

Can We Learn Heuristics For Graphical Model Inference Using Reinforcement Learning?

Safa Messaoud, Maghav Kumar, Alexander G. Schwing

University of Illinois at Urbana-Champaign

{messau2, mkumar10, aschwing}@illinois.edu

Abstract

Combinatorial optimization is frequently used in computer vision. For instance, in applications like semantic segmentation, human pose estimation and action recognition, programs are formulated for solving inference in Conditional Random Fields (CRFs) to produce a structured output that is consistent with visual features of the image. However, solving inference in CRFs is in general intractable, and approximation methods are computationally demanding and limited to unary, pairwise and hand-crafted forms of higher order potentials. In this paper, we show that we can learn program heuristics, i.e., policies, for solving inference in higher order CRFs for the task of semantic segmentation, using reinforcement learning. Our method solves inference tasks efficiently without imposing any constraints on the form of the potentials. We show compelling results on the Pascal VOC and MOTS datasets.

1. Introduction

Graphical model inference is an important combinatorial optimization task for robotics and autonomous systems. Despite significant progress in recent years due to increasingly accurate deep net models, challenges such as inconsistent bounding box detection, segmentation or image classification remain. Those inconsistencies can be addressed with Conditional Random Fields (CRFs), albeit requiring to solve an inference task which is of combinatorial complexity.

Classical algorithms to address combinatorial problems come in three paradigms: exact, approximate and heuristic. Exact algorithms are often based on solving an Integer Linear Program (ILP) using a combination of a Linear Programming (LP) relaxation and a branch-and-bound framework. Particularly for large problems, repeated solving of linear programs is computationally expensive and therefore prohibitive. Approximation algorithms address this concern, however, often at the expense of weak optimality guarantees. Moreover, approximation algorithms often involve manual construction for each problem. Seemingly easier to develop are heuristics which are generally computationally fast but

guarantees are hardly provided. In addition, tuning of hyper-parameters for a particular problem instance may be required. A fourth paradigm has been considered since the early 2000s and gained popularity again recently [93, 6, 85, 5, 27, 18]: learned algorithms. This fourth paradigm is based on the intuition that data governs the properties of the combinatorial algorithm. For instance, semantic image segmentation always deals with similarly sized problem structures or semantic patterns. It is therefore conceivable that learning to solve the problem on a given dataset uncovers strategies which are close to optimal but hard to find manually, since it is much more effective for a learning algorithm to sift through large amounts of sample problems. To achieve this, in a series of work, reinforcement learning techniques were developed [93, 6, 85, 5, 27, 18] and shown to perform well on a variety of combinatorial tasks from the traveling salesman problem and the knapsack formulation to maximum cut and minimum vertex cover.

While the aforementioned learning based techniques have been shown to perform extremely well on classical benchmarks, we are not aware of results for inference algorithms in CRFs for semantic segmentation. We hence wonder whether we can learn heuristics to address graphical model inference in semantic segmentation problems? To study this we develop a new framework for higher order CRF inference for the task of semantic segmentation using a Markov Decision Process (MDP). To solve the MDP, we assess two reinforcement learning algorithms: a Deep Q-Net (DQN) [58] and a deep net guided Monte Carlo Tree Search (MCTS) [82].

The proposed approach has two main advantages: (1) Unlike traditional approaches, it does not impose any constraints on the form of the CRF terms to facilitate effective inference. We demonstrate our claim by designing detection based higher order potentials that result in computationally intractable classical inference approaches. (2) Our method is more efficient than traditional approaches as inference complexity is linear in arbitrary potential orders while classical methods have exponential dependence on the largest clique size in general. This is due to the fact that semantic segmentation is reduced to sequentially inferring the labels

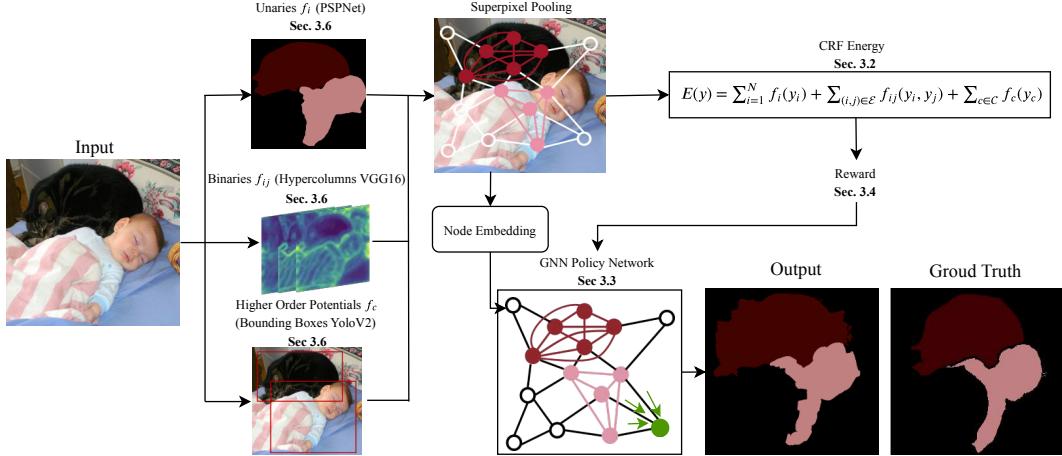


Figure 1: Pipeline of the proposed approach. Inference in a higher order CRF is solved using reinforcement learning for the task of semantic segmentation. For Pascal VOC, unaries are obtained from PSPNet [94], pairwise potentials are computed using hypercolumns from VGG16 [30] and higher order potentials are based on detection bounding boxes from YoloV2 [69]. The policy network is modeled as a graph embedding network [17] following the CRF graph structure. It sequentially produces the labeling of every node (superpixel).

of every variable based on a learned policy, without use of any iterative or search procedure.

We evaluate the proposed approach on two benchmarks: (1) the Pascal VOC semantic segmentation dataset [19], and (2) the MOTS multi-object tracking and segmentation dataset [86]. We demonstrate that our method outperforms traditional inference algorithms while being more efficient.

2. Related Work

We first review work on semantic segmentation before discussing learning of combinatorial optimizers.

Semantic Segmentation: In early 2000, classifiers were locally applied to images to generate segmentations [42] which resulted in a noisy output. To address this concern, as early as 2004, He *et al.* [33] applied Conditional Random Fields (CRFs) [43] and multi-layer perceptron features. For inference, Gibbs sampling was used, since MAP inference is NP-hard due to the combinatorial nature of the program. Progress in combinatorial optimization for flow-based problems in the 1990s and early 2000s [21, 23, 26, 9, 7, 8, 10, 40] showed that min-cut solvers can find the MAP solution of sub-modular energy functions of graphical models for binary segmentation. Approximation algorithms like swap-moves and α -expansion [10] were developed to extend applicability of min-cut solvers to more than two labels. Semantic segmentation was further popularized by combining random forests with CRFs [81]. Recently, the performance on standard semantic segmentation benchmarks like Pascal VOC 2012 [19] has been dramatically boosted by convolutional networks. Both deeper [48] and wider [61, 71, 92] network architectures have been proposed. Advances like spatial pyramid pooling [94] and atrous spatial pyramid pooling [15] emerged to remedy limited receptive fields. Other approaches jointly train deep nets with

CRFs [16, 78, 28, 79, 52, 14, 96] to better capture the rich structure present in natural scenes.

CRF Inference: Algorithmically, to find the MAP configuration, LP relaxations have been extensively studied in the 2000s [74, 13, 41, 39, 22, 88, 34, 83, 35, 68, 89, 54, 53, 36, 75, 76, 77, 55, 56]. Also, CRF inference was studied as a differentiable module within a deep net [95, 51, 57, 24, 25]. However, both directions remain computationally demanding, particularly if high order potentials are involved. We therefore wonder whether recent progress in learning based combinatorial optimization yields effective algorithms for high order CRF inference in semantic segmentation.

Learning-based Combinatorial Optimization: Decades of research on combinatorial optimization, often also referred to as discrete optimization, uncovered a large amount of valuable exact, approximation and heuristic algorithms. Already in the early 2000s, but more prominently recently [93, 6, 85, 5, 27, 18], learning based algorithms have been suggested for combinatorial optimization. They are based on the intuition that instances of similar problems are often solved repeatedly. While humans have uncovered impressive heuristics, data driven techniques are likely to uncover even more compelling mechanisms. It is beyond the scope of this paper to review the vast literature on combinatorial optimization. Instead, we subsequently focus on learning based methods. Among the first is work by Boyan and Moore [6], discussing how to learn to predict the outcome of a local search algorithm in order to bias future search trajectories. Around the same time, reinforcement learning techniques were used to solve resource-constrained scheduling tasks [93]. Reinforcement learning is also the technique of choice for recent approaches addressing NP-hard tasks [5, 27, 18, 45] like the traveling salesman, knapsack, maximum cut, and minimum vertex cover problems. Similarly, promising results exist for structured predic-

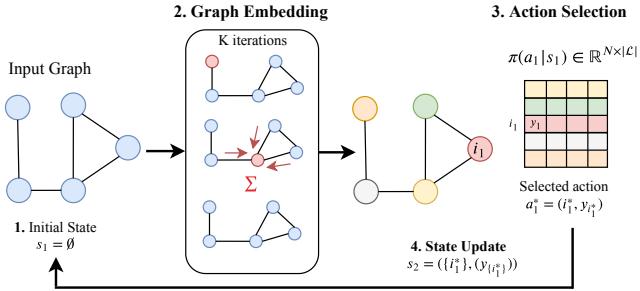


Figure 2: Illustration of one iteration of reinforcement learning for the inference task. The policy network samples an action $a_1 = (i_1^*, y_{i_1^*})$ from the learned distribution $\pi(a_1|s_1) \in \mathbb{R}^{N \times |\mathcal{L}|}$ at iteration $t = 1$.

tion problems like dialog generation [46, 90, 31], program synthesis [12, 50, 65], semantic parsing [49], architecture search [97], chunking and parsing [80], machine translation [67, 62, 4], summarization [63], image captioning [70], knowledge graph reasoning [91], query rewriting [60, 11] and information extraction [59, 66]. Instead of directly learning to solve a given program, machine learning techniques have also been applied to parts of combinatorial solvers, *e.g.*, to speed up branch-and-bound rules [44, 73, 32, 38]. We also want to highlight recent work on learning to optimize for continuous problems [47, 2].

Given those impressive results on challenging real-world problems, we wonder: can we learn programs for solving higher order CRFs for semantic image segmentation? Since CRF inference is typically formulated as a combinatorial optimization problem, we want to know how recent advances in learning based combinatorial optimization can be leveraged.

3. Approach

We first present an overview of our approach before we discuss the individual components in greater detail.

3.1. Overview

Graphical models factorize a global energy function as a sum of local functions of two types: (1) local evidence; and (2) co-occurrence information. Both cues are typically obtained from deep net classifiers which are combined in a joint energy formulation. Finding the optimal semantic segmentation configuration, *i.e.*, finding the minimizing argument of the energy, generally involves solving an NP-hard combinatorial optimization problem. Notable exceptions include energies with sub-modular co-occurrence terms.

Instead of using classical directions, *i.e.*, heuristics, exhaustive search, or relaxations, here, we assess suitability of learning based combinatorial optimization. Intuitively, we argue that CRF inference for the task of semantic segmentation exhibits an inherent similarity which can be exploited by

learning based algorithms. In spirit, this mimics the design of heuristic rules. However, different from hand-crafting those rules, we use a learning based approach. To the best of our knowledge, this is the first work to successfully apply learning based combinatorial optimization to CRF inference for semantic segmentation. We therefore first provide an overview of the developed approach, outlined in Fig. 1.

Just like classical approaches, we also use local evidence and co-occurrence information, obtained from deep nets. This information is consequently used to form an energy function defined over a Conditional Random Field (CRF). An example of a CRF with variables corresponding to superpixels (circles), pairwise potentials (edges) and higher order potentials obtained from object detections (fully connected cliques) is illustrated in Fig. 1. However, different from classical methods, we find the minimizing configuration of the energy by repeatedly applying a learned policy network. In every iteration, the policy network selects a random variable, *i.e.*, the pixel and its label by computing a probability distribution over all currently unlabeled pixels and their labels. Specifically, the pixel and label are determined by choosing the highest scoring entry in a matrix where the number of rows and columns correspond to the currently unlabeled pixels and the available labels respectively, as illustrated in Fig. 2.

3.2. Problem Formulation

Formally, given an image x , we are interested in predicting the semantic segmentation $y = (y_1, \dots, y_N) \in \mathcal{Y}$. Hereby, N denotes the total number of pixels or superpixels, and the semantic segmentation of a superpixel $i \in \{1, \dots, N\}$ is referred to via $y_i \in \mathcal{L} = \{1, \dots, |\mathcal{L}|\}$, which can be assigned one out of $|\mathcal{L}|$ possible discrete labels from the set of possible labels \mathcal{L} . The output space is denoted $\mathcal{Y} = \mathcal{L}^N$.

Classical techniques obtain local evidence $f_i(y_i)$ for every pixel or superpixel, and co-occurrence information in the form of pairwise potentials $f_{ij}(y_i, y_j)$ and higher order potentials $f_c(y_c)$. The latter assigns an energy to a clique $c \subseteq \{1, \dots, N\}$ of variables $y_c = (y_i)_{i \in c}$. For readability, we drop the dependence of the energies f_i , f_{ij} and f_c on the image x and the parameters of the employed deep nets. The goal of energy based semantic segmentation is to find the configuration y^* which has the lowest energy $E(y)$, *i.e.*,

$$y^* = \arg \min_{y \in \mathcal{Y}} E(y) \triangleq \sum_{i=1}^N f_i(y_i) + \sum_{(i,j) \in \mathcal{E}} f_{ij}(y_i, y_j) + \sum_{c \in \mathcal{C}} f_c(y_c). \quad (1)$$

Hereby, the sets \mathcal{E} and \mathcal{C} subsume respectively the captured set of pairwise and higher order co-occurrence patterns. Details about the potentials are presented in Sec. 3.6.

Solving the combinatorial program given in Eq. (1), *i.e.*, inferring the optimal configuration y^* is generally NP-hard.

Algorithm 1: Inference Procedure

```

1:  $s_1 = \emptyset$ ;
2: for  $t = 1$  to  $N$  do
3:    $a_t^* = \arg \max_{a_t \in \mathcal{A}_t} \pi(a_t | s_t)$ 
4:    $(i_t^*, y_{i_t^*}) \leftarrow a_t^*$ 
5:    $s_{t+1} = s_t \oplus (i_t^*, y_{i_t^*})$ 
6: end for
7: Return:  $\hat{y} \leftarrow s_{N+1}$ 

```

Different from existing methods, we develop a learning based combinatorial optimization heuristic for semantic segmentation with the intention to better capture the intricacies of energy minimization than can be done by hand-crafting rules. The developed heuristic sequentially labels one variable y_i , $i \in \{1, \dots, N\}$, at a time.

Formally, selection of one superpixel at a time can be formulated in a reinforcement learning context, as shown in Fig. 2. Specifically, an agent operates in $t \in \{1, \dots, N\}$ time-steps according to a policy $\pi(a_t | s_t)$ which encodes a probability distribution over actions $a_t \in \mathcal{A}_t$ given the current state s_t . The current state subsumes in selection order the indices of all currently labeled variables $I_t \subseteq \{1, \dots, N\}$ as well as their labels $y_{I_t} = (y_i)_{i \in I_t}$, *i.e.*, $s_t \in \{(I_t, y_{I_t}) : I_t \subseteq \{1, \dots, N\}, y_{I_t} \in \mathcal{L}^{|I_t|}\}$. We start with $s_1 = \emptyset$. The set of possible actions \mathcal{A}_t is the concatenation of the label spaces \mathcal{L} for all currently unlabeled pixels $j \in \{1, \dots, N\} \setminus I_t$, *i.e.*, $\mathcal{A}_t = \bigoplus_{j \in \{1, \dots, N\} \setminus I_t} \mathcal{L}$. We emphasize the difference between the concatenation operator and the product operator used to obtain the semantic segmentation output space $\mathcal{Y} = \mathcal{L}^N$, *i.e.*, the proposed approach does not operate in the product space.

As mentioned before, the policy $\pi(a_t | s_t)$ results in a probability distribution over actions $a_t \in \mathcal{A}_t$ from which we greedily select the most probable action

$$a_t^* = \arg \max_{a_t \in \mathcal{A}_t} \pi(a_t | s_t).$$

The most probable action a_t^* can be decomposed into the index for the selected variable, *i.e.*, i_t^* and its state $y_{i_t^*} \in \mathcal{L}$. We obtain the subsequent state s_{t+1} by combining the extracted variable index i_t^* and its labeling with the previous state s_t . Specifically, we obtain $s_{t+1} = s_t \oplus (i_t^*, y_{i_t^*})$ by slightly abusing the \oplus -operator to mean concatenation to a set and a list maintained within a state.

Formally, we summarize the developed reinforcement learning based semantic segmentation algorithm used for inferring a labeling \hat{y} in Alg. 1. In the following, we describe the policy function $\pi_\theta(a_t | s_t)$, which we found to work well for semantic segmentation, and different variants to learn its parameters θ .

3.3. Policy Function

We model the policy function $\pi_\theta(a_t | s_t)$ using a graph embedding network [17]. The input to the network is a

weighted graph $G(V, \mathcal{E}, w)$, where nodes $V = \{1, \dots, N\}$, correspond to variables, *i.e.*, in our case superpixels, \mathcal{E} is a set of edges connecting neighboring superpixels, as illustrated in Fig. 1 and $w : \mathcal{E} \rightarrow \mathbb{R}^+$ is the edge weight function. The weights $\{w(i, j)\}_{\{j : (i, j) \in \mathcal{E}\}}$ on the edges between a given node i and its neighbors $\{j : (i, j) \in \mathcal{E}\}$ form a distribution, obtained by normalizing the dot product between the hypercolumns [30] g_i and g_j via a softmax across neighbors. At every iteration, the state s_t is encoded in the graph G by tagging node $i \in V$ with a scalar $h_i = 1$ if the node is part of the already labeled set I_t , *i.e.*, if $i \in I_t$ and 0 otherwise. Moreover, a one-hot encoding $\tilde{y}_i \in \{0, 1\}^{|\mathcal{L}|}$ encodes the selected label of nodes $i \in I_t$. We set \tilde{y}_i to equal the all zeros vector if node i has not been selected yet.

Every node $i \in V$ is represented by a p -dimensional embedding, where p is a hyperparameter. The embedding is composed of \tilde{y}_i , h_i as well as superpixel features $b_i \in \mathbb{R}^F$ which encode appearance and bounding box characteristics that we discuss in detail in Sec. 4.

The output of the network is a $|\mathcal{L}|$ -dimensional vector π_i for each node $i \in V$, representing the scores of the $|\mathcal{L}|$ different labels for variable i .

The network iteratively generates a new representation $\mu_i^{(k+1)}$ for every node $i \in V$ by aggregating the current embeddings $\mu_i^{(k)}$ according to the graph structure \mathcal{E} starting from $\mu_i^{(0)} = \mathbf{0}$, $\forall i \in V$. After K steps, the embedding captures long range interactions between the graph features as well as the graph properties necessary to minimize the energy function E . Formally, the update rule for node i is

$$\mu_i^{(k+1)} \leftarrow \text{Relu} \left(\theta_1^{(k)} h_i + \theta_2^{(k)} \tilde{y}_i + \theta_3^{(k)} b_i + \theta_4^{(k)} \sum_{j : (i, j) \in \mathcal{E}} w(i, j) \mu_j^{(k)} \right), \quad (2)$$

where $\theta_1^{(k)} \in \mathbb{R}^p$, $\theta_2^{(k)} \in \mathbb{R}^{p \times |\mathcal{L}|}$, $\theta_3^{(k)} \in \mathbb{R}^{p \times F}$ and $\theta_4^{(k)} \in \mathbb{R}^{p \times p}$ are trainable parameters. After K steps, π_i for every unlabeled node $i \in \{1, \dots, N\} \setminus I_t$ is obtained via

$$\pi_i = \theta_5 \mu_i^{(K)} \quad \forall i \in \{1, \dots, N\} \setminus I_t, \quad (3)$$

where $\theta_5 \in \mathbb{R}^{|\mathcal{L}| \times p}$ is another trainable model parameter. We illustrate the policy function $\pi_\theta(a_t | s_t)$ and one iteration of inference in Fig. 2.

3.4. Reward Function:

To train the policy, ideally, the reward function $r_t(s_t, a_t)$ is designed such that the cumulative reward coincides exactly with the objective function that we aim at maximizing, *i.e.*, $\sum_{t=1}^N r_t(s_t, a_t) = -E(\hat{y})$, where \hat{y} is extracted from s_{N+1} . Hence, at step t , we define the reward as the difference between the value of the negative new energy E_t and the negative energy from the previous step E_{t-1} , *i.e.*, $r_t(s_t, a_t) = E_{t-1}(y_{I_{t-1}}) - E_t(y_{I_t})$, where $E_0 = 0$. Poten-

Table 1: Illustration of the energy reward computation following the two proposed reward schemes on a fully connected graph with 3 nodes.

t	i_t	E_t	$r_t = -(E_t - E_{t-1})$	$r_t = \pm 1$	Graph
0	—	0	—	—	
1	1	$f_1(y_1)$	$-f_1(y_1)$	$-1 + 2 \cdot \mathbb{1}_{\{(E_t(y_1) < E_t(\hat{y}_1)) \vee \hat{y}_1\}}$	
2	2	$f_1(y_1) + f_2(y_2) + f_{12}(y_1, y_2)$	$-f_2(y_2) - f_{12}(y_1, y_2)$	$-1 + 2 \cdot \mathbb{1}_{\{(E_t(y_1, y_2) < E_t(\hat{y}_2)) \vee \hat{y}_2\}}$	
3	3	$f_1(y_1) + f_2(y_2) + f_3(y_3) + f_{12}(y_1, y_2) + f_{23}(y_2, y_3) + f_{13}(y_1, y_3) + f_{\{1,2,3\}}(y_1, y_2, y_3)$	$-f_3(y_3) - f_{23}(y_2, y_3) - f_{13}(y_1, y_3) - f_{\{1,2,3\}}(y_1, y_2, y_3)$	$-1 + 2 \cdot \mathbb{1}_{\{(E_t(y_1, y_2, y_3) < E_t(y_1, y_2, \hat{y}_3)) \vee \hat{y}_3\}}$	

tials depending on variables that are not labeled at time t are not incorporated in the evaluation of $E_t(y_{I_t})$.

We also study a second scheme, where the reward is truncated to $+1$ or -1 , *i.e.*, $r_t(s_t, a_t) \in \{-1, 1\}$. For every selected node i_t , with label y_{i_t} , we compare the energy function $E_t(y_{I_t})$ with the one obtained when using all other labels $\hat{y}_{i_t} \in \mathcal{L} \setminus y_{i_t}$. If the chosen label y_{i_t} results in the lowest energy, the obtained reward is $+1$, otherwise it is -1 .

Note that the unary potentials result in a reward for every time step. Pairwise and high order potentials result in a sparse reward as their value is only available once all the superpixels forming the pair or clique are labeled. We illustrate the energy and reward computation on a graph with three fully connected nodes in Tab. 1.

3.5. Learning Policy Parameters

To learn the parameters θ of the policy function $\pi_\theta(a_t|s_t)$, a plethora of reinforcement learning algorithms are applicable. To provide a careful assessment of the developed approach, we study two different techniques, Q-learning and Monte-Carlo Tree Search, both of which we describe next.

Q-learning: In the context of Q-learning, we interpret the $|\mathcal{L}|$ -dimensional policy network output vector corresponding to a currently unlabeled node $i \in \{1, \dots, N\} \setminus I_t$ as the Q-values $Q(s_t, a_t; \theta)$ associated to the action a_t of selecting node i and assigning label $y_i \in \mathcal{L}$. Since we only consider actions to label currently unlabeled nodes we obtain a total of $|\mathcal{A}_t|$ different Q-values.

We perform standard Q-learning and minimize the squared loss $(z - Q(s_t, a_t; \theta))^2$, where we use target $z = \gamma \max_{a'} Q(s_{t+1}, a'; \theta) + r_t(s_t, a_t)$ for a non-terminal state. The reward is denoted r_t and detailed above. The terminal state is reached when all the nodes are labeled.

Instead of updating the Q-function based on the current sample $(s_t, a_t, r_t(s_t, a_t), s_{t+1})$, we use a replay memory populated with samples (graphs) from previous episodes. In every iteration, we select a batch of samples and perform stochastic gradient descent on the squared loss.

During the exploration phase, beyond random actions, we encourage the following three different sets of actions to generate more informative rewards for training: (1) \mathcal{M}_1 : Choosing nodes that are adjacent to the already selected ones in the graph. Otherwise, the reward will only be based on the unary terms as the pairwise term is only evaluated if the

neighbors are labeled ($t = 2$ in Tab. 1); (2) \mathcal{M}_2 : Selecting nodes with the lowest unary distribution entropy. The low entropy indicates a high confidence of the unary deep net. Hence, the labels of the corresponding nodes are more likely to be correct and provide useful information to neighbors with higher entropy in the upcoming iterations. (3) \mathcal{M}_3 : Assigning the same label to nodes forming the same higher order potential. Further details are in **Appendix B**.

Monte-Carlo Tree Search: While DQN tries to learn a policy from looking at samples representing one action at a time, MCTS has the inherent ability to update its policy after looking multiple steps ahead via a tree search procedure. At training time, through extensive simulations, MCTS builds a *statistics tree*, reflecting an empirical distribution $\pi^{\text{MCTS}}(a_t|s_t)$. Specifically, for a given image, a node in the search tree corresponds to the state s_t in our formulation and an edge corresponds to a possible action a_t . The root node is initialized to $s_1 = \emptyset$. The statistics stored at every node correspond to (1) $N(s_t)$: the number of times state s_t has been reached, (2) $N(a_t|s_t)$: the number of times action a_t was chosen in state s_t in all previous simulations, as well as (3) $\tilde{r}_t(s_t, a_t)$: the averaged reward across all simulations starting at s_t and taking action a_t . The MCTS policy is defined as $\pi^{\text{MCTS}}(a_t|s_t) = \frac{N(a_t|s_t)}{N(s_t)}$. The simulations follow an exploration-exploitation procedure modeled by a variant of the Probabilistic Upper Confidence Bound (PUCB) [82]: $U(a_t, s_t) = \frac{\tilde{r}_t(s_t, a_t)}{N(a_t|s_t)} + \pi_\theta(a_t|s_t) \frac{\sqrt{N(s_t)}}{1+N(a_t|s_t)}$. During exploration, we additionally encourage the same action sets \mathcal{M}_1 , \mathcal{M}_2 and \mathcal{M}_3 used for DQN. Also, similarly to DQN, the generated experiences (s_t, π^{MCTS}) are stored in a replay buffer. The policy network is then trained through a cross entropy loss

$$L(\theta) = - \sum_s \sum_a \pi^{\text{MCTS}}(a|s) \log \pi_\theta(a|s). \quad (4)$$

to match the empirically constructed distribution. Here, the second sum is over all valid actions from a state s sampled from the replay buffer and $\pi^{\text{MCTS}}(a|s)$ is the corresponding empirically estimated distribution.

A more detailed description of the MCTS search process, including pseudo-code is available in **Appendix B**. At inference time, we use a low budget of simulations. Final actions are taken according to the constructed $\pi^{\text{MCTS}}(a_t|s_t)$.

Table 2: Performance results for the minimizing the energy function E_t under reward scheme 1 ($R_t^1 = -(E_t - E_{t-1})$) and reward scheme 2 ($R_t^2 = \pm 1$).

Nodes	Metrics	Supervised	Unary						Unary + Pairwise						Unary + Pairwise + HOP1						Unary + Pairwise + HOP1 + HOP2															
			BP		R_t^1		R_t^2		BP		TBP		DD		L-Flip		α -Exp		R_t^1		R_t^2		BP		TBP		DD		L-Flip		α -Exp		R_t^1		R_t^2	
			DQN	MCTS	DQN	MCTS	DQN	MCTS	DQN	MCTS	DQN	MCTS	DQN	MCTS	DQN	MCTS	DQN	MCTS	DQN	MCTS	DQN	MCTS	DQN	MCTS	DQN	MCTS	DQN	MCTS								
Pascal VOC	IoU (sp)	85.21	88.59	88.04	88.19	88.59	88.59	88.73	88.73	88.73	88.73	88.72	43.31	66.51	87.91	88.73	89.26	89.27	89.27	89.27	88.58	57.43	73.37	89.55	89.66	58.34	73.85	90.05	90.09							
	IoU (p)	69.05	72.56	70.77	71.99	72.56	72.56	72.43	72.43	72.43	72.43	72.35	51.88	53.71	72.83	72.43	72.35	72.59	72.59	72.60	72.35	59.73	81.07	81.05	74.94	74.77	81.23	81.33								
MOTS	IoU (sp)	83.54	88.01	87.29	88.01	88.01	88.01	88.10	88.10	88.10	88.10	88.10	88.06	88.22	88.56	88.52	88.53	88.55	88.54	88.07	60.60	64.82	88.94	88.91	82.19	81.89	89.30	89.57								
	IoU (p)	75.88	80.64	80.47	80.64	80.64	80.64	80.68	80.68	80.68	80.68	80.68	80.54	80.86	80.75	80.91	80.91	80.93	80.91	80.65	57.36	59.73	81.07	81.05	74.94	74.77	81.23	81.33								
2008	IoU (sp)	84.91	87.39	87.34	87.39	87.39	87.39	87.55	87.55	87.56	87.55	87.55	82.23	83.67	87.80	87.84	87.95	87.96	87.96	87.95	87.54	37.80	57.66	88.73	88.69	43.99	45.67	88.43	88.21							
	IoU (p)	77.93	82.35	82.20	82.35	82.35	82.35	82.48	82.48	82.47	82.47	82.47	77.36	79.14	82.64	82.70	82.72	82.72	82.72	82.71	82.47	36.65	52.73	83.05	82.95	41.74	42.44	82.79	82.67							

The replay-memory for both DQN and MCTS is divided into two chunks. The first chunk corresponds to the unary potential, while the second chunk corresponds to the overall energy function. A node is assigned to the second chunk if its associated reward is higher than the one obtained from its unary labeling. This ensures positive rewards from all the potentials during training. Every chunk is further divided into $|\mathcal{L}|$ categories corresponding to the $|\mathcal{L}|$ classes of the selected node. This guarantees a balanced sampling of the label classes in every batch during training. Beyond DQN and MCTS, we experimented with policy gradients but could not get it to work as it is an on-policy algorithm. Reusing experiences for the structured replay buffer was crucial for success of the learning algorithm.

3.6. Energy Function

Finally we provide details on the energy function E given in Eq. (1). The unary potentials $f_i(y_i) \in \mathbb{R}^{|\mathcal{L}|}$ are obtained from a semantic segmentation deep net. The pairwise potential encodes smoothness and is computed as follows:

$$f_{i,j}(y_i, y_j) = \psi(y_i, y_j) \cdot \alpha_p \cdot \mathbb{1}_{|g_i^T g_j| < \beta_p}, \quad (5)$$

where $\psi(y_i, y_j)$ is the label compatibility function describing the co-occurrence of two classes in adjacent locations and is given by the Potts model:

$$\psi(y_i, y_j) = \begin{cases} 1 & \text{if } y_i \neq y_j \\ 0 & \text{otherwise} \end{cases}. \quad (6)$$

Moreover, $|g_i^T g_j|$ is the above defined unnormalized weight $w(i, j)$ for the edge connecting the i^{th} and j^{th} nodes, *i.e.*, superpixels. Intuitively, if the dot product between the hypercolumns g_i and g_j is smaller than a threshold β_p and the two superpixels are labeled differently, a penalty of value α_p incurs.

While the pairwise term mitigates boundary errors, we address recognition errors with two detection-based higher order potentials [3]. For this purpose, we use the YoloV2 bounding box object detector [69] as it ensures a good trade-off between speed and accuracy. Every bounding box b is presented by a tuple (l_b, c_b, I_b) , where l_b is the class label of the detected object, c_b is the confidence score of the detection

and $I_b \subseteq \{1, \dots, N\}$ is the set of superpixels that belong to the foreground detection obtained via Grab-Cut [72].

The first higher order potential (HOP1) encourages superpixels within a bounding box to take the bounding box label, while enabling recovery from false detections that do not agree with other energy types. For this purpose, we add an auxiliary variable z_b for every bounding box b . We use $z_b = 1$, if the bounding box is inferred to be valid, otherwise $z_b = 0$. Formally,

$$f(y_{I_b}, z_b) = \begin{cases} w_b \cdot c_b \cdot \sum_{i \in I_b} \mathbb{1}_{y_i = l_b} & \text{if } z_b = 0 \\ w_b \cdot c_b \cdot \sum_{i \in I_b} \mathbb{1}_{y_i \neq l_b} & \text{if } z_b = 1 \end{cases}, \quad (7)$$

where, $w_b \in \mathbb{R}$ is a weight parameter. This potential can be simplified into a sum of pairwise potentials between z_b and each y_i with $i \in I_b$, *i.e.*, $f(y_{I_b}, z_b) = \sum_{i \in I_b} f_{i,b}(y_i, z_b)$, where:

$$f_{i,b}(y_i, z_b) = \begin{cases} w_b \cdot c_b \cdot \mathbb{1}_{y_i = l_b} & \text{if } z_b = 0 \\ w_b \cdot c_b \cdot \mathbb{1}_{y_i \neq l_b} & \text{if } z_b = 1 \end{cases}. \quad (8)$$

This simplification enables solving the higher order potential using traditional techniques like mean field inference [3].

To show the merit of the RL framework, we introduce another higher order potential (HOP2) that can not be seamlessly reduced to a pairwise one:

$$f(y_{I_b}) = \lambda_b \cdot \mathbb{1}_{(\sum_{i \in I_b} y_i = l) < \frac{|I_b|}{C}}, \quad (9)$$

with λ_b and C being scalar parameters. This potential is evaluated for bounding boxes with special characteristics to encourage the superpixels in the bounding box to be of label l . Intuitively, if the number of superpixels $i \in I_b$ having label l is less than a threshold $\frac{|I_b|}{C}$, a penalty λ_b incurs. For Pascal VOC, we evaluate the potential on bounding boxes b included in larger bounding boxes, as we noticed that the unaries frequently miss small objects overlapping with other larger objects in the image ($l = l_b$). For MOTS, we evaluate this potential on bounding boxes of type ‘pedestrians’ overlapping with bounding boxes of type ‘bicycle.’ As cyclists should not be labeled as pedestrians, we set l to be the background class. Transforming this term into a pairwise one to enable using traditional inference techniques requires an exponential number of auxiliary variables.

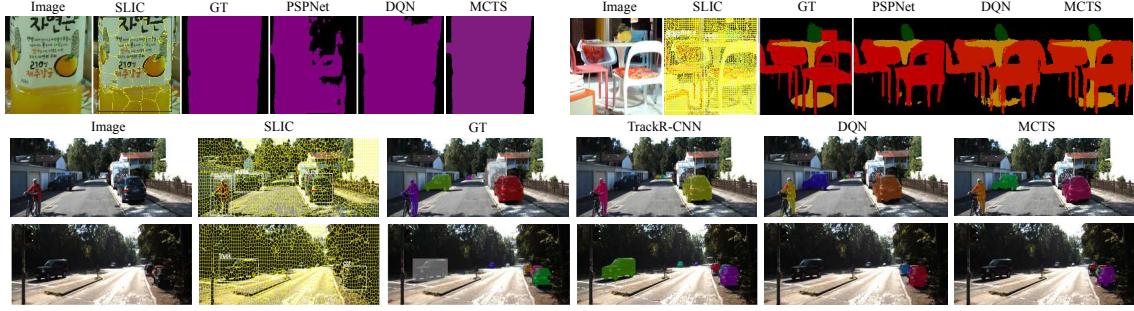


Figure 3: Success cases.

4. Experiments

In the following, we evaluate our learning based inference algorithm on Pascal VOC [19] and MOTS [19] datasets. The original Pascal VOC dataset contains 1464 training and 1449 validation images. In addition to this data, we make use of the annotations provided by [29], resulting in a total of 10582 training instances. The number of classes is 21. MOTS is a multi-object tracking and segmentation dataset for cars and pedestrians (2 classes). It consists of 12 training sequences (5027 frames) and 9 validation ones (2981 frames). In this work, we perform semantic segmentation at the level of superpixels, generated using SLIC [84]. Every superpixel corresponds to a node i in the graph as illustrated in Fig. 1. The unary potentials at the pixel level are obtained from PSPNet [94] for Pascal VOC and TrackR-CNN [86] for MOTS. The superpixels' unaries are the average of the unaries of all the pixels that belong to that superpixel. The higher order potential is based on the YoloV2 [69] bounding box detector. Additional training and implementation details are described in Appendix A.

Evaluation Metrics: As evaluation metrics, we use intersection over union (IoU) computed at the level of both superpixels (sp) and pixels (p). IoU (p) is obtained after mapping the superpixel level labels to the corresponding set of pixels.

Baselines: We compare our results to the segmentations obtained by five different solvers from three categories: (1) message passing algorithms, *i.e.*, Belief propagation (BP) [64] and Tree-reweighted Belief Propagation (TBP) [87], (2) a Lagrangian relaxation method, *i.e.*, Dual Decomposition Subgradient (DD) [37], and (3) move making algorithms, *i.e.*, Lazy Flipper (LFlip) [1] and α -expansion as implemented in [20]. Note that these solvers can not optimize our HOP2 potential. Besides, we train a supervised model that predicts the node label from the provided node features.

Performance Evaluation: We show the results of solving the program given in Eq. (1) in Tab. 2, for unary (Col. 5), unary plus pairwise (Col. 6), unary plus pairwise plus HOP1 (Col. 7) and unary plus pairwise plus HOP1 and HOP2 (Col. 8) potentials. For every potential type, we report results on graphs with superpixel numbers 50, 250 and 500 for

Pascal VOC and 2000 for MOTS, obtained from DQN and MCTS, trained each with the two reward schemes discussed in Sec. 3.5. Since MOTS has small objects, we opt for a higher number of superpixels. It is remarkable to observe that DQN and MCTS are able to learn heuristics which outperform the baselines. Interestingly, the policy has learned to produce better semantic segmentations than the ones obtained via MRF energy minimization. Guided by a reward derived just from the energy function, the graph neural network (the policy) learns characteristic node embeddings for every semantic class by leveraging the hypercolumn and bounding box features as well as the neighborhood characteristics. The supervised baseline shows low performance, which proves the merit of the learned policies. Overall MCTS performance is comparable to the DQN one. This is mainly due to the learned policies being somewhat local and focusing on object boundaries, not necessitating a large multi-step look-ahead, as we will show in the following.

In Fig. 3, we report success cases of the RL algorithms. Smoothness modeled by our energy fixed the bottle segmentation in the first image. Furthermore, our model detects missing parts of the table in the second image in the first row and a car in the image in the second row, that were missed by the unaries. Also, we show that we fix a mis-labeling of a truck as a car in the image in the third row.

Flexibility of Potentials: In Fig. 4, we show examples of improved segmentations when using the pairwise, HOP1 and HOP2 potentials respectively. The motorcycle driver segmentation improved incrementally with every potential (first image) and the cyclist is not detected anymore as a pedestrian (second image).

Generalization and scalability: The graph embedding network enables training and testing on graphs with different number of nodes, since the same parameters are used. We investigate how models trained on graphs with few nodes perform on larger graphs. As shown in Tab. 3, compelling accuracy and IoU values for generalization to graphs with up to 500, 1000, 2000, and 10000 nodes are observed when using a policy trained on graphs of 250 nodes for Pascal VOC, and to graphs with up to 5000 and 10000 nodes when using

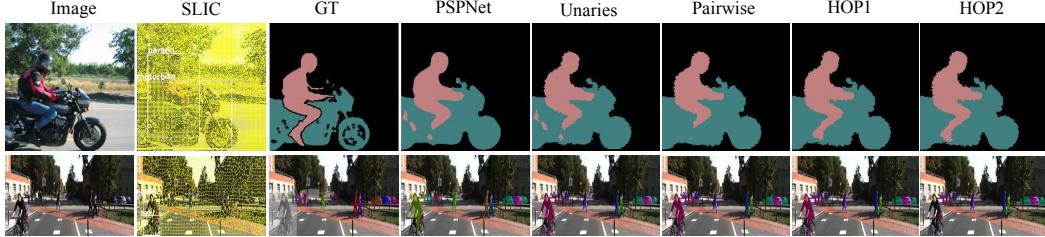


Figure 4: Output of our method for different potentials.

Table 3: Generalization of the learned policy.

Pascal VOC		PSPNet		500		1000		2000		10000	
				DQN	MCTS	DQN	MCTS	DQN	MCTS	DQN	MCTS
		IoU (sp)	—	88.74	88.73	87.58	87.61	86.36	86.39	84.66	84.67
		IoU (p)	82.61	83.06	83.01	83.71	83.73	83.74	83.78	83.80	83.82
MOTS		TrackR-CNN		5000		10000					
				DQN	MCTS	DQN	MCTS				
		IoU (sp)	—	79.81	79.80	76.73	76.69				
		IoU (p)	84.98	83.49	83.46	84.69	84.63				

Table 4: Run-time during inference in seconds for Pascal VOC dataset.

Nodes	U+P						U+P+HOP1						U+P+HOP1 +HOP2											
	BP	TBP	DD	L-Flip	α -Exp	DQN	MCTS	BP	TBP	DD	L-Flip	α -Exp	DQN	MCTS	BP	TBP	DD	L-Flip	α -Exp	DQN	MCTS			
50	0.14	0.52	0.28	0.12	0.01	0.04	0.20	0.15	0.62	0.31	0.187	0.01	0.04	0.23	0.04	0.24								
250	1.56	2.13	1.26	0.53	0.04	0.22	0.22	1.70	2.77	1.65	0.59	0.07	0.22	2.89	0.22	3.01								
500	3.26	4.76	2.82	1.07	0.12	0.52	0.77	3.37	5.37	3.70	0.97	0.22	0.53	9.17	0.52	9.69								
1000	6.63	9.65	6.84	1.80	0.30	0.78	18.5	7.22	10.4	7.47	2.25	0.36	0.78	21.6	0.78	22.8								
2000	12.3	19.9	14.8	3.57	0.70	1.70	38.3	12.7	23.9	15.1	4.47	0.72	1.70	43.2	1.72	46.2								
10000	72.8	130.9	143.7	22.6	4.81	8.23	202.1	88.7	140.1	106.9	23.5	4.72	8.25	209.7	8.20	210.3								

a policy trained on 2000 nodes for MOTS. Here, we consider the energy consisting of the combined potentials (unary, pairwise, HOP1 and HOP2). Note that we outperform PSPNet at the pixel level for Pascal VOC.

Runtime efficiency: In Tab. 4, we show the inference run-time for respectively the baselines, DQN and MCTS. The run-time scales *linearly* with the number of nodes and does not even depend on the potential type/order in case of DQN, as inference is reduced to a forward pass of the policy network at every iteration (Fig. 2). DQN is faster than all the solvers apart from α -exp. However, performance-wise, α -expansion has worse results (Tab. 2). MCTS is slower as it performs multiple simulations per node and requires computation of the reward at every step.

Learned Policies: In Fig. 5, we show the probability map across consecutive time steps. The selected nodes are colored in white. The darker the superpixel, the smaller the probability of selecting it next. We found that the heuristic learns a notion of smoothness, choosing nodes that are in close proximity and of the same label as the selected ones. Also, the policy learns to start labeling the nodes with low unary distribution entropy first, then decides on the ones with higher entropy.

Limitations: Our method is based on super-pixels, hence datasets with small objects require a large number of nodes and a longer run-time (MOTS vs. Pascal VOC). Also, our method is sensitive to bounding box class errors, as illustrated in Fig. 6 (first example), and to the parameters calibra-

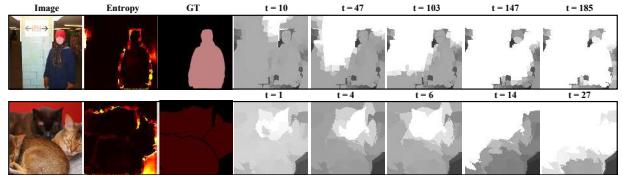


Figure 5: Visualization of the learned policy.

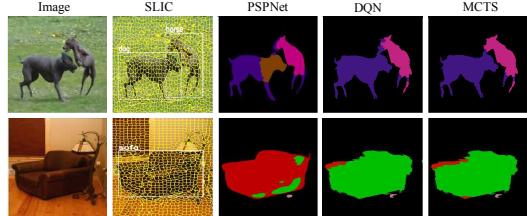


Figure 6: Failure cases.

tion of the energy function, as shown in the second example of the same figure. We plan to address the latter concern in future work via end-to-end training. Furthermore, little is known about deep reinforcement learning convergence. Nevertheless, it has been successfully applied to solve combinatorial programs by leveraging the structure in the data. We show that in our case as well, it converges to reasonable policies.

5. Conclusion

We study how to solve higher order CRF inference for semantic segmentation with reinforcement learning. The approach is able to deal with potentials that are too expensive to optimize using conventional techniques and outperforms traditional approaches while being more efficient. Hence, the proposed approach offers more flexibility for energy functions while scaling linearly with the number of nodes and the potential order. To answer our question: can we learn heuristics for graphical model inference? We think we can but we also want to note that a lot of manual work is required to find suitable features and graph structures. For this reason we think more research is needed to truly automate learning of heuristics for graphical model inference. We hope the research community will join us in this quest.

Acknowledgements: This work is supported in part by NSF under Grant No. 1718221 and MRI #1725729, UIUC, Samsung, 3M, and Cisco Systems Inc. (Gift Award CG 1377144). We thank Cisco for access to the Arcetri cluster and Iou-Jen Liu for initial discussions.

References

[1] B. Andres, J. H. Kappes, T. Beier, U. Köthe, and F. A. Hamprecht. The lazy flipper: Efficient depth-limited exhaustive search in discrete graphical models. In *ECCV*, 2012. 7

[2] M. Andrychowicz, M. Denil, S. Gomez, M. W. Hoffman, D. Pfau, T. Schaul, B. Shillingford, and N. De Freitas. Learning to learn by gradient descent by gradient descent. In *Proc. NeurIPS*, 2016. 3

[3] A. Arnab, S. Jayasumana, S. Zheng, and P. Torr. Higher order conditional random fields in deep neural networks. In *Proc. ECCV*, 2016. 6

[4] D. Bahdanau, P. Brakel, K. Xu, A. Goyal, R. Lowe, J. Pineau, A. Courville, and Y. Bengio. An actor-critic algorithm for sequence prediction. In *Proc. ICLR*, 2017. 3

[5] I. Bello, H. Pham, Q. V. Le, M. Norouzi, and S. Bengio. Neural Combinatorial Optimization with Reinforcement Learning. In <https://arxiv.org/abs/1611.09940>, 2016. 1, 2

[6] J. Boyan and A. W. Moore. Learning evaluation functions to improve optimization by local search. *JMLR*, 2000. 1, 2

[7] Y. Boykov and M. P. Jolly. Interactive graph cuts for optimal boundary and region segmentation of objects in n-d images. In *Proc. ICCV*, 2001. 2

[8] Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. In *Proc. EMMCVPR*, 2001. 2

[9] Y. Boykov, O. Veksler, and R. Zabih. Markov Random Fields with Efficient Approximations. In *Proc. CVPR*, 1998. 2

[10] Y. Boykov, O. Veksler, and R. Zabih. Fast Approximate Energy Minimization via Graph Cuts. *PAMI*, 2001. 2

[11] C. Buck, J. Bulian, M. Ciaramita, W. Gajewski, A. Gesmundo, N. Houlsby, and W. Wang. Ask the right questions: Active question reformulation with reinforcement learning. In *Proc. ICLR*, 2018. 3

[12] R. Bunel, M. Hausknecht, J. Devlin, R. Singh, and P. Kohli. Leveraging grammar and reinforcement learning for neural program synthesis. In *Proc. ICLR*, 2018. 3

[13] C. Chekuri, S. Khanna, J. Naor, and L. Zosin. Approximation algorithms for the metric labeling problem via a new linear programming formulation. In *Proc. SODA*, 2001. 2

[14] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic Image Segmentation with Deep Convolutional Nets and Fully Connected CRFs. In *Proc. ICLR*, 2015. 2

[15] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 2

[16] L. C. Chen, A. G. Schwing, A. Yuille, and R. Urtasun. Learning Deep Structured Models. In *Proc. ICML*, 2015. * equal contribution. 2

[17] H. Dai, B. Dai, and L. Song. Discriminative embeddings of latent variable models for structured data. In *Proc. ICML*, 2016. 2, 4

[18] H. Dai, E. B. Khalil, Y. Zhang, B. Dilkina, and L. Song. Learning Combinatorial Optimization Algorithms over Graphs. In *Proc. NeurIPS*, 2017. 1, 2

[19] M. Everingham, L. van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes (VOC) Challenge. *IJCV*, 2010. 2, 7

[20] A. Fix, A. Gruber, E. Boros, and R. Zabih. A graph cut algorithm for higher-order markov random fields. In *ICCV*, 2011. 7

[21] L. R. Ford and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, 1962. 2

[22] A. Globerson and T. Jaakkola. Fixing max-product: Convergent message passing algorithms for MAP LP-relaxations. In *Proc. NeurIPS*, 2007. 2

[23] A. Goldberg and R. Tarjan. A new approach to the maximum flow problem. *JACM*, 1988. 2

[24] C. Graber, O. Meshi, and A. G. Schwing. Deep structured prediction with nonlinear output transformations. In *NeurIPS*, 2018. 2

[25] C. Graber and A. G. Schwing. Graph structured prediction energy networks. In *NeurIPS*, 2019. 2

[26] D. Greig, B. Porteous, and A. Seheult. Exact maximum a posteriori estimation for binary images. *J. of the Royal Statistical Society*, 1989. 2

[27] S. Gu, T. Lillicrap, Z. Ghahramani, R. E. Turner, and S. Levine. Q-Prop: Sample-Efficient Policy Gradient with An Off-Policy Critic. In *Proc. ICLR*, 2017. 1, 2

[28] A. Guiasti, D. Ciresan, J. Masci, L. Gambardella, and J. Schmidhuber. Fast image scanning with deep max-pooling convolutional neural networks. In *Proc. ICIP*, 2013. 2

[29] B. Hariharan, P. Arbelaez, L. Bourdev, S. Maji, and J. Malik. Semantic Contours from Inverse Detectors. In *Proc. ICCV*, 2011. 7

[30] B. Hariharan, P. Arbelaez, R. Girshick, and J. Malik. Hypercolumns for object segmentation and fine-grained localization. In *Proc. CVPR*, 2015. 2, 4

[31] D. He, Y. Xia, T. Qin, L. Wang, N. Yu, T.-Y. Liu, and W.-Y. Ma. Dual learning for machine translation. In *Proc. NeurIPS*, 2016. 3

[32] H. He, H. Daume, and J. M. Eisner. Learning to search in branch and bound algorithms. In *Proc. NeurIPS*, 2014. 3

[33] X. He, R. S. Zemel, and M. Á. Carreira-Perpiñán. Multiscale Conditional Random Fields for Image Labeling. In *Proc. CVPR*, 2004. 2

[34] J. K. Johnson. *Convex relaxation methods for graphical models: Lagrangian and maximum entropy approaches*. PhD thesis, MIT, 2008. 2

[35] V. Jojic, S. Gould, and D. Koller. Accelerated dual decomposition for MAP inference. In *Proc. ICML*, 2010. 2

[36] J. H. Kappes, B. Savchynskyy, and C. Schnörr. A Bundle Approach To Efficient MAP-Inference by Lagrangian Relaxation. In *Proc. CVPR*, 2012. 2

[37] J. H. Kappes, B. Savchynskyy, and C. Schnörr. A bundle approach to efficient map-inference by lagrangian relaxation. In *CVPR*, 2012. 7

[38] E. B. Khalil, P. Le Bodic, L. Song, G. L. Nemhauser, and B. N. Dilkina. Learning to branch in mixed integer programming. In *Proc. AAAI*, 2016. 3

[39] V. Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *PAMI*, 2006. 2

[40] V. Kolmogorov and R. Zabih. What Energy Functions Can Be Minimized via Graph Cuts? *PAMI*, 2004. 2

[41] V. N. Kolmogorov and M. J. Wainwright. On the optimality of tree-reweighted max-product message-passing. In *Proc. UAI*, 2005. 2

[42] S. Konishi and A. L. Yuille. Statistical cues for domain specific image segmentation with performance analysis. In *Proc. CVPR*, 2000. 2

[43] J. Lafferty, A. McCallum, and F. Pereira. Conditional Random Fields: Probabilistic Models for segmenting and labeling sequence data. In *Proc. ICML*, 2001. 2

[44] M. G. Lagoudakis and M. L. Littman. Learning to select branching rules in the dpll procedure for satisfiability. *ENDM*, 2001. 3

[45] A. Laterre, Y. Fu, M. K. Jabri, A.-S. Cohen, D. Kas, K. Hajjar, T. S. Dahl, A. Kerkeni, and K. Beguir. Ranked reward: Enabling self-play reinforcement learning for combinatorial optimization. In *Proc. Deep RL Workshop NeurIPS*, 2018. 2

[46] J. Li, W. Monroe, A. Ritter, M. Galley, J. Gao, and D. Jurafsky. Deep reinforcement learning for dialogue generation. In *Proc. EMNLP*, 2016. 3

[47] K. Li and J. Malik. Learning to Optimize. In *Proc. ICLR*, 2017. 3

[48] X. Li, Z. Liu, P. Luo, C. Change, and X. Tang. Not all pixels are equal: Difficulty-aware semantic segmentation via deep layer cascade. In *CVPR*, 2017. 2

[49] C. Liang, J. Berant, Q. Le, K.D. Forbus, and N. Lao. Neural symbolic machines: Learning semantic parsers on freebase with weak supervision. In *Proc. ACL*, 2016. 3

[50] C. Liang, M. Norouzi, J. Berant, Q. Le, and N. Lao. Memory augmented policy optimization for program synthesis with generalization. In *Proc. NeurIPS*, 2017. 3

[51] Z. Liu, X. Li, P. Luo, C. C. Loy, and X. Tang. Semantic image segmentation via deep parsing network. In *ICCV*, 2015. 2

[52] J. Long, E. Shelhamer, and T. Darrell. Fully Convolutional Networks for Semantic Segmentation. In *Proc. CVPR*, 2015. 2

[53] A. F. T. Martins, M. A. T. Figueiredo, P. M. Q. Aguiar, N. A. Smith, and E. P. Xing. An Augmented Lagrangian Approach to Constrained MAP Inference. In *Proc. ICML*, 2011. 2

[54] O. Meshi and A. Globerson. An Alternating Direction Method for Dual MAP LP Relaxation. In *Proc. ECML PKDD*, 2011. 2

[55] O. Meshi, M. Mahdavi, and A. Schwing. Smooth and Strong: MAP Inference with Linear Convergence. In *Proc. NIPS*, 2015. 2

[56] O. Meshi and A. G. Schwing. Asynchronous Parallel Coordinate Minimization for MAP Inference. In *Proc. NIPS*, 2017. 2

[57] S. Messaoud, D. Forsyth, and A. Schwing. Structural consistency and controllability for diverse colorization. In *ECCV*, 2018. 2

[58] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 2015. 1

[59] K. Narasimhan, A. Yala, and R. Barzilay. Improving information extraction by acquiring external evidence with reinforcement learning. In *Proc. EMNLP*, 2016. 3

[60] R. Nogueira and K. Cho. Task-oriented query reformulation with reinforcement learning. In *Proc. EMNLP*, 2017. 3

[61] H. Noh, S. Hong, and B. Han. Learning deconvolution network for semantic segmentation. In *ICCV*, 2015. 2

[62] M. Norouzi, S. Bengio, N. Jaitly, M. Schuster, Y. Wu, D. Schuurmans, et al. Reward augmented maximum likelihood for neural structured prediction. In *Proc. NeurIPS*, 2016. 3

[63] R. Paulus, C. Xiong, and R. Socher. A deep reinforced model for abstractive summarization. In *Proc. ICLR*, 2018. 3

[64] J. Pearl. Reverend bayes on inference engines: a distributed hierarchical approach. In *Proc. AAAI*, 1982. 7

[65] T. Pierrot, G. Ligner, S. E. Reed, O. Sigaud, N. Perrin, A. Laterre, D. Kas, K. Beguir, and N. Freitas. Learning compositional neural programs with recursive tree search and planning. *NeurIPS*, 2019. 3

[66] P. Qin, W. Xu, and W. Y. Wang. Robust distant supervision relation extraction via deep reinforcement learning. In *Proc. ACL*, 2018. 3

[67] M. Ranzato, S. Chopra, M. Auli, and W. Zaremba. Sequence level training with recurrent neural networks. In *Proc. ICLR*, 2016. 3

[68] P. Ravikumar, A. Agarwal, and M. J. Wainwright. Message-passing for graph-structured linear programs: Proximal methods and rounding schemes. *JMLR*, 2010. 2

[69] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *Proc. CVPR*, 2016. 2, 6, 7

[70] S.J. Rennie, E. Marcheret, Y. Mroueh, J. Ross, and V. Goel. Self-critical sequence training for image captioning. In *Proc. CVPR*, 2017. 3

[71] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 2015. 2

[72] C. Rother, V. Kolmogorov, and A. Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. In *Proc. TOG*, 2004. 6

[73] H. Samulowitz and R. Memisevic. Learning to solve QBF. In *Proc. AAAI*, 2007. 3

[74] M. I. Schlesinger. Sintaksicheskiy analiz dvumernykh zritel'nikh signalov v usloviyakh pomekh (Syntactic analysis of two-dimensional visual signals in noisy conditions). *Kibernetika*, 1976. 2

[75] A. Schwing, T. Hazan, M. Pollefeys, and R. Urtasun. Distributed Message Passing for Large Scale Graphical Models. In *Proc. CVPR*, 2011. 2

[76] A. G. Schwing, T. Hazan, M. Pollefeys, and R. Urtasun. Globally Convergent Dual MAP LP Relaxation Solvers using Fenchel-Young Margins. In *Proc. NeurIPS*, 2012. 2

[77] A. G. Schwing, T. Hazan, M. Pollefeys, and R. Urtasun. Globally Convergent Parallel MAP LP Relaxation Solver using the Frank-Wolfe Algorithm. In *Proc. ICML*, 2014. 2

[78] A. G. Schwing and R. Urtasun. Fully Connected Deep Structured Networks. In <https://arxiv.org/abs/1503.02351>, 2015. 2

[79] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun. OverFeat: Integrated Recognition, Localization and Detection using Convolutional Networks. In *Proc. ICLR*, 2014. 2

[80] A. Sharaf and H. Daumé III. Structured prediction via learning to search under bandit feedback. In *Proc. Workshop on Structured Prediction for NLP ACL*, 2017. 3

[81] J. Shotton, J. Winn, C. Rother, and A. Criminisi. TextronBoost: Joint Appearance, Shape and Context Modeling for Multi-Class Object Recognition and Segmentation. In *Proc. ECCV*, 2006. 2

[82] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 2016. 1, 5

[83] D. Sontag, T. Meltzer, A. Globerson, and T. Jaakkola. Tightening LP Relaxations for MAP using Message Passing. In *Proc. NeurIPS*, 2008. 2

[84] SLIC Superpixels Compared to State-of-the Art Superpixel Methods. Slic superpixels. *TPAMI*, 2012. 7

[85] O. Vinyals, M. Fortunato, and N. Jaitly. Pointer networks. In *Proc. NeurIPS*, 2015. 1, 2

[86] P. Voigtlaender, M. Krause, A. Osep, J. Luiten, B. B. G. Sekar, A. Geiger, and B. Leibe. Mots: Multi-object tracking and segmentation. In *CVPR*, 2019. 2, 7

[87] M. J. Wainwright, T. S. Jaakkola, and A. S. Willsky. Tree-reweighted belief propagation algorithms and approximate ml estimation by pseudo-moment matching. In *AISTATS*, 2003. 7

[88] T. Werner. A Linear Programming Approach to Max-sum Problem: A Review. *PAMI*, 2007. 2

[89] T. Werner. Revisiting the linear programming relaxation approach to Gibbs energy minimization and weighted constraint satisfaction. *PAMI*, 2010. 2

[90] J.D. Williams, K. Asadi, and G. Zweig. Hybrid code networks: practical and efficient end-to-end dialog control with supervised and reinforcement learning. In *Proc. ACL*, 2017. 3

[91] W. Xiong, T. Hoang, and W. Y. Wang. Deeppath: A reinforcement learning method for knowledge graph reasoning. In *Proc. EMNLP*, 2017. 3

[92] F. Yu and V. Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015. 2

[93] W. Zhang and T. G. Dietterich. Solving combinatorial optimization tasks by reinforcement learning: A general methodology applied to resource-constrained scheduling. *JAIR*, 2000. 1, 2

[94] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia. Pyramid scene parsing network. In *CVPR*, 2017. 2, 7

[95] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. Torr. Conditional random fields as recurrent neural networks. In *ICCV*, 2015. 2

[96] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. H. S. Torr. Conditional Random Fields as Recurrent Neural Networks. In *Proc. ICCV*, 2015. 2

[97] B. Zoph and Q. V. Le. Neural architecture search with reinforcement learning. In *Proc. ICLR*, 2017. 3