems from the test set on which all self-supervised methods, namely,  $SSL_{pen}$  (Donti et al., 2021), PDL (Park and Van Hentenryck, 2022) and our method  $SS-CMPE$  and  $SS-CMPE_{pen}$ , obtained feasible solutions and this was done for each possible value of  $q$ . We then computed their gaps and compare them via figure SF4. Among the three methods analyzed,  $SS-CMPE$  and  $SS-CMPE_{pen}$  consistently exhibits superior performance across the majority of cases. Its optimality gaps are significantly smaller compared to the other two methods. This finding suggests that  $SS-CMPE$  is more effective in minimizing the objective value and achieving solutions closer to optimality for the given examples.



Table ST5: Summary: best SS-CMPE vs other supervised methods including MSE, SL+Penalty, MAE, and MAE+Penalty. Bold represents the minimum gap, while underlined means the least violations

q		10		30		60		90	
Models /Dataset	Gap / Violations	best SS-CMPE	best supervised	best SS-CMPE	best supervised	best SS-CMPE	best supervised	best SS-CMPE	best supervised
Segment-12	gap	0.057	<b>0.054</b>	<b>0.051</b>	0.052	0.053	<b>0.051</b>	<b>0.050</b>	0.051
	violations	<u>0.152</u>	0.511	<u>0.166</u>	0.500	<u>0.084</u>	0.348	<u>0.021</u>	0.131
Segment-14	gap	0.050	<b>0.049</b>	<b>0.047</b>	0.049	0.048	0.048	0.051	0.051
	violations	<u>0.134</u>	0.691	<u>0.088</u>	0.616	<u>0.066</u>	0.410	<u>0.046</u>	0.207
Segment-15	gap	0.051	0.051	<b>0.050</b>	0.051	0.052	<b>0.051</b>	<b>0.051</b>	0.052
	violations	<u>0.060</u>	0.570	<u>0.060</u>	0.417	<u>0.049</u>	0.248	<u>0.001</u>	0.061
Grids-17	gap	0.072	<b>0.054</b>	0.069	<b>0.057</b>	<b>0.066</b>	0.067	0.063	<b>0.058</b>
	violations	<u>0.035</u>	0.304	<u>0.013</u>	0.125	<u>0.002</u>	0.044	<u>0.001</u>	0.002
Grids-18	gap	0.064	<b>0.056</b>	0.067	<b>0.060</b>	<b>0.060</b>	0.065	0.065	<b>0.064</b>
	violations	<u>0.017</u>	0.210	<u>0.019</u>	0.087	<u>0.000</u>	0.025	<u>0.000</u>	0.015
DNA	gap	0.138	<b>0.135</b>	0.138	<b>0.136</b>	0.137	<b>0.136</b>	0.139	<b>0.137</b>
	violations	<u>0.013</u>	0.434	<u>0.002</u>	0.448	<u>0.001</u>	0.281	<u>0.001</u>	0.089
20NewsGr.	gap	<b>0.043</b>	0.044	<b>0.045</b>	0.046	<b>0.044</b>	0.046	0.044	0.044
	violations	<u>0.069</u>	0.455	<u>0.054</u>	0.176	<u>0.007</u>	0.046	<u>0.001</u>	0.001
WebKB	gap	0.059	<b>0.054</b>	0.058	<b>0.054</b>	0.056	<b>0.054</b>	<b>0.053</b>	0.054
	violations	<u>0.074</u>	0.471	<u>0.029</u>	0.378	<u>0.001</u>	0.174	<u>0.001</u>	0.018
BBC	gap	0.038	<b>0.036</b>	0.043	<b>0.036</b>	0.042	<b>0.037</b>	0.040	<b>0.037</b>
	violations	<u>0.074</u>	0.657	<u>0.067</u>	0.557	<u>0.056</u>	0.384	<u>0.002</u>	0.151
Ad	gap	<b>0.129</b>	0.204	<b>0.131</b>	0.201	<b>0.130</b>	0.204	<b>0.131</b>	0.213
	violations	<u>0.017</u>	0.085	<u>0.004</u>	0.041	<u>0.000</u>	0.021	<u>0.000</u>	0.005
Average	gap	<b>0.070</b>	0.074	<b>0.070</b>	0.074	<b>0.069</b>	0.076	<b>0.069</b>	0.076
	violations	<u>0.065</u>	0.439	<u>0.050</u>	0.335	<u>0.027</u>	0.198	<u>0.007</u>	0.068

Table ST6: Summary: best SS-CMPE vs other self-supervised methods including  $SL_{pen}$ , and PDL. Bold represents the minimum gap, while underlined means the least violations

q		10		30		60		90	
Models/ Dataset	Gap/ Violations	best SS-CMPE	best SSL	best SS-CMPE	best SSL	best SS-CMPE	best SSL	best SS-CMPE	best SSL
Segment-12	gap	0.057	<b>0.054</b>	<b>0.051</b>	0.052	0.053	<b>0.051</b>	<b>0.050</b>	0.051
	violations	<u>0.152</u>	0.545	<u>0.166</u>	0.622	<u>0.084</u>	0.486	<u>0.021</u>	0.163
Segment-14	gap	<b>0.050</b>	0.058	<b>0.047</b>	0.050	<b>0.048</b>	0.050	<b>0.051</b>	0.053
	violations	<u>0.134</u>	0.507	<u>0.088</u>	0.414	<u>0.066</u>	0.394	<u>0.046</u>	0.207
Segment-15	gap	0.051	0.051	<b>0.050</b>	0.053	0.052	<b>0.051</b>	<b>0.051</b>	0.054
	violations	<u>0.060</u>	0.676	<u>0.060</u>	0.360	<u>0.049</u>	0.274	<u>0.001</u>	0.086
Grids-17	gap	<b>0.072</b>	0.086	<b>0.069</b>	0.079	<b>0.066</b>	0.093	<b>0.063</b>	0.087
	violations	<u>0.035</u>	0.043	<u>0.013</u>	0.003	0.002	<u>0.001</u>	<u>0.001</u>	0.014
Grids-18	gap	<b>0.064</b>	0.105	<b>0.067</b>	0.074	<b>0.060</b>	0.093	<b>0.065</b>	0.118
	violations	<u>0.017</u>	0.060	0.019	<u>0.001</u>	<u>0.000</u>	0.003	<u>0.000</u>	0.005
DNA	gap	<b>0.138</b>	0.140	<b>0.138</b>	0.141	<b>0.137</b>	0.139	<b>0.139</b>	0.143
	violations	<u>0.013</u>	0.048	<u>0.002</u>	0.062	0.001	0.003	<u>0.001</u>	0.004
20NewsGr	gap	0.043	0.043	<b>0.045</b>	0.046	<b>0.044</b>	0.045	<b>0.044</b>	0.046
	violations	<u>0.069</u>	0.278	<u>0.054</u>	0.129	<u>0.007</u>	0.024	0.001	0.001
WebKB	gap	0.059	<b>0.058</b>	0.058	<b>0.057</b>	<b>0.056</b>	0.057	<b>0.053</b>	0.057
	violations	<u>0.074</u>	0.149	<u>0.029</u>	0.096	<u>0.001</u>	0.056	<u>0.001</u>	0.013
BBC	gap	<b>0.038</b>	0.041	<b>0.043</b>	0.042	0.042	<b>0.038</b>	0.040	<b>0.039</b>
	violations	<u>0.074</u>	0.336	<u>0.067</u>	0.160	<u>0.056</u>	0.149	<u>0.002</u>	0.029
Ad	gap	<b>0.129</b>	0.135	<b>0.131</b>	0.140	<b>0.130</b>	0.142	<b>0.131</b>	0.134
	violations	<u>0.017</u>	0.055	<u>0.004</u>	0.006	<u>0.000</u>	0.013	<u>0.000</u>	0.004
Average	gap	<b>0.070</b>	0.077	<b>0.070</b>	0.073	<b>0.069</b>	0.076	<b>0.069</b>	0.078
	violations	<u>0.065</u>	0.270	<u>0.050</u>	0.185	<u>0.027</u>	0.140	<u>0.007</u>	0.053

Table ST7: Summary: best SS-CMPE has significantly lower violations compared to other methods on all the problems and over all the chosen  $q$  values. It has comparable gap to the other methods.

q		10		30		60		90	
Models / Datasets	Gap / Violations	best SS-CMPE	others	best SS-CMPE	others	best SS-CMPE	others	best SS-CMPE	others
Segment-12	Gap	0.057	<b>0.054</b>	<b>0.051</b>	0.052	0.053	<b>0.051</b>	<b>0.050</b>	0.051
	Violations	<b>0.152</b>	0.511	<b>0.166</b>	0.500	<b>0.084</b>	0.348	<b>0.021</b>	0.131
Segment-14	Gap	0.050	<b>0.049</b>	<b>0.047</b>	0.049	0.048	0.048	0.051	0.051
	Violations	<b>0.134</b>	0.507	<b>0.088</b>	0.414	<b>0.066</b>	0.394	<b>0.046</b>	0.207
Segment-15	Gap	0.051	0.051	<b>0.050</b>	0.051	0.052	<b>0.051</b>	<b>0.051</b>	0.052
	Violations	<b>0.060</b>	0.570	<b>0.060</b>	0.360	<b>0.049</b>	0.248	<b>0.001</b>	0.061
Grids-17	Gap	0.072	<b>0.054</b>	0.069	<b>0.057</b>	<b>0.066</b>	0.067	0.063	<b>0.058</b>
	Violations	<b>0.035</b>	0.043	<b>0.013</b>	0.003	<b>0.002</b>	0.001	<b>0.001</b>	0.002
Grids-18	Gap	0.064	<b>0.056</b>	0.067	<b>0.060</b>	<b>0.060</b>	0.065	0.065	<b>0.064</b>
	Violations	<b>0.017</b>	0.060	<b>0.019</b>	0.001	<b>0.000</b>	0.003	<b>0.000</b>	0.005
DNA	Gap	0.138	<b>0.135</b>	0.138	<b>0.136</b>	0.137	<b>0.136</b>	0.139	<b>0.137</b>
	Violations	<b>0.013</b>	0.048	<b>0.002</b>	0.062	<b>0.001</b>	0.003	<b>0.001</b>	0.004
20NewsGr	Gap	0.043	0.043	<b>0.045</b>	0.046	<b>0.044</b>	0.045	<b>0.044</b>	0.044
	Violations	<b>0.069</b>	0.278	<b>0.054</b>	0.129	<b>0.007</b>	0.024	<b>0.001</b>	0.001
WebKB	Gap	0.059	<b>0.054</b>	0.058	<b>0.054</b>	0.056	<b>0.054</b>	<b>0.053</b>	0.054
	Violations	<b>0.074</b>	0.149	<b>0.029</b>	0.096	<b>0.001</b>	0.056	<b>0.001</b>	0.013
BBC	Gap	0.038	<b>0.036</b>	0.043	<b>0.036</b>	0.042	<b>0.037</b>	0.040	<b>0.037</b>
	Violations	<b>0.074</b>	0.336	<b>0.067</b>	0.160	<b>0.056</b>	0.149	<b>0.002</b>	0.029
Ad	Gap	<b>0.129</b>	0.135	<b>0.131</b>	0.140	<b>0.130</b>	0.142	<b>0.131</b>	0.134
	Violations	<b>0.017</b>	0.055	<b>0.004</b>	0.006	<b>0.000</b>	0.013	<b>0.000</b>	0.004
Average	Gap	0.070	<b>0.067</b>	0.070	<b>0.068</b>	<b>0.069</b>	0.070	0.069	<b>0.068</b>
	Violations	<b>0.065</b>	0.256	<b>0.050</b>	0.173	<b>0.027</b>	0.124	<b>0.007</b>	0.046

## Learning to Solve the Constrained Most Probable Explanation Task in Probabilistic Graphical Models



Figure SF4: Illustration of the optimality gap for self-supervised methods (on feasible examples only) for all approaches. Lower is better.

Table ST8: Average gap and constraint violations over test samples for models applied to the Segmentation12 Dataset for different q values. The plot displays the mean values of the average gap and constraint violations, with standard deviations denoted by  $\pm$

q		10	30	60	90
ILP Obj		491.150	476.654	467.913	461.967
MAE	Gap	0.064 $\pm$ 0.051	0.061 $\pm$ 0.049	0.053 $\pm$ 0.042	0.052 $\pm$ 0.040
	Violations	0.569 $\pm$ 0.495	0.545 $\pm$ 0.498	0.430 $\pm$ 0.495	0.131 $\pm$ 0.337
MSE	Gap	0.054 $\pm$ 0.042	0.053 $\pm$ 0.042	0.051 $\pm$ 0.041	0.051 $\pm$ 0.041
	Violations	0.776 $\pm$ 0.417	0.580 $\pm$ 0.494	0.486 $\pm$ 0.500	0.186 $\pm$ 0.390
MAE+Penalty	Gap	0.064 $\pm$ 0.050	0.061 $\pm$ 0.049	0.059 $\pm$ 0.046	0.052 $\pm$ 0.043
	Violations	0.511 $\pm$ 0.500	0.500 $\pm$ 0.500	0.348 $\pm$ 0.476	0.140 $\pm$ 0.347
SL+Penalty	Gap	0.054 $\pm$ 0.042	0.052 $\pm$ 0.041	0.051 $\pm$ 0.040	0.051 $\pm$ 0.042
	Violations	0.651 $\pm$ 0.477	0.505 $\pm$ 0.500	0.486 $\pm$ 0.500	0.186 $\pm$ 0.390
SSL <sub>pen</sub>	Gap	0.054 $\pm$ 0.043	0.052 $\pm$ 0.040	0.051 $\pm$ 0.041	0.051 $\pm$ 0.040
	Violations	0.790 $\pm$ 0.407	0.622 $\pm$ 0.485	0.486 $\pm$ 0.500	0.186 $\pm$ 0.390
PDL	Gap	0.063 $\pm$ 0.050	0.052 $\pm$ 0.041	0.052 $\pm$ 0.041	0.051 $\pm$ 0.041
	Violations	0.545 $\pm$ 0.498	0.622 $\pm$ 0.485	0.517 $\pm$ 0.500	0.163 $\pm$ 0.369
SS-CMPE	Gap	0.057 $\pm$ 0.044	0.051 $\pm$ 0.040	0.053 $\pm$ 0.043	0.050 $\pm$ 0.041
	Violations	0.503 $\pm$ 0.293	0.346 $\pm$ 0.485	0.257 $\pm$ 0.437	0.104 $\pm$ 0.305
SS-CMPE <sub>pen</sub>	Gap	0.058 $\pm$ 0.043	0.052 $\pm$ 0.040	0.053 $\pm$ 0.043	0.051 $\pm$ 0.041
	Violations	0.152 $\pm$ 0.359	0.166 $\pm$ 0.372	0.084 $\pm$ 0.277	0.021 $\pm$ 0.143

Table ST9: Average gap and constraint violations over test samples for models applied to the Segmentation14 Dataset for different q values. The plot displays the mean values of the average gap and constraint violations, with standard deviations denoted by  $\pm$ .

q		10	30	60	90
ILP Obj		493.647	482.837	476.145	470.485
MAE	Gap	0.067 $\pm$ 0.051	0.062 $\pm$ 0.048	0.062 $\pm$ 0.047	0.051 $\pm$ 0.040
	Violations	0.691 $\pm$ 0.462	0.623 $\pm$ 0.485	0.435 $\pm$ 0.496	0.271 $\pm$ 0.444
MSE	Gap	0.051 $\pm$ 0.039	0.049 $\pm$ 0.038	0.048 $\pm$ 0.037	0.053 $\pm$ 0.041
	Violations	0.810 $\pm$ 0.392	0.818 $\pm$ 0.386	0.606 $\pm$ 0.489	0.252 $\pm$ 0.434
MAE+Penalty	Gap	0.065 $\pm$ 0.051	0.061 $\pm$ 0.048	0.060 $\pm$ 0.046	0.054 $\pm$ 0.043
	Violations	0.693 $\pm$ 0.461	0.616 $\pm$ 0.487	0.410 $\pm$ 0.492	0.207 $\pm$ 0.405
SL+Penalty	Gap	0.049 $\pm$ 0.039	0.049 $\pm$ 0.039	0.049 $\pm$ 0.037	0.054 $\pm$ 0.041
	Violations	0.810 $\pm$ 0.392	0.803 $\pm$ 0.397	0.601 $\pm$ 0.490	0.215 $\pm$ 0.411
SSL <sub>pen</sub>	Gap	0.061 $\pm$ 0.047	0.050 $\pm$ 0.038	0.050 $\pm$ 0.039	0.053 $\pm$ 0.042
	Violations	0.590 $\pm$ 0.492	0.618 $\pm$ 0.486	0.394 $\pm$ 0.489	0.207 $\pm$ 0.405
PDL	Gap	0.058 $\pm$ 0.045	0.068 $\pm$ 0.051	0.056 $\pm$ 0.043	0.054 $\pm$ 0.043
	Violations	0.507 $\pm$ 0.500	0.414 $\pm$ 0.493	0.403 $\pm$ 0.491	0.207 $\pm$ 0.405
SS-CMPE	Gap	0.050 $\pm$ 0.038	0.047 $\pm$ 0.037	0.048 $\pm$ 0.037	0.051 $\pm$ 0.040
	Violations	0.502 $\pm$ 0.295	0.444 $\pm$ 0.401	0.309 $\pm$ 0.497	0.150 $\pm$ 0.358
SS-CMPE <sub>pen</sub>	Gap	0.050 $\pm$ 0.039	0.048 $\pm$ 0.038	0.050 $\pm$ 0.037	0.052 $\pm$ 0.042
	Violations	0.134 $\pm$ 0.340	0.088 $\pm$ 0.284	0.066 $\pm$ 0.248	0.046 $\pm$ 0.211

Table ST10: Average gap and constraint violations over test samples for models applied to the Segmentation15 Dataset for different  $q$  values. The plot displays the mean values of the average gap and constraint violations, with standard deviations denoted by  $\pm$ .

q		10	30	60	90
ILP Obj		531.436	520.647	516.797	514.276
MAE	Gap	0.053 $\pm$ 0.043	0.054 $\pm$ 0.043	0.053 $\pm$ 0.042	0.052 $\pm$ 0.041
	Violations	0.570 $\pm$ 0.495	0.417 $\pm$ 0.493	0.248 $\pm$ 0.432	0.061 $\pm$ 0.239
MSE	Gap	0.051 $\pm$ 0.038	0.052 $\pm$ 0.041	0.052 $\pm$ 0.040	0.053 $\pm$ 0.041
	Violations	0.833 $\pm$ 0.373	0.715 $\pm$ 0.452	0.450 $\pm$ 0.498	0.076 $\pm$ 0.265
MAE+Penalty	Gap	0.056 $\pm$ 0.043	0.054 $\pm$ 0.042	0.053 $\pm$ 0.041	0.053 $\pm$ 0.042
	Violations	0.616 $\pm$ 0.486	0.457 $\pm$ 0.498	0.265 $\pm$ 0.441	0.075 $\pm$ 0.263
SL+Penalty	Gap	0.052 $\pm$ 0.040	0.051 $\pm$ 0.040	0.051 $\pm$ 0.040	0.052 $\pm$ 0.041
	Violations	0.833 $\pm$ 0.373	0.715 $\pm$ 0.452	0.450 $\pm$ 0.498	0.076 $\pm$ 0.265
SSL <sub>pen</sub>	Gap	0.053 $\pm$ 0.042	0.059 $\pm$ 0.047	0.052 $\pm$ 0.041	0.054 $\pm$ 0.043
	Violations	0.676 $\pm$ 0.468	0.360 $\pm$ 0.480	0.274 $\pm$ 0.446	0.097 $\pm$ 0.295
PDL	Gap	0.051 $\pm$ 0.040	0.053 $\pm$ 0.043	0.051 $\pm$ 0.040	0.054 $\pm$ 0.043
	Violations	0.698 $\pm$ 0.459	0.461 $\pm$ 0.499	0.392 $\pm$ 0.488	0.086 $\pm$ 0.280
SS-CMPE	Gap	0.051 $\pm$ 0.040	0.050 $\pm$ 0.041	0.052 $\pm$ 0.040	0.051 $\pm$ 0.040
	Violations	0.366 $\pm$ 0.474	0.298 $\pm$ 0.499	0.225 $\pm$ 0.417	0.059 $\pm$ 0.236
SS-CMPE <sub>pen</sub>	Gap	0.052 $\pm$ 0.040	0.052 $\pm$ 0.042	0.052 $\pm$ 0.041	0.051 $\pm$ 0.041
	Violations	0.060 $\pm$ 0.238	0.060 $\pm$ 0.238	0.049 $\pm$ 0.215	0.001 $\pm$ 0.032

Table ST11: Average gap and constraint violations over test samples for models applied to the Grids17 Dataset for different  $q$  values. The plot displays the mean values of the average gap and constraint violations, with standard deviations denoted by  $\pm$ .

q		10	30	60	90
ILP Obj		2892.585191	2884.506703	2877.311872	2878.147272
MAE	Gap	0.119 $\pm$ 0.078	0.133 $\pm$ 0.088	0.125 $\pm$ 0.084	0.114 $\pm$ 0.078
	Violations	0.565 $\pm$ 0.496	0.376 $\pm$ 0.485	0.122 $\pm$ 0.328	0.020 $\pm$ 0.140
MSE	Gap	0.054 $\pm$ 0.041	0.057 $\pm$ 0.045	0.067 $\pm$ 0.054	0.075 $\pm$ 0.056
	Violations	0.314 $\pm$ 0.464	0.152 $\pm$ 0.359	0.044 $\pm$ 0.205	0.013 $\pm$ 0.111
MAE+Penalty	Gap	0.129 $\pm$ 0.081	0.144 $\pm$ 0.089	0.137 $\pm$ 0.087	0.127 $\pm$ 0.085
	Violations	0.534 $\pm$ 0.499	0.376 $\pm$ 0.485	0.142 $\pm$ 0.350	0.019 $\pm$ 0.137
SL+Penalty	Gap	0.054 $\pm$ 0.041	0.059 $\pm$ 0.045	0.069 $\pm$ 0.054	0.058 $\pm$ 0.044
	Violations	0.304 $\pm$ 0.460	0.125 $\pm$ 0.331	0.044 $\pm$ 0.205	0.002 $\pm$ 0.045
SSL <sub>pen</sub>	Gap	0.104 $\pm$ 0.060	0.079 $\pm$ 0.056	0.093 $\pm$ 0.059	0.087 $\pm$ 0.058
	Violations	0.149 $\pm$ 0.357	0.013 $\pm$ 0.113	0.004 $\pm$ 0.067	0.026 $\pm$ 0.159
PDL	Gap	0.086 $\pm$ 0.056	0.087 $\pm$ 0.058	0.109 $\pm$ 0.063	0.112 $\pm$ 0.062
	Violations	0.043 $\pm$ 0.202	0.003 $\pm$ 0.055	0.001 $\pm$ 0.022	0.014 $\pm$ 0.118
SS-CMPE	Gap	0.072 $\pm$ 0.051	0.071 $\pm$ 0.052	0.066 $\pm$ 0.048	0.063 $\pm$ 0.049
	Violations	0.146 $\pm$ 0.353	0.032 $\pm$ 0.176	0.024 $\pm$ 0.152	0.001 $\pm$ 0.022
SS-CMPE <sub>pen</sub>	Gap	0.073 $\pm$ 0.052	0.069 $\pm$ 0.051	0.073 $\pm$ 0.052	0.066 $\pm$ 0.050
	Violations	0.035 $\pm$ 0.185	0.013 $\pm$ 0.113	0.002 $\pm$ 0.039	0.001 $\pm$ 0.022

Table ST12: Average gap and constraint violations over test samples for models applied to the Grids18 Dataset for different  $q$  values. The plot displays the mean values of the average gap and constraint violations, with standard deviations denoted by  $\pm$ .

q		10	30	60	90
ILP Obj		4185.600003	4166.310635	4158.737261	4167.11386
MAE	Gap	0.138 $\pm$ 0.087	0.142 $\pm$ 0.091	0.122 $\pm$ 0.081	0.102 $\pm$ 0.073
	Violations	0.606 $\pm$ 0.489	0.389 $\pm$ 0.488	0.178 $\pm$ 0.383	0.019 $\pm$ 0.138
MSE	Gap	0.056 $\pm$ 0.044	0.060 $\pm$ 0.047	0.069 $\pm$ 0.056	0.064 $\pm$ 0.050
	Violations	0.332 $\pm$ 0.471	0.194 $\pm$ 0.395	0.178 $\pm$ 0.383	0.015 $\pm$ 0.122
MAE+Penalty	Gap	0.143 $\pm$ 0.087	0.154 $\pm$ 0.093	0.145 $\pm$ 0.092	0.114 $\pm$ 0.080
	Violations	0.551 $\pm$ 0.498	0.364 $\pm$ 0.481	0.172 $\pm$ 0.378	0.021 $\pm$ 0.145
SL+Penalty	Gap	0.061 $\pm$ 0.047	0.064 $\pm$ 0.049	0.065 $\pm$ 0.050	0.133 $\pm$ 0.082
	Violations	0.210 $\pm$ 0.407	0.087 $\pm$ 0.282	0.025 $\pm$ 0.158	0.025 $\pm$ 0.158
SSL <sub>pen</sub>	Gap	0.115 $\pm$ 0.065	0.074 $\pm$ 0.055	0.093 $\pm$ 0.060	0.123 $\pm$ 0.065
	Violations	0.060 $\pm$ 0.238	0.182 $\pm$ 0.386	0.013 $\pm$ 0.115	0.005 $\pm$ 0.074
PDL	Gap	0.105 $\pm$ 0.063	0.126 $\pm$ 0.067	0.101 $\pm$ 0.062	0.118 $\pm$ 0.064
	Violations	0.097 $\pm$ 0.295	0.001 $\pm$ 0.032	0.003 $\pm$ 0.050	0.005 $\pm$ 0.071
SS-CMPE	Gap	0.064 $\pm$ 0.047	0.072 $\pm$ 0.053	0.060 $\pm$ 0.046	0.065 $\pm$ 0.049
	Violations	0.200 $\pm$ 0.400	0.029 $\pm$ 0.169	0.000 $\pm$ 0.000	0.001 $\pm$ 0.032
SS-CMPE <sub>pen</sub>	Gap	0.067 $\pm$ 0.051	0.067 $\pm$ 0.051	0.078 $\pm$ 0.055	0.075 $\pm$ 0.053
	Violations	0.017 $\pm$ 0.129	0.019 $\pm$ 0.137	0.002 $\pm$ 0.039	0.000 $\pm$ 0.000

Table ST13: Average gap and constraint violations over test samples for models applied to the AD Dataset for different  $q$  values. The plot displays the mean values of the average gap and constraint violations, with standard deviations denoted by  $\pm$ .

q		10	30	60	90
ILP Obj		2535.424	2526.023	2521.957	2519.917
MAE	Gap	0.288 $\pm$ 0.061	0.276 $\pm$ 0.063	0.283 $\pm$ 0.061	0.270 $\pm$ 0.063
	Violations	0.671 $\pm$ 0.470	0.457 $\pm$ 0.498	0.257 $\pm$ 0.437	0.046 $\pm$ 0.210
MSE	Gap	0.204 $\pm$ 0.061	0.201 $\pm$ 0.063	0.204 $\pm$ 0.061	0.213 $\pm$ 0.059
	Violations	0.336 $\pm$ 0.472	0.063 $\pm$ 0.243	0.069 $\pm$ 0.254	0.013 $\pm$ 0.111
MAE +Penalty	Gap	0.294 $\pm$ 0.062	0.290 $\pm$ 0.066	0.277 $\pm$ 0.061	0.274 $\pm$ 0.061
	Violations	0.600 $\pm$ 0.490	0.468 $\pm$ 0.499	0.271 $\pm$ 0.444	0.039 $\pm$ 0.194
SL+Penalty	Gap	0.216 $\pm$ 0.061	0.213 $\pm$ 0.063	0.220 $\pm$ 0.061	0.229 $\pm$ 0.060
	Violations	0.085 $\pm$ 0.278	0.041 $\pm$ 0.197	0.021 $\pm$ 0.142	0.005 $\pm$ 0.074
SSL <sub>pen</sub>	Gap	0.135 $\pm$ 0.055	0.140 $\pm$ 0.057	0.142 $\pm$ 0.054	0.134 $\pm$ 0.055
	Violations	0.244 $\pm$ 0.430	0.143 $\pm$ 0.350	0.054 $\pm$ 0.226	0.005 $\pm$ 0.074
PDL	Gap	0.148 $\pm$ 0.056	0.152 $\pm$ 0.056	0.146 $\pm$ 0.055	0.139 $\pm$ 0.054
	Violations	0.055 $\pm$ 0.228	0.006 $\pm$ 0.080	0.013 $\pm$ 0.113	0.004 $\pm$ 0.063
SS-CMPE	Gap	0.135 $\pm$ 0.055	0.131 $\pm$ 0.057	0.131 $\pm$ 0.055	0.131 $\pm$ 0.054
	Violations	0.102 $\pm$ 0.302	0.025 $\pm$ 0.155	0.005 $\pm$ 0.071	0.003 $\pm$ 0.055
SS-CMPE <sub>pen</sub>	Gap	0.129 $\pm$ 0.054	0.136 $\pm$ 0.057	0.130 $\pm$ 0.054	0.133 $\pm$ 0.054
	Violations	0.017 $\pm$ 0.127	0.004 $\pm$ 0.063	0.000 $\pm$ 0.000	0.000 $\pm$ 0.000

Table ST14: Average gap and constraint violations over test samples for models applied to the BBC Dataset for different  $q$  values. The plot displays the mean values of the average gap and constraint violations, with standard deviations denoted by  $\pm$ .

$q$		10	30	60	90
ILP Obj		890.289	880.270	875.668	872.118
MAE	Gap	$0.053 \pm 0.037$	$0.046 \pm 0.034$	$0.043 \pm 0.032$	$0.045 \pm 0.034$
	Violations	$0.779 \pm 0.415$	$0.624 \pm 0.485$	$0.414 \pm 0.493$	$0.165 \pm 0.371$
MSE	Gap	$0.036 \pm 0.027$	$0.036 \pm 0.028$	$0.037 \pm 0.028$	$0.037 \pm 0.029$
	Violations	$0.924 \pm 0.265$	$0.854 \pm 0.354$	$0.578 \pm 0.494$	$0.204 \pm 0.403$
MAE+Penalty	Gap	$0.047 \pm 0.036$	$0.044 \pm 0.034$	$0.041 \pm 0.031$	$0.040 \pm 0.031$
	Violations	$0.657 \pm 0.475$	$0.557 \pm 0.497$	$0.384 \pm 0.486$	$0.151 \pm 0.358$
SL+Penalty	Gap	$0.036 \pm 0.028$	$0.036 \pm 0.028$	$0.037 \pm 0.030$	$0.037 \pm 0.029$
	Violations	$0.919 \pm 0.272$	$0.854 \pm 0.354$	$0.578 \pm 0.494$	$0.204 \pm 0.403$
SSL <sub>pen</sub>	Gap	$0.041 \pm 0.032$	$0.042 \pm 0.032$	$0.038 \pm 0.030$	$0.039 \pm 0.030$
	Violations	$0.516 \pm 0.500$	$0.393 \pm 0.488$	$0.408 \pm 0.492$	$0.130 \pm 0.336$
PDL	Gap	$0.043 \pm 0.033$	$0.051 \pm 0.036$	$0.045 \pm 0.034$	$0.044 \pm 0.033$
	Violations	$0.336 \pm 0.472$	$0.160 \pm 0.366$	$0.149 \pm 0.356$	$0.029 \pm 0.169$
SS-CMPE	Gap	$0.038 \pm 0.029$	$0.043 \pm 0.032$	$0.042 \pm 0.033$	$0.040 \pm 0.031$
	Violations	$0.316 \pm 0.495$	$0.239 \pm 0.427$	$0.108 \pm 0.310$	$0.044 \pm 0.206$
SS-CMPE <sub>pen</sub>	Gap	$0.038 \pm 0.030$	$0.043 \pm 0.032$	$0.044 \pm 0.033$	$0.040 \pm 0.031$
	Violations	$0.074 \pm 0.263$	$0.067 \pm 0.250$	$0.056 \pm 0.229$	$0.002 \pm 0.045$

Table ST15: Average gap and constraint violations over test samples for models applied to the 20 Newsgroup Dataset for different  $q$  values. The plot displays the mean values of the average gap and constraint violations, with standard deviations denoted by  $\pm$ .

$q$		10	30	60	90
ILP Obj		928.386	924.439	923.173	921.754
MAE	Gap	$0.044 \pm 0.034$	$0.046 \pm 0.036$	$0.047 \pm 0.035$	$0.048 \pm 0.037$
	Violations	$0.470 \pm 0.499$	$0.176 \pm 0.381$	$0.049 \pm 0.215$	$0.001 \pm 0.022$
MSE	Gap	$0.050 \pm 0.038$	$0.053 \pm 0.039$	$0.051 \pm 0.038$	$0.051 \pm 0.037$
	Violations	$0.639 \pm 0.480$	$0.403 \pm 0.491$	$0.142 \pm 0.349$	$0.008 \pm 0.089$
MAE+Penalty	Gap	$0.044 \pm 0.035$	$0.047 \pm 0.036$	$0.047 \pm 0.036$	$0.047 \pm 0.035$
	Violations	$0.455 \pm 0.498$	$0.181 \pm 0.386$	$0.046 \pm 0.210$	$0.001 \pm 0.022$
SL+Penalty	Gap	$0.046 \pm 0.036$	$0.047 \pm 0.035$	$0.046 \pm 0.036$	$0.044 \pm 0.034$
	Violations	$0.573 \pm 0.495$	$0.384 \pm 0.486$	$0.161 \pm 0.367$	$0.015 \pm 0.122$
SSL <sub>pen</sub>	Gap	$0.045 \pm 0.036$	$0.046 \pm 0.036$	$0.045 \pm 0.035$	$0.046 \pm 0.035$
	Violations	$0.386 \pm 0.487$	$0.139 \pm 0.346$	$0.024 \pm 0.152$	$0.002 \pm 0.039$
PDL	Gap	$0.043 \pm 0.035$	$0.046 \pm 0.036$	$0.046 \pm 0.036$	$0.046 \pm 0.036$
	Violations	$0.278 \pm 0.448$	$0.129 \pm 0.335$	$0.028 \pm 0.165$	$0.001 \pm 0.032$
SS-CMPE	Gap	$0.043 \pm 0.034$	$0.045 \pm 0.035$	$0.044 \pm 0.034$	$0.044 \pm 0.034$
	Violations	$0.317 \pm 0.465$	$0.086 \pm 0.280$	$0.019 \pm 0.137$	$0.001 \pm 0.032$
SS-CMPE <sub>pen</sub>	Gap	$0.044 \pm 0.033$	$0.045 \pm 0.035$	$0.046 \pm 0.035$	$0.045 \pm 0.034$
	Violations	$0.069 \pm 0.254$	$0.054 \pm 0.227$	$0.007 \pm 0.083$	$0.001 \pm 0.022$




Table ST16: Average gap and constraint violations over test samples for models applied to the Webkb Dataset for different  $q$  values. The plot displays the mean values of the average gap and constraint violations, with standard deviations denoted by  $\pm$ .

q		10	30	60	90
ILP Obj		828.463	824.361	825.517	823.917
MAE	Gap	0.057 $\pm$ 0.043	0.057 $\pm$ 0.043	0.058 $\pm$ 0.043	0.062 $\pm$ 0.044
	Violations	0.613 $\pm$ 0.487	0.502 $\pm$ 0.500	0.249 $\pm$ 0.433	0.046 $\pm$ 0.210
MSE	Gap	0.065 $\pm$ 0.046	0.069 $\pm$ 0.046	0.066 $\pm$ 0.045	0.065 $\pm$ 0.045
	Violations	0.695 $\pm$ 0.461	0.473 $\pm$ 0.499	0.210 $\pm$ 0.407	0.018 $\pm$ 0.133
MAE+Penalty	Gap	0.058 $\pm$ 0.042	0.058 $\pm$ 0.042	0.059 $\pm$ 0.043	0.061 $\pm$ 0.044
	Violations	0.471 $\pm$ 0.499	0.395 $\pm$ 0.489	0.211 $\pm$ 0.408	0.041 $\pm$ 0.197
SL+Penalty	Gap	0.054 $\pm$ 0.042	0.054 $\pm$ 0.041	0.054 $\pm$ 0.040	0.054 $\pm$ 0.040
	Violations	0.584 $\pm$ 0.493	0.378 $\pm$ 0.485	0.174 $\pm$ 0.380	0.018 $\pm$ 0.131
SSL <sub>pen</sub>	Gap	0.058 $\pm$ 0.043	0.057 $\pm$ 0.041	0.057 $\pm$ 0.042	0.057 $\pm$ 0.042
	Violations	0.360 $\pm$ 0.480	0.217 $\pm$ 0.413	0.082 $\pm$ 0.274	0.015 $\pm$ 0.120
PDL	Gap	0.063 $\pm$ 0.045	0.060 $\pm$ 0.044	0.058 $\pm$ 0.042	0.057 $\pm$ 0.042
	Violations	0.149 $\pm$ 0.357	0.096 $\pm$ 0.295	0.056 $\pm$ 0.229	0.013 $\pm$ 0.115
SS-CMPE	Gap	0.059 $\pm$ 0.043	0.058 $\pm$ 0.043	0.056 $\pm$ 0.042	0.056 $\pm$ 0.042
	Violations	0.169 $\pm$ 0.374	0.050 $\pm$ 0.218	0.026 $\pm$ 0.159	0.004 $\pm$ 0.063
SS-CMPE <sub>pen</sub>	Gap	0.062 $\pm$ 0.045	0.061 $\pm$ 0.044	0.057 $\pm$ 0.042	0.053 $\pm$ 0.040
	Violations	0.074 $\pm$ 0.263	0.029 $\pm$ 0.169	0.001 $\pm$ 0.032	0.001 $\pm$ 0.022


Table ST17: Average gap and constraint violations over test samples for models applied to the DNA Dataset for different  $q$  values. The plot displays the mean values of the average gap and constraint violations, with standard deviations denoted by  $\pm$ .

q		10	30	60	90
ILP Obj		222.848	221.635	221.114	220.625
MAE	Gap	0.138 $\pm$ 0.109	0.142 $\pm$ 0.109	0.136 $\pm$ 0.109	0.141 $\pm$ 0.111
	Violations	0.444 $\pm$ 0.497	0.448 $\pm$ 0.497	0.286 $\pm$ 0.452	0.114 $\pm$ 0.317
MSE	Gap	0.138 $\pm$ 0.112	0.140 $\pm$ 0.112	0.139 $\pm$ 0.111	0.139 $\pm$ 0.110
	Violations	0.506 $\pm$ 0.500	0.565 $\pm$ 0.496	0.322 $\pm$ 0.467	0.113 $\pm$ 0.317
MAE+Penalty	Gap	0.140 $\pm$ 0.111	0.136 $\pm$ 0.106	0.143 $\pm$ 0.111	0.137 $\pm$ 0.113
	Violations	0.444 $\pm$ 0.497	0.448 $\pm$ 0.497	0.286 $\pm$ 0.452	0.114 $\pm$ 0.317
SL+Penalty	Gap	0.135 $\pm$ 0.109	0.140 $\pm$ 0.112	0.141 $\pm$ 0.111	0.143 $\pm$ 0.115
	Violations	0.434 $\pm$ 0.496	0.494 $\pm$ 0.500	0.281 $\pm$ 0.450	0.089 $\pm$ 0.285
SSL <sub>pen</sub>	Gap	0.140 $\pm$ 0.115	0.141 $\pm$ 0.111	0.146 $\pm$ 0.116	0.143 $\pm$ 0.118
	Violations	0.048 $\pm$ 0.214	0.062 $\pm$ 0.241	0.014 $\pm$ 0.118	0.004 $\pm$ 0.067
PDL	Gap	0.140 $\pm$ 0.113	0.141 $\pm$ 0.113	0.139 $\pm$ 0.112	0.144 $\pm$ 0.120
	Violations	0.287 $\pm$ 0.452	0.129 $\pm$ 0.335	0.003 $\pm$ 0.055	0.006 $\pm$ 0.077
SS-CMPE	Gap	0.138 $\pm$ 0.113	0.138 $\pm$ 0.108	0.137 $\pm$ 0.106	0.139 $\pm$ 0.109
	Violations	0.046 $\pm$ 0.210	0.017 $\pm$ 0.129	0.012 $\pm$ 0.109	0.008 $\pm$ 0.089
SS-CMPE <sub>pen</sub>	Gap	0.139 $\pm$ 0.116	0.139 $\pm$ 0.113	0.140 $\pm$ 0.112	0.139 $\pm$ 0.113
	Violations	0.013 $\pm$ 0.115	0.002 $\pm$ 0.045	0.001 $\pm$ 0.022	0.001 $\pm$ 0.022


**D.B Optimality Gap And Violations in Self-Supervised Methods for Different q Values**



(a) Optimality Gap (avg %) and Average Violations for Grids UAI networks



(b) Opt. Gap (avg %) and Avg. Violations for Segmentation UAI networks



(c) Optimality Gap (avg %) and Average Violations for Tractable Models

Figure SF5: Visualization of Optimality Gap (average %) and Average Violations for Self-Supervised Methods across different q values. Points closer to the origin indicate better performance.

In the scatter plots depicted in Figure SF5, three distinct evaluations of the optimality gap against the average violations for various self-supervised methods across different q values are visualized. Points positioned closer to the origin indicate better performance, with reduced optimality gaps and fewer violations. In Figure SF5(a), focused on Grids UAI networks, the  $SS-CMPE_{pen}$  method generally occupies a position near the origin, indicating its commendable performance in this setting. The  $SS-CMPE$  method exhibits a comparable performance to the  $SS-CMPE_{pen}$  method, with occasional high levels of violations observed in two instances.

In Figure SF5(b), showcasing the Segmentation UAI networks, the  $SS-CMPE$  and  $SS-CMPE_{pen}$  methods again demonstrate superiority, particularly evident by their prevalence near the origin. Finally, in Figure SF5(c) related to tractable models, the  $SS-CMPE_{pen}$  method often achieves optimal placement close to the origin, reflecting a balanced performance. These evaluations provide critical insights into the effectiveness and robustness of the proposed self-supervised methods across different problems.