

Adversarial Constraint Learning for Structured Prediction

Hongyu Ren, Russell Stewart, Jiaming Song, Volodymyr Kuleshov, Stefano Ermon

Department of Computer Science, Stanford University
 {hyren, stewartr, tsong, kuleshov, ermon}@cs.stanford.edu

Abstract

Constraint-based learning reduces the burden of collecting labels by having users specify general properties of structured outputs, such as constraints imposed by physical laws. We propose a novel framework for simultaneously learning these constraints and using them for supervision, bypassing the difficulty of using domain expertise to manually specify constraints. Learning requires a black-box simulator of structured outputs, which generates valid labels, but need not model their corresponding inputs or the input-label relationship. At training time, we constrain the model to produce outputs that cannot be distinguished from simulated labels by adversarial training. Providing our framework with a small number of labeled inputs gives rise to a new semi-supervised structured prediction model; we evaluate this model on multiple tasks — tracking, pose estimation and time series prediction — and find that it achieves high accuracy with only a small number of labeled inputs. In some cases, no labels are required at all.

1 Introduction

Large labeled datasets are a key component for building state-of-the-art systems in many applications of machine learning, including image recognition, machine translation, and speech recognition. Collecting such datasets can be expensive, which has driven significant research interest in unsupervised, semi-supervised, and weakly supervised learning approaches [Radford *et al.*, 2015; Kingma *et al.*, 2014; Papandreou *et al.*, 2015; Ratner *et al.*, 2016].

Constraint-based learning is a recently proposed form of weak supervision which aims to reduce the need for labeled inputs by having users supervise algorithms through general properties that hold over the label space [Shcherbatyi and Andres, 2016; Stewart and Ermon, 2017]. Examples of such properties include logical rules [Richardson and Domingos, 2006; Chang *et al.*, 2007; Choi *et al.*, 2015; Xu *et al.*, 2017] or physical laws [Stewart and Ermon, 2017; Ermon *et al.*, 2015].

Unlike labels — which only apply to their corresponding inputs — properties used in a constraint-based learning ap-

proach are specified once for the entire dataset, providing an opportunity for more cost-efficient supervision. Algorithms supervised with explicit constraints have shown promising results in object detection [Stewart and Ermon, 2017], preference learning [Choi *et al.*, 2015], materials science [Ermon *et al.*, 2012], and semantic segmentation [Pathak *et al.*, 2015].

However, describing the high level invariants of a dataset may also require a non-trivial amount of effort. First, designing constraints requires strong domain expertise. Second, in the case of high dimensional labels, it is difficult to encode the constraints using simple formulas. For example, suppose we want to constrain a pedestrian joint detector to produce skeletons that “look like a walking person”; in this case, it is difficult to capture invariants over human poses with simple logical or algebraic formulas that an annotator could specify. Third, constraints may change over time and across tasks; designing new constraints for new tasks may not scale in many practical applications.

In this paper, we propose an *implicit* approach to constraint learning, in which invariants are automatically learned from a small set of representative label samples (see Figure 1). These samples do not need to be tied to corresponding inputs (as in supervised learning) and may come from a black-box simulator that abstracts away physics-based formulas or produces examples of labels collected by humans. Such simulators include physics engines, humanoid simulators from robotics, or driving simulators [Li *et al.*, 2017].

Inspired by recent advances in implicit (likelihood-free) generative modeling, we capture the distribution of outputs using an approach based on adversarial learning [Goodfellow *et al.*, 2014]. Specifically, we train two distinct learners: a primary model for the task at hand and an auxiliary classification algorithm called discriminator. During training, we constrain the main model such that its outputs cannot be distinguished by the discriminator from representative (true) label samples, thus forcing it to capture the structure of the label space. This approach forms a novel adversarial framework for performing weak supervision with learned constraints, which we call adversarial constraint learning.

Although constraint learning does not require input-label pairs, providing such pairs can improve performance and turns our problem into an instance of semi-supervised learning. In this setting, our approach combines supervised learning on a small labeled dataset with constraint learning on a




Figure 1: Constraint learning allows us to learn a conditional probabilistic model $p_\theta(\mathbf{y}|\mathbf{x})$ (parameterized by θ) without direct labels by specifying properties h that holds over the output space. In prior work (left), h is defined as a formula describing known invariants. In this paper (right), we propose to instead learn h through an auxiliary classifier D_ϕ (parameterized by ϕ) that discriminates \mathbf{y} (provided by $p_\theta(\mathbf{y}|\mathbf{x})$) from $\hat{\mathbf{y}}$ (provided by an additional source unrelated to \mathbf{x} , such as a simulator).

large unlabeled set, where constraint learning enforces that the structure of predictions on unlabeled data matches the structure observed in the labeled data. Experimental results demonstrate that this method performs better than state-of-the-art semi-supervised learning methods on a variety of structured prediction problems.

2 Background

In this section, we introduce structured prediction and constraint-based learning. The next section will expand upon these subjects to introduce the proposed adversarial constraint learning framework.

2.1 Structured Prediction

Our work focuses on structured prediction, a form of supervised learning, in which the outputs $\mathbf{y} \in \mathcal{Y}$ can be a complex object such as a vector, a tree, or a graph [Koller and Friedman, 2009]. We capture the distribution of \mathbf{y} using a conditional probabilistic model $p_\theta(\mathbf{y}|\mathbf{x})$ parameterized by $\theta \in \Theta$. A model $p_\theta(\mathbf{y}|\mathbf{x})$ maps each input $\mathbf{x} \in \mathcal{X}$ to the corresponding output distribution $p_\theta(\mathbf{y}) \in \mathcal{P}(\mathcal{Y})$, where $\mathcal{P}(\mathcal{Y})$ denotes all the probability distributions over \mathcal{Y} . For example, we may take $p_\theta(\mathbf{y}|\mathbf{x})$ to be a Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}_\theta(\mathbf{x}), \boldsymbol{\Sigma}_\theta(\mathbf{x}))$ with mean $\boldsymbol{\mu}_\theta(\mathbf{x})$ and variance $\boldsymbol{\Sigma}_\theta(\mathbf{x})$.

A standard approach to learning $p_\theta(\mathbf{y}|\mathbf{x})$ (or p_θ as an abbreviation) is to solve an optimization problem of the form

$$\theta^* = \arg \min_{\theta \in \Theta} \sum_{i=1}^n \ell(p_\theta(\mathbf{y}|\mathbf{x}_i), \mathbf{y}_i) + R(p_\theta) \quad (1)$$

over a labeled dataset $\mathbf{D} = \{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_n, \mathbf{y}_n)\}$. A typical supervised learning objective is comprised of a loss function $\ell : \mathcal{P}(\mathcal{Y}) \times \mathcal{Y} \rightarrow \mathbb{R}$ and a regularization term $R : \mathcal{P}(\mathcal{Y}) \rightarrow \mathbb{R}$ that encourages non-degenerate solutions or solutions that incorporate prior knowledge [Stewart and Ermon, 2017].

2.2 Constraint-Based Learning

Collecting a large labeled dataset for supervised learning can often be tedious. Constraint-based learning is a form of weak supervision which instead asks users to specify high-level constraints over the output space, such as logical rules or physical laws [Shcherbatyi and Andres, 2016;

Stewart and Ermon, 2017; Richardson and Domingos, 2006; Xu *et al.*, 2017]. For example, in an object tracking task where \mathcal{Y} corresponds to the space of joint positions over time, we expect correct outputs to be consistent with the laws of physical mechanics.

Let $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ be an unlabeled dataset of inputs. Formally, constraints can be specified via a function $h : \mathcal{P}(\mathcal{Y}) \rightarrow \mathbb{R}$, which penalizes conditional probabilistic models $p_\theta(\mathbf{y}|\mathbf{x})$ that are inconsistent with known high-level structure of the label space. Learning from constraints proceeds by optimizing the following objective:

$$\hat{\theta}^* = \arg \min_{\theta \in \Theta} \sum_{i=1}^m h(p_\theta(\mathbf{y}|\mathbf{x}_i)) + R(p_\theta) \quad (2)$$

over \mathbf{X} . By solving this optimization problem, we look for a probabilistic model parameterized by $\hat{\theta}^*$ that satisfies known constraints when applied to the unlabeled dataset \mathbf{X} (through the h term), and is likely a priori (through the $R(p_\theta)$ term). Note that although the constraint h is data-dependent, *it does not require explicit labels*. For example, in object tracking we could ask that when making predictions on \mathbf{X} , joint positions over time are consistent with known kinematic equations, with h measuring how the output distribution from p_θ deviates from those equations. The regularization term can be used to avoid overly complex and/or degenerate solutions, and may include $L1$, $L2$, or entropy regularization terms. Stewart and Ermon [Stewart and Ermon, 2017] have shown that a model learned with the objective described in Eq. 2 can learn to track objects.

3 Adversarial Constraint Learning

The process of manually specifying high level constraints, h , can be time-consuming and may require significant domain expertise. Such is the case in pose estimation, where it is difficult to describe high dimensional rules for joints movements precisely; but the large availability of unpaired videos and motion capture data makes constraint learning attractive in spite of the difficulty of providing high dimensional constraints.

In the sciences, discovering general invariants is often a data-driven approach; for example, the laws of physics are often discovered by validating hypotheses with experimental results. Motivated by this, we propose in this section a novel framework for learning constraints from data.

3.1 Learning Constraints from Data

Suppose we have a dataset of inputs $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_m\}$, a dataset of labels $\mathbf{Y} = \{\mathbf{y}_1, \dots, \mathbf{y}_k\}$, and a set $\mathbf{D} = \{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_n, \mathbf{y}_n)\}$ that describes correspondence between some elements of \mathbf{X} and \mathbf{Y} . We denote the empirical distributions of \mathbf{X} , \mathbf{Y} and \mathbf{D} as $p(\mathbf{x})$, $p(\mathbf{y})$ and $p(\mathbf{x}, \mathbf{y})$ respectively. Note that \mathbf{Y} can come from either a simulator (such as one based on physical rules), or from some other source of data (such as motion captures of people for which we have no corresponding videos).

Let us first consider the setting where $\mathbf{D} = \emptyset$; i.e. there are inputs and labels but no correspondence between them.

In spite of the lack of correspondences, we will see that constraints h can be learned from the prior knowledge that the same underlying distribution generates both the empirical labels \mathbf{Y} and the structured predictions obtained from applying our model to \mathbf{X} . These learned constraints can then be used for supervision. Let structured predictions be given by the following *implicit* sampling procedure:

$$\mathbf{x} \sim p(\mathbf{x}) \quad , \quad \mathbf{y} \sim p_\theta(\mathbf{y}|\mathbf{x}) \quad (3)$$

where $p_\theta(\mathbf{y}|\mathbf{x})$ is a (parameterized) conditional distribution of outputs given inputs. Discarding \mathbf{x} , the above procedure corresponds to sampling from the marginal distribution over \mathcal{Y} , $p_\theta(\mathbf{y}) = \int p_\theta(\mathbf{y}|\mathbf{x})p(\mathbf{x})d\mathbf{x}$.

Labels drawn from $p(\mathbf{y})$ should have high likelihood values in $p_\theta(\mathbf{y})$, but optimizing this objective directly is computationally infeasible; evaluating the marginal likelihood $p_\theta(\mathbf{y})$ exactly is expensive due to the integration over $p(\mathbf{x})$. Instead, we formulate the task of learning a constraint loss h from $p(\mathbf{y})$ through a likelihood-free approach using the framework of generative adversarial learning [Goodfellow *et al.*, 2014], which only requires samples from $p_\theta(\mathbf{y})$ and $p(\mathbf{y})$.

We introduce an auxiliary classifier D_ϕ (parametrized by ϕ) called discriminator which scores outputs in the label space \mathcal{Y} . It is trained to assign high scores to representative output labels from $p(\mathbf{y})$, while assigning low scores to samples from $p_\theta(\mathbf{y})$. It learns to effectively extract latent constraints that hold over the output space and that are implicitly encoded in the samples from $p(\mathbf{y})$. The goal of $p_\theta(\mathbf{y}|\mathbf{x})$ is to produce outputs result in higher scores in the discriminator, satisfying the constraints imposed by D_ϕ in the process.

For practical reasons, we consider $p_\theta(\mathbf{y}|\mathbf{x})$ to be a Dirac-delta distribution $\delta(\mathbf{y} - f_\theta(\mathbf{x}))$, and thus we refer to the conditional probabilistic model as the mapping $f_\theta(\mathbf{x}) : \mathcal{X} \rightarrow \mathcal{Y}$ in the experiment section for simplicity. We train D_ϕ and $p_\theta(\mathbf{y}|\mathbf{x})$ for the following objective [Arjovsky *et al.*, 2017]

$$\min_\theta \max_\phi \mathcal{L}^A \quad (4)$$

$$\mathcal{L}^A = \mathbb{E}_{\mathbf{y} \sim p(\mathbf{y})} [D_\phi(\mathbf{y})] - \mathbb{E}_{\mathbf{y} \sim p_\theta(\mathbf{y}|\mathbf{x}), \mathbf{x} \sim p(\mathbf{x})} [D_\phi(\mathbf{y})]$$

Assuming infinite capacity, Theorem 1 of [Goodfellow *et al.*, 2014] shows that at the optimal solution of Eq. 4, D_ϕ cannot distinguish between the given set of labels and those predicted by the model p_θ , suggesting that the latter satisfy the set of constraints defined by D_ϕ . Unlike in constraint-based learning where a (possibly incomplete) set of constraints is manually specified, convergence in the adversarial setting implies that the label and output distributions match on all possible discriminator projections. Figure 2(a) shows an overview of the adversarial constraint learning framework in the context of trajectory estimation.

3.2 Constraint Learning via Matching Distributions

Generative Adversarial Networks (GANs) are a prominent example of implicit probabilistic models [Mohamed and Lakshminarayanan, 2016] which are defined through a stochastic sampling procedure instead of an explicitly defined likelihood function. One advantage of implicit generative models is that

they can be trained with methods that do not require likelihood evaluations.

Hence, our approach to learning constraints for structured prediction can also be interpreted as learning an implicit generative model $p_\theta(\mathbf{y})$ that matches the empirical label distribution $p(\mathbf{y})$. Specifically, our adversarial constraint learning approach optimizes over an approximation to the optimal transport from $p_\theta(\mathbf{y})$ to $p(\mathbf{y})$ [Arjovsky *et al.*, 2017]; thus our constraint can be implicitly defined as “ θ minimizes the optimal transport from $p_\theta(\mathbf{y})$ to $p(\mathbf{y})$ ”.

3.3 Semi-Supervised Structured Prediction

Models	(\mathbf{x}, \mathbf{y})	(\mathbf{x}, \cdot)	(\cdot, \mathbf{y})
SL	✓		
SSL	✓	✓	
ACL		✓	✓
SSACL	✓	✓	✓

Table 1: Settings in different learning paradigms. Supervised Learning (SL) requires a dataset with paired (\mathbf{x}, \mathbf{y}). Semi-Supervised Learning (SSL) utilizes additional unlabeled inputs (\mathbf{x}, \cdot). Adversarial Constraint Learning (ACL) requires inputs (\mathbf{x}, \cdot) and labels (\cdot, \mathbf{y}) but without correspondences between them. Semi-Supervised Adversarial Constraint Learning (SSACL) extends ACL by also considering labeled pairs (\mathbf{x}, \mathbf{y}).

Although our framework does not require datasets containing input-label pairs $\mathbf{D} \supseteq \emptyset$, providing it with such data gives rise to a new semi-supervised structured prediction method.

When given a set of labeled examples, we may extend our constraint learning objective (over both labeled and unlabeled data) with a standard classification loss term (over labeled data):

$$\mathcal{L}^{SS} = \mathcal{L}^A + \alpha \mathbb{E}_{\mathbf{x}_i, \mathbf{y}_i \sim p(\mathbf{x}, \mathbf{y})} [\ell(p_\theta(\mathbf{y}|\mathbf{x}_i), \mathbf{y}_i)] \quad (5)$$

where \mathcal{L}^A is the adversarial constraint learning objective defined in Eq. 4, and α is a hyperparameter that balances between fitting to the general (implicit) label distribution (first term) and fitting to the explicit labeled dataset (second term).

Our semi-supervised constraint learning framework is different from traditional semi-supervised learning approaches, as listed in Table 1. In particular, traditional semi-supervised learning methods assume there is a large source of *inputs* and tend to impose regularization over \mathcal{X} , such as through latent variables [Kingma *et al.*, 2014], through outputs [Miyato *et al.*, 2017], or through another network [Salimans *et al.*, 2016]. We consider the case where there exists a source, e.g., a simulator that can provide abundant *samples from the label space that are not matched to particular inputs*, and impose regularization over \mathcal{Y} by exploiting a discriminator that provides an implicit constraint over the predicted \mathbf{y} values. Therefore, we can also utilize sample labels that are not associated with particular inputs, instead of merely restricting to standard labeled (\mathbf{x}, \mathbf{y}) pairs. Moreover, our method can be easily combined with existing semi-supervised learning approaches [Kingma *et al.*, 2014; Li *et al.*, 2016; Miyato *et al.*, 2017] to further boost performance.

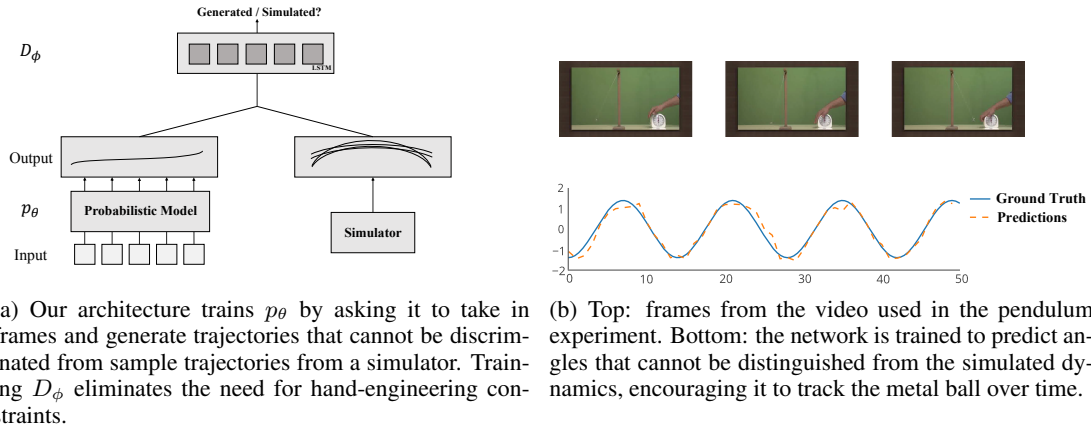


Figure 2: Architecture and results of the pendulum tracking experiment.

4 Experimental Results

We evaluate the proposed framework on three structured prediction problems. First, we aim to track the angle of a pendulum in a video without labels using supervision provided by a physics-based simulator. Next, we extend the output space to higher dimensions and perform human pose estimation in a semi-supervised setting. Lastly, we evaluate our approach on multivariate time series prediction, where the goal is to predict future temperature and humidity.

A label simulator is provided for each experiment in place of hand-written constraints. Although explicit constraints for the pendulum case can be written down analytically, we demonstrate that our adversarial framework is capable of learning the constraint from data. In the other two experiments, we consider structured prediction settings where the outputs are high dimensional; in these settings, the correct constraints are very complex and hand-engineering them would be difficult. Instead, our model learns these constraints from a small number of samples provided by the simulator.

4.1 Pendulum Tracking

For this task, we aim to predict the angle of the pendulum from images in a YouTube video ¹, i.e., learn a regression mapping $r_\theta : \mathbb{R}^{h \times w \times 3} \rightarrow \mathbb{R}$, where h and w are the height and width of the input image. Since the outputs of r_θ over consecutive frames are constrained by temporal structure (a sine wave in this case), we concatenate consecutive outputs of r_θ and form a high dimensional trajectory, thus defining $f_\theta([\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]) = [r_\theta(\mathbf{x}_1), r_\theta(\mathbf{x}_2), \dots, r_\theta(\mathbf{x}_n)]$. Critically, r_θ must make a separate prediction for each image, preventing f_θ from simply memorizing the output structure. Unlike previous methods [Stewart and Ermon, 2017], no explicit formulas are provided for supervision, and the (implicit) constraints are learned through the discriminator D_ϕ using samples provided by the physics simulator.

Training Details The video contains a total of 170 images, and we hold out 34 images for evaluation. We manually ob-

serve that the pendulum completes one full oscillation approximately every 12 frames. Based on this observation, we write a simulator of these dynamics with a simple harmonic oscillator having a fixed amplitude and random sample period of 10 to 14 frames. D_ϕ is trained to distinguish between the output of r_θ across $n = 5$ continuous images and a random trajectory sampled from the simulator. We implement r_θ as a 5 layer convolutional neural network with ReLU nonlinearities, and D_ϕ as a 5-cell LSTM. We use $\alpha = 10$ in Eq. 5, and the same training procedure and hyperparameters as [Gulrajani *et al.*, 2017] across our experiments.

Evaluation We manually label the horizontal position of the ball of the pendulum for each frame *in the test set*, and measure the correlation between the predicted positions and the ground truth labels. Since the same r_θ is applied to each input frame independently, f_θ cannot just memorize valid (i.e. simple harmonic) trajectory sequences and produce them while ignoring inputs. The model must learn to track the pendulum in order to fool the discriminator and subsequently achieve a high correlation on the test set.

Our adversarial constraint learning approach achieves a correlation of 96.3%, whereas training with hand-crafted constraints achieves a marginally higher correlation of 96.6%. Both approaches are trained without labels. Example predictions on the test data are shown in Figure 2. This real-world experiment demonstrates the effectiveness of constraint-based learning in the absence of labels, and suggests that using learned constraints from data is almost as effective as using ideal hand-crafted constraints.

4.2 Pose Estimation

In this experiment, we evaluate the proposed model on pose estimation, which has a significantly larger output space. We aim to learn a regression network $r_\theta : \mathbb{R}^{h \times w \times 3} \rightarrow \mathbb{R}^{k \times 2}$, where k denotes the number of joints to detect, and each joint has 2 coordinates. As in the pendulum tracking experiment, r_θ is mapped across several frames to produce a trajectory $f_\theta([\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]) = [r_\theta(\mathbf{x}_1), r_\theta(\mathbf{x}_2), \dots, r_\theta(\mathbf{x}_n)]$ that is indistinguishable from samples provided by the simulator.

¹https://www.youtube.com/watch?v=02w9ISii_Hs



Figure 3: Pose estimation using the proposed semi-supervised adversarial constraint learning approach. r_θ takes in single image and outputs the 2-D location of 6 joints (in green). Lines (in red) are added automatically. The images show results across 4 test groups (horizontal strips) when only 3 out of 28 training groups were directly labeled.

We evaluate the model with videos and joint trajectories from the CMU multi-modal action database (MAD) [Huang *et al.*, 2014]. MAD contains videos of 20 subjects performing a sequence of actions in each video. We extract frames from subjects performing the “Jump and Side-Kick” action and train r_θ to detect the location of the left/right hip/knee/foot ($k = 6$) in each frame. The processed dataset contains 35 groups (549 valid frames in total).

Training Details We divide the 35 groups of motion data into training and testing sets of 28 groups and 7 groups, respectively, where direct labels will be provided for a subset of the 28 training groups. Each group contains 14 to 17 frames, and we train on randomly selected contiguous intervals of length $n = 5$. Using the metric of PCK@0.1 [Yang and Ramanan, 2013] for evaluation, a prediction is considered correct if it lies within $\beta \max(h, w)$ pixels from the true location, where h and w denote the height and width of the subject’s body. We evaluate with $\beta = 0.1$.

We first design a simulator of valid labels (joint positions) based on known kinematics of skeletons. Specifically, the anatomical shape of the subject’s legs approximately forms an expanding isosceles trapezoid when they jump and side kick. We simulate a large range of trapezoidal motions capturing these trajectories, which requires much less effort than hand engineering precise mathematical formulas to express explicit constraints. r_θ takes a single image as input and produces a 12 dimensional vector, representing the location of $k = 6$ joints. Critically, as in the pendulum experiment, r_θ is applied to each frame independently, and has *no knowledge* of the neighboring frames. The outputs of r_θ are concatenated and passed to the discriminator D_ϕ for training.

Evaluation We construct 50 random train/test splits of the dataset and report the averaged PCK@0.1 scores for evaluation. The results are summarized in Table 2 and Figure 3, where we compare three forms of learning when labels are only available for $i \leq 28$ of the training groups:

- “L(i)”: vanilla supervised learning on labeled groups

- “L(i)+VAT”: a baseline form of semi-supervised learning leveraging virtual adversarial training on unlabeled groups (VAT, [Miyato *et al.*, 2017])
- “L(i)+ADV”: semi-supervised learning with adversarial constraint learning (Eq. 5)

When no labels are provided (“L(0)+ADV”; i.e., optimizing just \mathcal{L}^A), r_θ is able to find the correct “shape” of the joints for each frame, but the predictions are biased. Since the subjects are not strictly acting in the center of the image, a constant minor shift ($\Delta x, \Delta y$) for all predicted joint locations still meets the requirements imposed by D_ϕ , which encodes the structure of the label space. This problem is addressed when providing even a very small (“i=1”) number of labeled training groups and using the semi-supervised objective \mathcal{L}^{SS} . Availability of labels fixes the constant bias and we note that using adversarial training produces a massive (25-30%) boost over both the supervised and “VAT” baselines when only 1 group of labeled data is available.

With only 3 groups of labeled data (“L(3)+ADV”), adversarial constraint learning achieves a comparable performance to standard supervised learning with 7 groups of labeled inputs (“L(7)”). Adversarial constraint learning “L(i)+ADV” consistently outperforms the virtual adversarial training “L(i)+VAT” baseline for different values of i . When further combined with VAT regularization in the \mathcal{L}^{SS} objective, our method achieves slightly better performance.

The strong performance of our model over baselines on the pose estimation task with a few or no labels demonstrates that constraint learning can work well over high-dimensional outputs when using our proposed adversarial framework. Designing precise constraints in high-dimensional spaces is often tedious, error-prone, and restricted to one particular domain. Our method avoids these downsides by learning these constraints implicitly through data generated from a simulator, even though the simulator can be a noisy (or even slightly biased) description of the true label distribution.

4.3 Time Series Prediction

Lastly, we validate our model on another structured prediction problem: multi-step multivariate time series prediction.

PCK@0.1(%)	Left Hip	Left Knee	Left Foot	Right Hip	Right Knee	Right Foot
L(0)+ADV	0.6813	0.7326	0.6047	0.6669	0.6729	0.5834
L(1)	0.5453	0.5728	0.5464	0.5360	0.4983	0.4362
L(1)+VAT	0.5795	0.6086	0.5797	0.5608	0.5016	0.4571
L(1)+ADV	0.8529	0.8510	0.8151	0.8482	0.8531	0.7394
L(3)	0.8275	0.7937	0.6716	0.8092	0.7529	0.6196
L(3)+VAT	0.8334	0.7866	0.7082	0.8281	0.7760	0.6420
L(3)+ADV	0.8760	0.9097	0.8328	0.8601	0.8746	0.7549
L(5)	0.8603	0.8483	0.7493	0.8309	0.8267	0.6626
L(5)+VAT	0.8750	0.8764	0.7411	0.8471	0.8398	0.6596
L(5)+ADV	0.9022	0.9160	0.8581	0.9192	0.8706	0.7894
L(7)	0.9088	0.8639	0.8217	0.8887	0.8387	0.7338
L(7)+VAT	0.9201	0.8665	0.8436	0.9074	0.8312	0.7526
L(7)+ADV	0.9469	0.9347	0.8418	0.9367	0.8988	0.8161
L(ALL)	0.9622	0.9633	0.9290	0.9464	0.9133	0.8936
L(ALL)+ADV	0.9758	0.9882	0.9627	0.9708	0.9522	0.8740

Table 2: PCK@0.1 results on MAD. “L(i)” indicates supervised learning (SL) where labeled data is provided for only i out of 28 groups in the training set. “L(i)+VAT” indicates SL with additional optimization over unlabeled groups using virtual adversarial training [Miyato *et al.*, 2017] (SSL). “L(i)+ADV” indicates SL with additional optimization over the entire training set with the ACL objective. Our approach outperforms the baselines, especially when very few labels are available.




Figure 4: Mean absolute error of the predictions on temperature (top) and humidity (bottom) during training (left) and testing (right). Our method SSACL (trained on \mathcal{L}^{SS} objective) consistently outperforms SL (supervised learning) on the test set with different numbers of “complete groups” used in training.

In this task, we aim to learn a mapping $f_\theta : \mathbb{R}^{t \times m} \rightarrow \mathbb{R}^{k \times m}$. Given t consecutive values of a series, $(\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_t)$, we aim to predict the following k values, $(\mathbf{y}_{t+1}, \dots, \mathbf{y}_{t+k})$, where each \mathbf{y} has m variables. In this task, D_ϕ learns the constraint that both holds across variables and time by distinguishing the output of f_θ from real label samples.

Training Details We conduct experiments on the SML2010 Dataset [Zamora-Martínez *et al.*, 2014], which contains humidity and temperature data of indoor and outdoor environments over 40 days at 15 minute intervals.

We hold out 8 consecutive days for testing and leave the rest for training. From the train and test set, we sample 480 and 120 groups of time series data respectively, each having length of 28 hours, and smooth each group into 7 data points at 4-hour intervals. Each group uses the first $t = 5$ data points as input, and leaves the final $k = 2$ values as targets for prediction, with each data point having $m = 4$ variables representing the indoor/outdoor temperature/humidity. We measure the mean absolute error (MAE) on the test set.

We further explore the setting when not all groups in the training set are “complete”; for example, in some groups we

may only have temperature information. This is reasonable in real-world scenarios where sensors fail to work properly from time to time. Hence, we use “complete groups” to denote groups with full information, and “incomplete groups” to denote groups with only temperature information. Without humidity records, we could not perform supervised learning on these “incomplete groups”, since the input requires all $m = 4$ values. Under the context of adversarial constraint learning, however, such data can facilitate learning constraints over the temperature series. In this task, the simulator is designed to produce humidity samples from only the “complete groups” and temperature samples from all the groups in the training set. Both f_θ and D_ϕ are 4-layer MLPs with 64 neurons per layer.

Evaluation We display quantitative results in Figure 4. The supervised baseline (trained only with labels) achieves lower training error but results in higher test error. Our model effectively avoids overfitting to the small portion of labeled data and consistently outperforms the baseline, achieving a MAE of 1.933 and 3.042 on the predictions of temperature and humidity when all groups are “completely” labeled.

5 Conclusion

We have proposed adversarial constraint learning, a new framework for structured prediction that replaces hand-crafted, domain specific constraints with implicit, domain agnostic ones learned through adversarial methods. Experimental results on multiple structured prediction tasks demonstrate that adversarial constraint learning works across many real-world applications with limited data, and fits naturally into semi-supervised structured prediction problems. Our success with matching distributions of labeled and unlabeled model outputs motivates future work exploring analogous opportunities for adversarially matching labeled and unlabeled distributions of learned intermediate representations.

Acknowledgments

This work was supported by a grant from the SAIL-Toyota Center for AI Research, TRI, Siemens, ONR, NSF grants #1651565, #1522054, #1733686.

References

- [Arjovsky *et al.*, 2017] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.
- [Chang *et al.*, 2007] Ming-Wei Chang, Lev Ratinov, and Dan Roth. Guiding semi-supervision with constraint-driven learning. In *ACL*, pages 280–287, 2007.
- [Choi *et al.*, 2015] Arthur Choi, Guy Van den Broeck, and Adnan Darwiche. Tractable learning for structured probability spaces: A case study in learning preference distributions. In *Proceedings of 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.
- [Ermon *et al.*, 2012] Stefano Ermon, Ronan Le Bras, Carla P Gomes, Bart Selman, and R Bruce Van Dover. Smt-aided combinatorial materials discovery. In *International Conference on Theory and Applications of Satisfiability Testing*, pages 172–185. Springer, 2012.
- [Ermon *et al.*, 2015] Stefano Ermon, Ronan Le Bras, Santosh K Suram, John M Gregoire, Carla P Gomes, Bart Selman, and Robert Bruce van Dover. Pattern decomposition with complex combinatorial constraints: Application to materials discovery. In *AAAI*, pages 636–643, 2015.
- [Goodfellow *et al.*, 2014] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [Gulrajani *et al.*, 2017] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans. *arXiv preprint arXiv:1704.00028*, 2017.
- [Huang *et al.*, 2014] Dong Huang, Shitong Yao, Yi Wang, and Fernando De La Torre. Sequential max-margin event detectors. In *European conference on computer vision*, pages 410–424. Springer, 2014.
- [Kingma *et al.*, 2014] Diederik P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In *Advances in Neural Information Processing Systems*, pages 3581–3589, 2014.
- [Koller and Friedman, 2009] Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- [Li *et al.*, 2016] Chongxuan Li, Jun Zhu, and Bo Zhang. Max-margin deep generative models for (semi-) supervised learning. *arXiv preprint arXiv:1611.07119*, 2016.
- [Li *et al.*, 2017] Yunzhu Li, Jiaming Song, and Stefano Ermon. Infogail: Interpretable imitation learning from visual demonstrations. In *Advances in Neural Information Processing Systems*, pages 3815–3825, 2017.
- [Miyato *et al.*, 2017] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *arXiv preprint arXiv:1704.03976*, 2017.
- [Mohamed and Lakshminarayanan, 2016] Shakir Mohamed and Balaji Lakshminarayanan. Learning in implicit generative models. *arXiv preprint arXiv:1610.03483*, 2016.
- [Papandreou *et al.*, 2015] George Papandreou, Liang-Chieh Chen, Kevin Murphy, and Alan L Yuille. Weakly- and semi-supervised learning of a dcnn for semantic image segmentation. *arXiv preprint arXiv:1502.02734*, 2015.
- [Pathak *et al.*, 2015] Deepak Pathak, Philipp Krahenbuhl, and Trevor Darrell. Constrained convolutional neural networks for weakly supervised segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1796–1804, 2015.
- [Radford *et al.*, 2015] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- [Ratner *et al.*, 2016] Alexander J Ratner, Christopher M De Sa, Sen Wu, Daniel Selsam, and Christopher Ré. Data programming: Creating large training sets, quickly. In *Advances in Neural Information Processing Systems*, pages 3567–3575, 2016.
- [Richardson and Domingos, 2006] Matthew Richardson and Pedro Domingos. Markov logic networks. *Machine learning*, 62(1):107–136, 2006.
- [Salimans *et al.*, 2016] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems*, pages 2234–2242, 2016.
- [Shcherbatyi and Andres, 2016] Iaroslav Shcherbatyi and Bjoern Andres. Convexification of learning from constraints. In *German Conference on Pattern Recognition*, pages 79–90. Springer, 2016.
- [Stewart and Ermon, 2017] Russell Stewart and Stefano Ermon. Label-free supervision of neural networks with physics and domain knowledge. In *AAAI*, pages 2576–2582, 2017.
- [Xu *et al.*, 2017] Jingyi Xu, Zilu Zhang, Tal Friedman, Yitao Liang, and Guy Van den Broeck. A semantic loss function for deep learning with symbolic knowledge. *arXiv preprint arXiv:1711.11157*, 2017.
- [Yang and Ramanan, 2013] Yi Yang and Deva Ramanan. Articulated human detection with flexible mixtures of parts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(12):2878–2890, 2013.
- [Zamora-Martínez *et al.*, 2014] F Zamora-Martínez, P Romeu, P Botella-Rocamora, and J Pardo. Online learning of indoor temperature forecasting models towards energy efficiency. *Energy and Buildings*, 83:162–172, 2014.