LEARNING APPROXIMATE INFERENCE NETWORKS FOR STRUCTURED PREDICTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Structured prediction energy networks (SPENs; Belanger & McCallum 2016) use neural network architectures to define energy functions that can capture arbitrary dependencies among parts of structured outputs. Prior work used gradient descent for inference, relaxing the structured output to a set of continuous variables and then optimizing the energy with respect to them. We replace this use of gradient descent with a neural network trained to approximate structured argmax inference. This "inference network" outputs continuous values that we treat as the output structure. We develop large-margin training criteria for joint training of the structured energy function and inference network. On multi-label classification we report speed-ups of 10-60x compared to (Belanger et al., 2017) while also improving accuracy. For sequence labeling, for simple structured energies, our approach performs comparably to exact inference while being much faster at test time. We also show how inference networks can replace dynamic programming for test-time inference in conditional random fields, suggestive for their general use for fast inference in structured settings.

1 INTRODUCTION

Energy-based modeling (LeCun et al., 2006) associates a scalar measure of compatibility to each configuration of input and output variables. Given an input $\boldsymbol{x}$, the predicted output $\hat{\boldsymbol{y}}$ is chosen by minimizing an **energy function** $E(\boldsymbol{x}, \boldsymbol{y})$. For structured prediction, the parameterization of the energy function can leverage domain knowledge about the structured output space. However, learning and prediction become complex.

Structured prediction energy networks (SPENs; Belanger & McCallum 2016) use an energy function to score structured outputs, and perform inference by using gradient descent to iteratively optimize the energy with respect to the outputs. Belanger et al. (2017) develop an "end-to-end" method that unrolls an approximate energy minimization algorithm into a fixed-size computation graph that is trainable by gradient descent. After learning the energy function, however, they still must use gradient descent for test-time inference.

We replace the gradient descent approach with a neural network trained to do inference, which we call an **inference network**. It can have any architecture such that it takes an input $\boldsymbol{x}$ and returns an output interpretable as a $\boldsymbol{y}$. As in prior work, we relax $\boldsymbol{y}$ from discrete to continuous. For multi-label classification, we use a feed-forward network that outputs a vector. We assign a single label to each dimension of the vector, interpreting its value as the probability of predicting that label. For sequence labeling, we output a distribution over predicted labels at each position in the sequence. We adapt the energy functions such that they can operate with both discrete ground truth outputs and outputs generated by our inference networks.

We define large-margin training objectives to jointly train the energy functions and inference networks. They avoid argmax computations during training, making training faster than standard SPENs. We experiment with multi-label classification using the same setup as Belanger & McCallum (2016). We show speed-ups of 10x in training time and 60x in test-time inference while also improving accuracy.

We then design a SPEN and inference network for sequence labeling by using recurrent neural networks (RNNs). We perform comparably to a conditional random field (CRF) when using the same energy function, with faster test-time inference. We also experiment with a richer energy that

includes a "label language model" that scores entire output label sequences using an RNN, showing it can improve handling of long-distance dependencies in part-of-speech tagging.

2 STRUCTURED PREDICTION ENERGY NETWORKS

We denote the space of inputs by $\mathcal{X}$. For a given input $\boldsymbol{x} \in \mathcal{X}$, we denote the space of legal structured outputs by $\mathcal{Y}(\boldsymbol{x})$. We denote the entire space of structured outputs by $\mathcal{Y} = \cup_{\boldsymbol{x} \in \mathcal{X}} \mathcal{Y}(\boldsymbol{x})$. A SPEN defines an **energy function** $E_\Theta : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ parameterized by Θ that uses a functional architecture to compute a scalar energy for an input/output pair.

We describe the SPEN for multi-label classification (MLC) from Belanger & McCallum (2016). Here, $\boldsymbol{x}$ is a fixed-length feature vector. We assume there are L labels, each of which can be on or off for each input, so $\mathcal{Y}(\boldsymbol{x}) = \{0, 1\}^L$ for all $\boldsymbol{x}$. The energy function is the sum of two terms: $E_\Theta(\boldsymbol{x}, \boldsymbol{y}) = E^{loc}(\boldsymbol{x}, \boldsymbol{y}) + E^{lab}(\boldsymbol{y})$. $E^{loc}(\boldsymbol{x}, \boldsymbol{y})$ is the sum of linear models:

$$E^{loc}(\boldsymbol{x}, \boldsymbol{y}) = \sum_{i=1}^{L} y_i b_i^\top F(\boldsymbol{x}) \tag{1}$$

where b_i is a parameter vector for label i and $F(\boldsymbol{x})$ is a multi-layer perceptron computing a feature representation for the input $\boldsymbol{x}$. $E^{lab}(\boldsymbol{y})$ scores $\boldsymbol{y}$ independent of $\boldsymbol{x}$:

$$E^{lab}(\boldsymbol{y}) = c_2^\top g(C_1 \boldsymbol{y}) \tag{2}$$

where c_2 is a parameter vector, g is an elementwise non-linearity function, and C_1 is a parameter matrix. After learning the energy function, prediction minimizes energy:

$$\hat{\boldsymbol{y}} = \operatorname*{argmin}_{\boldsymbol{y} \in \mathcal{Y}(\boldsymbol{x})} E_\Theta(\boldsymbol{x}, \boldsymbol{y}) \tag{3}$$

However, solving Eq. (3) requires combinatorial algorithms because $\mathcal{Y}$ is a discrete structured space. This becomes intractable when E_Θ does not decompose into a sum over small "parts" of $\boldsymbol{y}$. Belanger & McCallum (2016) relax this problem by allowing the discrete vector $\boldsymbol{y}$ to be continuous. We use $\mathcal{Y}_R$ to denote the relaxed output space. For MLC, $\mathcal{Y}_R(\boldsymbol{x}) = [0, 1]^L$. They solve the relaxed problem by using gradient descent to iteratively optimize the energy with respect to $\boldsymbol{y}$. Since they train with a structured large-margin objective, repeated inference is required during learning. They note that using gradient descent for this inference step is time-consuming and makes learning less stable. So Belanger et al. (2017) propose an "end-to-end" learning procedure inspired by Domke (2012). This approach performs backpropagation through each step of gradient descent, permitting more stable training but also evidently more overfitting. We compare to both methods in our experiments below.

3 INFERENCE NETWORKS FOR SPENS

Belanger & McCallum (2016) relaxed $\boldsymbol{y}$ from a discrete to a continuous vector and used gradient descent for inference. We also relax $\boldsymbol{y}$ but we use a different strategy to approximate inference. We define an **inference network** $\mathbf{A}_\Psi(\boldsymbol{x})$ parameterized by Ψ and train it with the goal that

$$\mathbf{A}_\Psi(\boldsymbol{x}) \approx \operatorname*{argmin}_{\boldsymbol{y} \in \mathcal{Y}_R(\boldsymbol{x})} E_\Theta(\boldsymbol{x}, \boldsymbol{y}) \tag{4}$$

Given an energy function E_Θ and a dataset X of inputs, we solve the following optimization problem:

$$\hat{\Psi} \leftarrow \operatorname*{argmin}_{\Psi} \sum_{\boldsymbol{x} \in X} E_\Theta(\boldsymbol{x}, \mathbf{A}_\Psi(\boldsymbol{x})) \tag{5}$$

The architecture of $\mathbf{A}_\Psi$ will depend on the task. For MLC, the same set of labels is applicable to every input, so $\boldsymbol{y}$ has the same length for all inputs. So, we can use a feed-forward network for $\mathbf{A}_\Psi$ with a vector output, treating each dimension as the prediction for a single label. For sequence labeling, each $\boldsymbol{x}$ (and therefore each $\boldsymbol{y}$) can have a different length, so we must use a network architecture for $\mathbf{A}_\Psi$ that permits different lengths of predictions. We use an RNN that returns a vector at each position of $\boldsymbol{x}$. We interpret this vector as a probability distribution over output labels at that position.

We note that the output of $\mathbf{A}_\Psi$ must be compatible with the energy function, which is typically defined in terms of the original discrete output space $\mathcal{Y}$. This may require generalizing the energy function to be able to operate both on elements of $\mathcal{Y}$ and $\mathcal{Y}_R$. For MLC, no change is required. For sequence labeling, the change is straightforward and is described below in Section 7.2.1.

4 JOINT TRAINING OF SPENS AND INFERENCE NETWORKS

Belanger & McCallum (2016) propose a structured hinge loss for training SPENs:

$$\min_{\Theta} \sum_{\langle \boldsymbol{x}_i, \boldsymbol{y}_i \rangle \in \mathcal{D}} \left[\max_{\boldsymbol{y} \in \mathcal{Y}_R(\boldsymbol{x})} (\triangle(\boldsymbol{y}, \boldsymbol{y}_i) - E_\Theta(\boldsymbol{x}_i, \boldsymbol{y}) + E_\Theta(\boldsymbol{x}_i, \boldsymbol{y}_i)) \right]_+ \tag{6}$$

where $\mathcal{D}$ is the set of training pairs, $[f]_+ = \max(0, f)$, and $\triangle(\boldsymbol{y}, \boldsymbol{y}')$ is a structured **cost** function that returns a nonnegative value indicating the difference between $\boldsymbol{y}$ and $\boldsymbol{y}'$. This loss is often referred to as "margin-rescaled" structured hinge loss (Taskar et al., 2004; Tsochantaridis et al., 2005).

However, this loss is expensive to minimize for structured models because of the "cost-augmented" inference step ($\max_{\boldsymbol{y} \in \mathcal{Y}_R(\boldsymbol{x})}$). In prior work with SPENs, this step used gradient descent. We replace this with a **cost-augmented inference network** $\mathbf{A}_\Phi(\boldsymbol{x})$. As suggested by the notation, the cost-augmented inference network $\mathbf{A}_\Phi$ and the inference network $\mathbf{A}_\Psi$ will typically have the same functional form, but use different parameters Φ and Ψ. We write our new optimization problem as:

$$\min_{\Theta} \max_{\Phi} \sum_{\langle \boldsymbol{x}_i, \boldsymbol{y}_i \rangle \in \mathcal{D}} [\triangle(\mathbf{A}_\Phi(\boldsymbol{x}_i), \boldsymbol{y}_i) - E_\Theta(\boldsymbol{x}_i, \mathbf{A}_\Phi(\boldsymbol{x}_i)) + E_\Theta(\boldsymbol{x}_i, \boldsymbol{y}_i)]_+ \tag{7}$$

We treat this optimization problem as a minimax game and find a saddle point for the game. Following Goodfellow et al. (2014), we implement this using an iterative numerical approach. We alternatively optimize Φ and Θ, holding the other fixed. Optimizing Φ to completion in the inner loop of training is computationally prohibitive and may lead to overfitting. So we alternate between one mini-batch for optimizing Φ and one for optimizing Θ. We also add L_2 regularization terms for Θ and Φ.

The objective for the cost-augmented inference network is:

$$\hat{\Phi} \leftarrow \operatorname*{argmax}_{\Phi} [\triangle(\mathbf{A}_\Phi(\boldsymbol{x}_i), \boldsymbol{y}_i) - E_\Theta(\boldsymbol{x}_i, \mathbf{A}_\Phi(\boldsymbol{x})_i) + E_\Theta(\boldsymbol{x}_i, \boldsymbol{y}_i)]_+ \tag{8}$$

That is, we update Φ so that $\mathbf{A}_\Phi$ yields an output that has low energy and high cost, in order to mimic cost-augmented inference. The energy parameters Θ are kept fixed. There is an analogy here to the generator in generative adversarial networks (Goodfellow et al., 2014): $\mathbf{A}_\Phi$ is trained to produce a high-cost structured output that is also appealing to the current energy function. To help stabilize training of Φ, we add several terms to this objective, discussed below in Section 5.

The objective for the energy function is:

$$\hat{\Theta} \leftarrow \operatorname*{argmin}_{\Theta} [\triangle(\mathbf{A}_\Phi(\boldsymbol{x}_i), \boldsymbol{y}_i) - E_\Theta(\boldsymbol{x}_i, \mathbf{A}_\Phi(\boldsymbol{x}_i)) + E_\Theta(\boldsymbol{x}_i, \boldsymbol{y}_i)]_+ + \lambda \|\Theta\|_2^2 \tag{9}$$

That is, we update Θ so as to widen the gap between the cost-augmented and ground truth outputs. There is an analogy here to the discriminator in generative adversarial networks. The energy function is updated so as to enable it to distinguish "fake" outputs produced by $\mathbf{A}_\Phi$ from real outputs $\boldsymbol{y}_i$.

Training iterates between updating Φ and Θ using the objectives above.

4.1 TEST-TIME INFERENCE

After training, we want to use an inference network $\mathbf{A}_\Psi$ defined in Eq. (4). However, training only gives us a cost-augmented inference network $\mathbf{A}_\Phi$. Since $\mathbf{A}_\Psi$ and $\mathbf{A}_\Phi$ have the same functional form, we can use Φ to initialize Ψ, then do additional training on $\mathbf{A}_\Psi$ as in Eq. (5) where X is the training or validation set. This step helps the resulting inference network to produce outputs with lower energy, as it is no longer affected by the cost function. Since this procedure does not use the output labels of the $\boldsymbol{x}$'s in X, it could also be applied to the test data in a transductive setting.

4.2 VARIATIONS AND SPECIAL CASES

This approach also permits us to use large-margin structured prediction with slack rescaling (Tsochantaridis et al., 2005). Slack rescaling can yield higher accuracies than margin rescaling, but requires "cost-scaled" inference during training which is intractable for many classes of output structures.

However, we can use our notion of inference networks to circumvent this tractability issue and approximately optimize the slack-rescaled hinge loss, yielding the following optimization problem:

$$\min_{\Theta} \max_{\Phi} \sum_{\langle \boldsymbol{x}_i, \boldsymbol{y}_i \rangle \in \mathcal{D}} \triangle(\mathbf{A}_\Phi(\boldsymbol{x}_i), \boldsymbol{y}_i)[1 - E_\Theta(\boldsymbol{x}_i, \mathbf{A}_\Phi(\boldsymbol{x}_i)) + E_\Theta(\boldsymbol{x}_i, \boldsymbol{y}_i)]_+ \tag{10}$$

Using the same argument as above, we can also break this into alternating optimization of Φ and Θ.

We can optimize a structured perceptron (Collins, 2002) version by using the margin-rescaled hinge loss (Eq. (7)) and fixing $\triangle(\mathbf{A}_\Phi(\boldsymbol{x}_i), \boldsymbol{y}_i) = 0$. When using this loss, the cost-augmented inference network is actually a test-time inference network, because the cost is always zero, so using this loss may lessen the need to retune the inference network after training.

When we fix $\triangle(\mathbf{A}_\Phi(\boldsymbol{x}_i), \boldsymbol{y}_i) = 1$, then margin-rescaled hinge is equivalent to slack-rescaled hinge. While using $\triangle = 1$ is not useful in standard max-margin training with exact argmax inference (because the cost has no impact on optimization when fixed to a positive constant), it is potentially useful in our setting. Consider our SPEN objectives with $\triangle = 1$:

$$[1 - E_\Theta(\boldsymbol{x}_i, \mathbf{A}_\Phi(\boldsymbol{x}_i)) + E_\Theta(\boldsymbol{x}_i, \boldsymbol{y}_i)]_+ \tag{11}$$

There will always be a nonzero difference between the two energies because $\mathbf{A}_\Phi(\boldsymbol{x}_i)$ will never exactly equal the discrete vector $\boldsymbol{y}_i$. Since there is no explicit minimization over all discrete vectors $\boldsymbol{y}$, this case is more similar to a "contrastive" hinge loss which seeks to make the energy of the true output lower than the energy of a particular "negative sample" by a margin of at least 1.

In our experiments, we will compare four hinge losses for training SPENs: margin-rescaled (Eq. (7)), slack-rescaled (Eq. (10)), perceptron (margin-rescaled with $\triangle = 0$), and contrastive ($\triangle = 1$).

5 Improving Training for Inference Networks

We found that the alternating nature of the optimization led to difficulties during training. Similar observations have been noted about other alternative optimization settings, especially those underlying generative adversarial networks (Salimans et al., 2016). Below we describe several techniques we found to help stabilize training, which are optional terms added to the objective in Eq. (8).

L_2 **Regularization:** We use L_2 regularization, adding the penalty term $\|\Phi\|_2^2$ with coefficient λ_1.

Entropy Regularization: We add an entropy-based regularizer $\text{loss}_\text{H}(\mathbf{A}_\Phi(\boldsymbol{x}))$ defined for the problem under consideration. For MLC, the output of $\mathbf{A}_\Phi(\boldsymbol{x})$ is a vector of scalars in $[0, 1]$, one for each label, where the scalar is interpreted as a label probability. The entropy regularizer loss_H is the sum of the entropies over these label binary distributions. For sequence labeling, where the length of $\boldsymbol{x}$ is N and where there are L unique labels, the output of $\mathbf{A}_\Phi(\boldsymbol{x})$ is a length-N sequence of length-L vectors, each of which represents the distribution over the L labels at that position in $\boldsymbol{x}$. Then, loss_H is the sum of entropies of these label distributions across positions in the sequence.

When tuning the coefficient λ_2 for this regularizer, we consider both positive and negative values, permitting us to favor either low- or high-entropy distributions as the task prefers. [1]

Local Cross Entropy Loss: We add a local (non-structured) cross entropy $\text{loss}_\text{CE}(\mathbf{A}_\Phi(\boldsymbol{x}_i), \boldsymbol{y}_i)$ defined for the problem under consideration. We only experiment with this loss for sequence labeling. It is the sum of the label cross entropy losses over all positions in the sequence. This loss provides more explicit feedback to the inference network, helping the optimization procedure to find a solution that minimizes the energy function while also correctly classifying individual labels. It can also be viewed as a multi-task loss for the inference network.

Regularization Toward Pretrained Inference Network: We add the penalty $\|\Phi - \Phi_0\|_2^2$ where Φ_0 is a pretrained network, e.g., a local classifier trained to independently predict each part of $\boldsymbol{y}$. We only experiment with this regularizer for the MLC tasks.

[1] MLC, encouraging lower entropy distributions worked better, while for sequence labeling, higher entropy was better. Further research is required to gain understanding of the role of entropy regularization in such alternating optimization settings.

Each additional term has its own tunable hyperparameter. Finally we obtain:

$$\hat{\Phi} \leftarrow \underset{\Phi}{\operatorname{argmax}} \; [\triangle(\mathbf{A}_\Phi(\boldsymbol{x}_i), \boldsymbol{y}_i) - E_\Theta(\boldsymbol{x}_i, \mathbf{A}_\Phi(\boldsymbol{x}_i)) + E_\Theta(\boldsymbol{x}_i, \boldsymbol{y}_i)]_+ \tag{12}$$

$$-\lambda_1 \|\Phi\|_2^2 + \lambda_2 \mathrm{loss}_{\mathrm{H}}(\mathbf{A}_\Phi(\boldsymbol{x}_i)) - \lambda_3 \mathrm{loss}_{\mathrm{CE}}(\mathbf{A}_\Phi(\boldsymbol{x}_i), \boldsymbol{y}_i) - \lambda_4 \|\Phi - \Phi_0\|_2^2 \tag{13}$$

6 Related Work

Our methods are reminiscent of other alternating optimization problems like that underlying generative adversarial networks (GANs; Goodfellow et al. 2014). GANs are based on a minimax game and have a value function that one agent (a discriminator D) seeks to maximize and another (a generator G) seeks to minimize. By their analysis, a log loss discriminator converges to a degenerate uniform solution. When using hinge loss, we can get a non-degenerate discriminator while matching the data distribution (Dai et al., 2017). Our formulation is closer to this hinge loss version of the GAN.

Our approach is also related to knowledge distillation (Ba & Caruana, 2014; Hinton et al., 2015), which refers to strategies in which one model (a "student") is trained to mimic another (a "teacher"). Typically, the teacher is a larger, more accurate model but which is too computationally expensive to use at test time. Urban et al. (2016) train shallow networks using image classification data labeled by an ensemble of deep teacher nets. Geras et al. (2016) train a convolutional network to mimic an LSTM for speech recognition. Others have explored knowledge distillation for sequence-to-sequence learning (Kim & Rush, 2016) and parsing (Kuncoro et al., 2016).

Since we train a single inference network for an entire dataset, our approach is also related to "amortized inference" (Srikumar et al., 2012; Gershman & Goodman, 2014; Paige & Wood, 2016). These methods precompute or save solutions to subproblems for faster overall computation. It is likely that inference networks devote more modeling capacity to the most frequent substructures in the data. A kind of inference network is used in variational autoencoders (Kingma & Welling, 2013) to approximate posterior inference for training generative models.

Our methods are also related to work in structured prediction that seeks to approximate structured models with factorized ones, e.g., mean-field approximations in graphical models (Koller & Friedman, 2009; Krähenbühl & Koltun, 2011).

There are other settings in which gradient descent is used for inference, e.g., image generation applications like DeepDream (Mordvintsev et al., 2015) and neural style transfer (Gatys et al., 2015), as well as machine translation (Hoang et al., 2017). In these and related settings, gradient descent has started to be replaced by inference networks. Our results below provide more evidence for making this transition. An alternative to what we pursue here would be to obtain an easier convex optimization problem for inference via input convex neural networks (Amos et al., 2017).

7 Experiments

In Sec. 7.1 we compare our approach to previous work on training SPENs for MLC. We compare accuracy and speed, finding our approach to outperform prior work. We then perform experiments with sequence labeling tasks in Sec. 7.2.

7.1 Multi-Label Classification

We use the MLC datasets used by Belanger & McCallum (2016): Bibtex, Delicious, and Bookmarks. Dataset statistics are shown in Table 8 in the Appendix. For Bibtex and Delicious, we follow Belanger and McCallum and tune the hyperparameters using a different sampling of train and test data, then use the standard train/test split for final experimentation using the tuned hyperparameters. For Bookmarks, we use the same train/dev/test split as (Belanger & McCallum, 2016). For evaluation, we report the example averaged (macro averaged) F1 measure.

We use the SPEN for MLC described in Section 2 and also used by Belanger & McCallum (2016). [2] We pretrain the feature networks $F(\boldsymbol{x})$ by minimizing independent-label cross entropy for 10 epochs

[2] The feature representation network $F(\boldsymbol{x})$, we use feed-forward networks with two hidden layers, using their same layer widths: 150 for Bibtex/Bookmarks and 250 for Delicious.

Table 1: Development F1 for Bookmarks when comparing losses and inference network retuning.

dataset	hinge loss	-retuning	+retuning
Bookmarks	margin rescaling	38.51	38.68
	slack rescaling	38.57	38.62
	perceptron (MR, $\triangle = 0$)	38.55	38.70
	contrastive ($\triangle = 1$)	38.80	38.88

Table 2: Test F1 when comparing methods on multi-label classification datasets.

	MLP	SPEN	SPEN (E2E)	SPEN (Ours)
Bibtex	38.9	42.2	38.1	42.2
Bookmarks	33.8	34.4	33.9	37.6
Delicious	37.8	37.5	34.4	37.5
average	36.8	38.0	35.5	39.1

using Adam (Kingma & Ba, 2014) with learning rate 0.001. While training SPENs, we only update the parameters of the energy function (Θ) and the inference network (Φ), keeping the feature network parameters $F(\boldsymbol{x})$ fixed. We use Adam with learning rate 0.001 to train Θ and Φ. We tune λ (the L_2 regularization strength for Θ) over the set $\{0.01, 0.001, 0.0001\}$.

The inference networks are feed-forward networks with two hidden layers, using the same architectures as the feature networks $F(\boldsymbol{x})$. This permits us to initialize inference network parameters Φ using pretrained feature network parameters. For the output, we use an affine transformation layer with a sigmoid nonlinearity function, so the output values are in the range $(0, 1)$. We interpret each value as the probability of predicting the corresponding label. We obtain discrete predictions by thresholding at a threshold τ tuned to maximize F1 on the development data.[3] We add three terms to the inference network objective from Section 5: L_2 regularization ($\lambda_1 \in \{0.01, 0.001, 0.0001\}$), entropy regularization ($\lambda_2 = 1$), and regularization toward the pretrained feature network ($\lambda_4 \in \{0, 1, 10\}$). [4]

Comparison of Loss Functions and Impact of Inference Network Retuning. We first report results comparing the four loss functions from Section 4.2. Table 1 shows results on the development set for Bookmarks, the largest of the three datasets. We find performance to be highly similar across the losses, with the contrastive loss appearing slightly better than the others.

After training, we "retune" the inference network as specified by Eq. (5) on the development set for 20 epochs using a smaller learning rate of 0.00001. Table 1 shows slightly higher F1 for all losses with retuning. We were surprised to see that the final cost-augmented inference network performs well as a test-time inference network. This suggests that by the end of training, the cost-augmented network may be approaching the argmin and that there may not be much need for retuning.

When using $\triangle = 0$ or 1, retuning leads to the same small gain as when using the margin-rescaled or slack-rescaled losses. Here the gain is presumably from adjusting the inference network for other inputs rather than from converting it from a cost-augmented to a test-time inference network.

Comparison to Prior Work. Table 2 shows results comparing to prior work. The MLP and SPEN baseline results are taken from (Belanger & McCallum, 2016). We obtained the "SPEN (E2E)" (Belanger et al., 2017) results by running the code available from the authors on these datasets. This method constructs a recurrent neural network that performs gradient-based minimization of the energy with respect to $\boldsymbol{y}$. They noted in their software release that, while this method is more stable, it is prone to overfitting and actually performs worse than the original SPEN. We indeed find this to be the case, as SPEN (E2E) underperforms SPEN on all three datasets.

Our method uses the contrastive hinge loss and retunes the inference network; these decisions were made based on the results in the previous section. Our approach achieves the best average performance across the three datasets. It performs especially well on Bookmarks, which is the largest of the three.

[3] τ is chosen from $[0, 0.01, 0.02, 0.03, 0.04, 0.05, 0.1, 0.15, 0.2, 0.25, 0.3, 0.35, 0.4, 0.45, 0.5, 0.55, 0.6, 0.65, 0.7, 0.75]$ as also done by Belanger & McCallum (2016).

[4] Margin rescaling and slack rescaling use squared L_2 distance for $\triangle$.

Table 3: Training and test-time inference speed comparison (examples/sec).

	Training Speed (examples/sec)			Testing Speed (examples/sec)		
	Bibtex	Bookmarks	Delicious	Bibtex	Bookmarks	Delicious
MLP	21670	19591	26158	90706	92307	113750
SPEN (E2E)	551	559	383	1420	1401	832
SPEN (Ours)	5533	5467	4667	94194	88888	112148

Speed Comparison. Table 3 compares training and test-time inference speed among the different methods. We only report speeds of methods that we ran.[5] The SPEN (E2E) times were obtained using code obtained from Belanger and McCallum. We suspect that SPEN training would be comparable to or slower than SPEN (E2E). Our method can process examples during training about 10 times as fast as the end-to-end SPEN, and 60-130 times as fast during test-time inference. In fact, at test time, our method is roughly the same speed as the MLP baseline, since our inference networks use the same architecture as the feature networks which form the MLP baseline. Compared to the MLP, the training of our method takes significantly more time overall because of joint training of the energy function and inference network, but fortunately the test-time inference is comparable.

7.2 Sequence Labeling

We also evaluate our methods on sequence labeling. We perform experiments with Twitter part-of-speech (POS) tagging here. Named entity recognition experiments are reported in the Appendix.

7.2.1 Energy Functions for Sequence Labeling

The input space $\mathcal{X}$ is now the set of all sequences of symbols drawn from a vocabulary. For an input sequence $\boldsymbol{x}$ of length N, where there are L possible output labels for each position in $\boldsymbol{x}$, the output space $\mathcal{Y}(\boldsymbol{x})$ is $[L]^N$, where the notation $[q]$ represents the set containing the first q positive integers. We define $\boldsymbol{y} = \langle y_1, y_2, .., y_N \rangle$ where each y_i ranges over possible output labels, i.e., $y_i \in [L]$.

When defining our energy for sequence labeling, we take inspiration from bidirectional LSTMs (BLSTMs; Hochreiter & Schmidhuber 1997) and conditional random fields (CRFs; Lafferty et al. 2001). A "linear chain" CRF uses two types of features: one capturing the connection between an output label and $\boldsymbol{x}$ and the other capturing the dependence between neighboring output labels. We use a BLSTM to compute feature representations for $\boldsymbol{x}$. We use $f(\boldsymbol{x}, t) \in \mathbb{R}^d$ to denote the "input feature vector" for position t, defining it to be the d-dimensional BLSTM hidden vector at t.

We then define the following energy function:

$$E_\Theta(\boldsymbol{x}, \boldsymbol{y}) = -\left(\sum_t U_{y_t}^\top f(\boldsymbol{x}, t) + \sum_t W_{y_{t-1}, y_t} \right) \tag{14}$$

where $U_i \in \mathbb{R}^d$ is a parameter vector for label i and the parameter matrix $W \in \mathbb{R}^{L \times L}$ contains label pair parameters. The full set of parameters Θ includes the U_i vectors, W, and the parameters of the BLSTM. The above energy only permits discrete $\boldsymbol{y}$. For the general case that permits relaxing $\boldsymbol{y}$ to be continuous, we treat each y_t as a vector. It will be one-hot for the ground truth $\boldsymbol{y}$ and will be a vector of label probabilities for relaxed $\boldsymbol{y}$'s. Then the general energy function is:

$$E_\Theta(\boldsymbol{x}, \boldsymbol{y}) = -\left(\sum_t \sum_{i=1}^{L} y_{t,i} \left(U_i^\top f(\boldsymbol{x}, t) \right) + \sum_t y_{t-1}^\top W y_t \right) \tag{15}$$

where $y_{t,i}$ is the ith entry of the vector y_t. In the discrete case, this entry is 1 for a single i and 0 for all others, so this energy reduces to Eq. (14) in that case. In the continuous case, this scalar indicates the probability of the tth position being labeled with label i. For the label pair terms in this general energy function, we use a bilinear product between the vectors y_{t-1} and y_t using parameter matrix W, which also reduces to Eq. (14) when they are one-hot vectors. Below we refer to Eq. (15) as the "pairwise energy" and we refer to a version of it with only the unary label terms as the "local energy".

[5]The MLP F1 scores above were taken from Belanger & McCallum (2016), but the MLP timing results reported in Table 3 are from our own experimental replication of their results.

Table 4: Accuracies on Twitter POS development set.

energy	model/training	test-time inference	development accuracy
local	BLSTM, cross entropy	local (independent label prediction)	87.6
pairwise	CRF, conditional log-likelihood	Viterbi algorithm	89.1
		Inference network + cross entropy	89.7
		Inference network + entropy	89.7
		Inference network + squared L2 distance	87.7
	SPEN, margin-rescaled hinge	Inference network + cross entropy	$\mathbf{A}_\Phi$: 89.1 $\mathbf{A}_\Psi$: 89.3
		Inference network + entropy	$\mathbf{A}_\Phi$: 84.2 $\mathbf{A}_\Psi$: 86.8

Table 5: Twitter POS accuracies when comparing losses and showing the impact of retuning the inference network (where the inference networks are always trained with the cross entropy term).

	dev		+retuning	
hinge loss	-retuning	+retuning	dev	test
margin rescaling	89.1	89.3	89.3	89.4
slack rescaling	89.4	89.6	89.6	89.8
perceptron (MR, $\triangle = 0$)	89.2	89.4	89.4	88.6
contrastive ($\triangle = 1$)	88.8	89.0	89.0	89.0

Tag Language Model. In order to capture long-distance dependencies in an entire sequence of labels, we train a "tag language model" on a large corpus of automatically-tagged tweets, then include a term in the energy function representing the log-probability of the given tag sequence under this tag language model. Full details are included in the Appendix.

7.2.2 Experimental Setup

For Twitter part-of-speech (POS) tagging, we use the annotated data from Gimpel et al. (2011) and Owoputi et al. (2013) which contains $L = 25$ POS tags. For training, we combine the 1000-tweet OCT27TRAIN set and the 327-tweet OCT27DEV set. For validation, we use the 500-tweet OCT27TEST set and for testing we use the 547-tweet DAILY547 test set. We pretrain 100-dimensional skip-gram word embeddings on 56 million English tweets using `word2vec` (Mikolov et al., 2013). Details for training the tag language model are provided in the Appendix.

We compare to a BLSTM baseline which consists simply of the local energy terms in the energy function. We use hidden vectors of dimensionality $d = 100$. For our method, we use BLSTMs for the inference networks. The output layer of the inference network is a softmax function, so at every position, the inference network produces a distribution over labels at that position. During training, we use early stopping based on validation accuracy. For $\triangle$, we use L_1 distance.

We train inference networks using stochastic gradient descent (SGD) with momentum and train the energy parameters using Adam. We tune hyperparameters on the validation set; full details of tuning are provided in the Appendix. We train CRF baselines with the standard CRF training objective (conditional log-likelihood) using the standard dynamic programming algorithms (forward-backward) to compute gradients during training. For optimization we use SGD with momentum.[6]

7.2.3 Results

Comparison to Exact Inference at Training and Test Time. In Table 4, we report results for BLSTM and CRF baselines. We then train inference networks given the fixed CRF energy function with various stabilization terms. When using either entropy or cross entropy, our inference networks outperform Viterbi. When using the squared L_2 distance term which regularizes the inference network toward the pretrained BLSTM, the accuracy reduces to be very close to that of the BLSTM.

[6]We tune the learning rate ($\{0.1, 0.05, 0.02, 0.01, 0.005, 0.001\}$) and L_2 regularization coefficient (over $\{0, 1\text{e}-3, 1\text{e}-4, 1\text{e}-5, 1\text{e}-6, 1\text{e}-7\}$).

Table 6: Twitter POS dev/test accuracies of BLSTM, CRF, and our methods.

energy	model/training	test-time inference	dev	test
local	BLSTM, cross entropy	local (independent label prediction)	87.6	88.2
pairwise	CRF, conditional log-likelihood	Viterbi algorithm	89.1	89.2
		Inference network (trained w/ entropy)	89.6	89.4
		Inference network (trained w/ entropy + TLM)	89.9	90.0

Table 7: Training and test-time inference speed comparison (examples/sec) for sequence labeling.

	Training Speed (examples/sec)		Testing Speed (examples/sec)	
	Twitter POS	NER	Twitter POS	NER
BLSTM	385	385	1250	1042
CRF (dynamic programming)	250	222	500	454
SPEN (inference network)	125	118	1250	1025

When using the CRF energy function, we can compare the use of exact and approximate inference during training.[7] The lower part of Table 4 shows results when training a SPEN with the CRF energy function. We find that, despite using approximate inference during training (i.e., jointly training the inference network and the energy function), the SPEN reaches comparable accuracies to Viterbi on the development set, at least when using cross entropy as an additional stabilizing term.

Loss Function Comparison. Table 5 shows results when comparing SPEN training objectives. We see a larger difference among losses here than for MLC tasks. When using the perceptron loss, there is no margin, which leads to overfitting (89.4 on dev, 88.6 on test). The contrastive loss, which strives to achieve a margin of 1, does better on test (89.0). We also see here that margin rescaling and slack rescaling both outperform the contrastive hinge, unlike the MLC tasks. We suspect that in the case in which each input/output has a different length, using a cost that captures length is more important.

Tag Language Models. The above results only use the pairwise energy; no results used the tag language model (TLM). Table 6 shows results when using the TLM as an additional term in training the inference network for the CRF energy, where we tune the weight of the TLM term over the set $\{0.1, 0.2, 0.5\}$. We find a gain of 0.6 on the test set. We investigated the improvements and found some to involve corrections that seemingly stem from handling non-local dependencies better. Table 9 in the Appendix shows examples in which the model with the tag language model appears to be better at using the broader context when making tagging decisions. These results suggest that our method of training inference networks can be used to add rich global features to structured prediction, though we leave a thorough exploration of features to future work.

Speed Comparison.

Table 7 compares the training and test-time inference speed of several methods. While our method is slower than the baselines during training, it is faster than the CRF at test time, operating at essentially the same speed as the BLSTM baseline while being more accurate.

8 Conclusions and Future Work

We presented ways to jointly train structured energy functions and inference networks using large-margin objectives. The energy function captures arbitrary dependencies among the labels, while the inference networks learns to capture the properties of the energy in an efficient manner, yielding fast test-time inference. Future work includes exploring the space of network architectures for inference networks to balance accuracy and efficiency, experimenting with additional global terms in structured energy functions, and exploring richer structured output spaces such as trees and sentences.

[7] We note the difference in training criteria between the CRF (log loss) and our methods (hinge losses); nonetheless, others have shown only small differences due to this (Gimpel & Smith, 2010).

REFERENCES

Brandon Amos, Lei Xu, and J. Zico Kolter. Input convex neural networks. *ICML*, 2017.

Jimmy Ba and Rich Caruana. Do deep nets really need to be deep? In *Advances in NIPS*. 2014.

David Belanger and Andrew McCallum. Structured prediction energy networks. In *Proceedings of the International Conference on Machine Learning*, 2016.

David Belanger, Bishan Yang, and Andrew McCallum. End-to-end learning for structured prediction energy networks. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.

Michael Collins. Discriminative training methods for hidden Markov models: Theory and experiments with perceptron algorithms. In *Proceedings of the ACL-02 conference on Empirical methods in natural language processing-Volume 10*, pp. 1–8. Association for Computational Linguistics, 2002.

Zihang Dai, Amjad Almahairi, Bachman Philip, Eduard Hovy, and Courville Aaron. Calibrating energy-based generative adversarial networks. In *ICLR*, 2017.

Justin Domke. Generic methods for optimization-based modeling. In Neil D. Lawrence and Mark A. Girolami (eds.), *Proc. of AISTATS*, 2012.

Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. A neural algorithm of artistic style. *CoRR*, abs/1508.06576, 2015.

Krzysztof J. Geras, Abdel rahman Mohamed, Rich Caruana, Gregor Urban, Shengjie Wang, Ozlem Aslan, Matthai Philipose, Matthew Richardson, and Charles Sutton. Blending lstms into cnns. In *ICLR (workshop track)*, 2016.

Samuel Gershman and Noah Goodman. Amortized inference in probabilistic reasoning. In *Proceedings of the Cognitive Science Society*, 2014.

Kevin Gimpel and Noah A. Smith. Softmax-margin CRFs: Training log-linear models with cost functions. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pp. 733–736, Los Angeles, California, June 2010. Association for Computational Linguistics.

Kevin Gimpel, Nathan Schneider, Brendan O'Connor, Dipanjan Das, Daniel Mills, Jacob Eisenstein, Michael Heilman, Dani Yogatama, Jeffrey Flanigan, and Noah A. Smith. Part-of-speech tagging for Twitter: annotation, features, and experiments. In *Proc. ACL*, 2011.

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 NIPS*, 2014.

Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.

Cong Duy Vu Hoang, Gholamreza Haffari, and Trevor Cohn. Towards decoding as continuous optimisation in neural machine translation. In *Proc. of EMNLP*, 2017.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 1997.

Yoon Kim and Alexander M. Rush. Sequence-level knowledge distillation. In *Proc. of EMNLP*, 2016.

Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2013.

Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques - Adaptive Computation and Machine Learning*. The MIT Press, 2009. ISBN 0262013193, 9780262013192.

Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *Advances in NIPS*. 2011.

Adhiguna Kuncoro, Miguel Ballesteros, Lingpeng Kong, Chris Dyer, and Noah A. Smith. Distilling an ensemble of greedy dependency parsers into one MST parser. In *Proc. of EMNLP*, 2016.

John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning*, ICML, 2001.

Yann LeCun, Sumit Chopra, Raia Hadsell, Marc'Aurelio Ranzato, and Fu-Jie Huang. A tutorial on energy-based learning. In *Predicting Structured Data*. MIT Press, 2006.

Xuezhe Ma and Eduard Hovy. End-to-end sequence labeling via bi-directional lstm-cnns-crf. In *Proc. of ACL*, 2016.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in NIPS*, 2013.

Alexander Mordvintsev, Christopher Olah, and Mike Tyka. Deepdream-a code example for visualizing neural networks. *Google Res*, 2015.

Olutobi Owoputi, Brendan O'Connor, Chris Dyer, Kevin Gimpel, Nathan Schneider, and Noah A. Smith. Improved part-of-speech tagging for online conversational text with word clusters. In *Proc. NAACL*, 2013.

Brooks Paige and Frank Wood. Inference networks for sequential monte carlo in graphical models. In *International Conference on Machine Learning*, pp. 3040–3049, 2016.

Jeffrey Pennington, Richard Socher, and Christopher D. Manning. Glove: Global vectors for word representation. In *Proc. of EMNLP*, 2014.

Lev Ratinov and Dan Roth. Design challenges and misconceptions in named entity recognition. In *Proc. of CoNLL*, 2009.

Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. Improved techniques for training gans. In *Advances in NIPS*. 2016.

Vivek Srikumar, Gourab Kundu, and Dan Roth. On amortizing inference cost for structured prediction. Proc. of EMNLP, 2012.

Ben Taskar, Carlos Guestrin, and Daphne Koller. Max-margin markov networks. In *Advances in NIPS*, 2004.

Erik F. Tjong Kim Sang and Fien De Meulder. Introduction to the conll-2003 shared task: Language-independent named entity recognition. In *Proc. of CONLL*, 2003.

Ioannis Tsochantaridis, Thorsten Joachims, Thomas Hofmann, and Yasemin Altun. Large margin methods for structured and interdependent output variables. *Journal of machine learning research*, 2005.

Gregor Urban, Krzysztof J. Geras, Samira Ebrahimi Kahou, Ozlem Aslan, Shengjie Wang, Rich Caruana, Abdel-rahman Mohamed, Matthai Philipose, and Matthew Richardson. Do deep convolutional nets really need to be deep? *arXiv preprint arXiv:1603.05691*, 2016.

9 Appendix

9.1 Multi-Label Classification

Table 8 shows dataset statistics for the multi-label classification datasets.

9.2 Twitter POS Tagging

Figure 1 shows the learned pairwise potential matrix W in Twitter POS tagging. We can see strong correlations between labels in neighborhoods. For example, an adjective (A) is more likely to be followed by a noun (N) than a verb (V) (see row labeled "A" in the figure).

Table 8: Statistics of the multi-label classification datasets.

	# labels	# features	# train	# dev	# test
Bibtex	159	1836	4836	-	2515
Bookmarks	208	2151	48000	12000	27856
Delicious	982	501	12896	-	3185

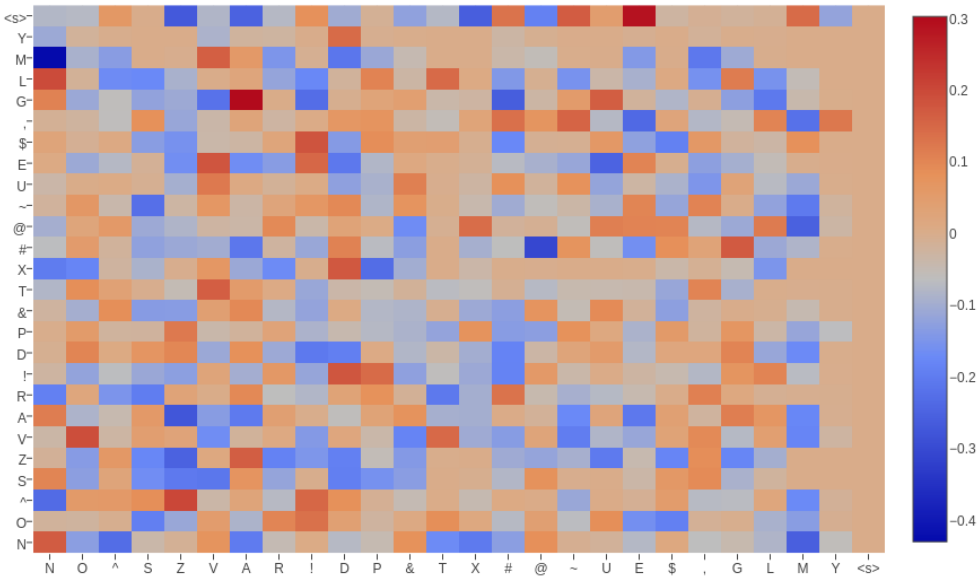

Figure 1: Learned pairwise potential matrix for Twitter POS tagging.

9.2.1 Hyperparameter Tuning

When training inference networks and SPENs for Twitter POS tagging, we use the following hyperparameter tuning. We tune the inference network learning rate ($\{0.1, 0.05, 0.02, 0.01, 0.005, 0.001\}$), L_2 regularization ($\lambda_1 \in \{0, 1\text{e}-3, 1\text{e}-4, 1\text{e}-5, 1\text{e}-6, 1\text{e}-7\}$), the entropy regularization term ($\lambda_2 \in \{0.1, 0.5, 1, 2, 5, 10\}$), the cross entropy regularization term ($\lambda_3 \in \{0.1, 0.5, 1, 2, 5, 10\}$), and the squared L2 distance ($\lambda_4 \in \{0, 0.1, 0.2, 0.5, 1, 2, 10\}$). We train the energy functions with Adam with a learning rate of 0.001 and L_2 regularization ($\lambda_1 \in \{0, 1\text{e}-3, 1\text{e}-4, 1\text{e}-5, 1\text{e}-6, 1\text{e}-7\}$).

9.2.2 Tag Language Model Description

Let Y_t denote the random variable for the tth label in the sequence Y with length m, y_t denote the value of Y_t and the distributed representation (embedding) of this label.

$$P(y_0, y_1, y_2, .., y_m, y_{m+1}) = \prod_{i=1}^{m+1} P(y_i | y_1, .., y_{m-1})$$
$$h_{t-1} = \text{LSTM}(h_{t-2}, y_{t-1})$$
$$P(y_t | y_1, .., y_{m-1}) = \text{softmax}(W h_{t-1} + b)$$

Where y_0 is the start symbol, y_{m+1} is the end symbol, W and b are the parameters of the model, LSTM is a long short-term memory network (Hochreiter & Schmidhuber, 1997), and h_t is the hidden unit of t-th label. When training the tag language model, y_i is a one-hot feature vector including the sequence end symbol.[8]

[8]We learn the representation for the sequence start symbol during the training process.

Table 9: Examples of improvements in Twitter POS tagging when using tag language model (TLM).

#	tweet (target word in bold)	predicted tags -TLM	predicted tags +TLM	gold standard
1	... we move urgently needed **#food** (wheat , flour) by truck convoy into western libya for 1st time ...	hashtag	noun	noun
2	... how can people choose **there** favourite performance , the whole thing was perfect !	adverb	determiner	determiner
3	<@mention> : the average **human** experiences about 70,000 thoughts a day .	adjective	noun	noun

After training the tag language model, we can compute the probability for any tag sequence, and we define the following term:

$$\overline{y_1}, .., \overline{y_{m+1}} = \mathrm{TLM}([y_0, \mathbf{A}_\Psi(\boldsymbol{x})_{1:m}])$$

$$E^{\mathrm{TLM}}(\mathbf{A}_\Psi(\boldsymbol{x})_{1:m}) = -\sum_{i=1}^{m+1} \log(\overline{y_i} * \mathbf{A}_\Psi(\boldsymbol{x})_i)$$

where $\overline{y_j}$ is a softmax distribution over tags for position j, $\mathbf{A}_\Psi(\boldsymbol{x})_{m+1}$ is the end-of-sequence symbol, $*$ is the inner product operation, and TLM is the function that computes the sequence of softmax distributions using the pretrained tag language model. We tune the relative weight for this extra energy term E^{TLM} over the set $\{0.1, 0.2, 0.5\}$.

For this tag language model feature, we automatically tag unlabeled tweets, then train the LSTM language model on the automatic tag sequences. In particular, we run the Twitter POS tagger from Owoputi et al. (2013) on a dataset of 303K randomly-sampled English tweets, using 300K for training and 3K for tuning hyperparameters and early stopping when training this tag language model. We train an LSTM language model on the tag sequences using stochastic gradient descent with momentum and early stopping on the validation set. We used a dropout rate of 0.5 for the LSTM hidden layer. We tune the learning rate ($\{0.1, 0.2, 0.5, 1.0\}$), the number of LSTM layers ($\{1, 2\}$), and the hidden layer size ($\{50, 100, 200\}$).

Table 9 shows examples in which our model with the global energy term (the tag language model) appears to be using broader context when making tagging decisions. Here we are comparing predictions in the test set between two settings in Table 6: "CRF, Inference network (trained w/ entropy)" vs. "CRF, Inference network (trained w/ entropy + TLM)".

In example 1, the token "#food" is tagged as a hashtag by the model without the TLM. However, it is annotated as a noun in the gold standard because it is functioning as a noun in this context. If only shown "urgently needed #food (", it is unclear whether "#food" is being used as a hashtag or as a noun in the tweet. Example 3 similarly benefits from the broader context, but this time in an adjective/noun ambiguity. In example 2, the model with the TLM is able to recognize that "there", while appearing to be an adverb based on its surface form, is actually a misspelling of the determiner "their" based on the context.

9.3 Named Entity Recognition

For named entity recognition (NER), we perform experiments on the English data from the CoNLL 2003 shared task (Tjong Kim Sang & De Meulder, 2003). This task contains sentences annotated with named entities and their types, containing 14987 training sentences, 3466 in the development set, and 3684 in the test set. There are four named entity types: PERSON, LOCATION, ORGANIZATION, and MISC. We use the BIOES tagging scheme instead of the original BIO2, following prior work (Ratinov & Roth, 2009; Ma & Hovy, 2016). There are $L = 17$ classes. We use 100-dimensional pretrained GloVe Pennington et al. (2014) embeddings trained on 6 billion words from Wikipedia and web text, which work better than other pretrained embeddings (Ma & Hovy, 2016).

Results are shown in Table 10. We see a large 4-point gap between the BLSTM and CRF, suggesting the importance of structured information for this problem. When training inference networks for the

Table 10: Named entity recognition F1 of BLSTM baselines, CRF baselines, and our methods. All inference networks below used the cross entropy stabilization term.

model	training	inference	dev	test
BLSTM	cross entropy	local	88.30	83.02
CRF	log loss	Viterbi	91.31	87.15
		inference network (1-layer)	89.70	84.52
		inference network (2-layer)	90.54	85.85
		inference network (3-layer)	91.02	86.73
SPEN	margin-rescaling	inference network (1-layer)	89.89	85.37
	slack rescaling	inference network (1-layer)	89.98	85.06

trained CRF (rows labelled "CRF + InfNet"), we find that we need multiple layers in the inference network in order to approach the result of the Viterbi algorithm. However, the 1-layer inference network result of 84.52 F1 is as fast in test-time inference as the BLSTM baseline which yields 83.02 F1. Also, while the SPEN losses underperform the CRF with Viterbi, they slightly outperform the CRF with the 1-layer inference network while having the same test-time inference speed.