Structured learning of tree potentials in CRF for image segmentation

Liu, Fayao; Lin, Guosheng; Qiao, Ruizhi; Shen, Chunhua

2017

Liu, F., Lin, G., Qiao, R., & Shen, C. (2018). Structured learning of tree potentials in CRF for image segmentation. *IEEE Transactions on Neural Networks and Learning Systems*, 29(6), 2631-2637. doi:10.1109/TNNLS.2017.2690453

<https://hdl.handle.net/10356/107576>

<https://doi.org/10.1109/TNNLS.2017.2690453>

© 2017 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at: <https://doi.org/10.1109/TNNLS.2017.2690453>.

Downloaded on 26 Apr 2023 07:11:56 SGT

Structured Learning of Tree Potentials in CRF for Image Segmentation

Fayao Liu, Guosheng Lin, Ruizhi Qiao, Chunhua Shen

Abstract—We propose a new approach to image segmentation, which exploits the advantages of both conditional random fields (CRFs) and decision trees. In the literature, the potential functions of CRFs are mostly defined as a *linear* combination of some pre-defined parametric models, and then methods like structured support vector machines (SSVMs) are applied to learn those linear coefficients. We instead formulate the unary and pairwise potentials as nonparametric forests—ensembles of decision trees, and learn the ensemble parameters and the trees in a unified optimization problem within the large-margin framework. In this fashion, we easily achieve *nonlinear* learning of potential functions on both unary and pairwise terms in CRFs. Moreover, we learn class-wise decision trees for each object that appears in the image. Due to the rich structure and flexibility of decision trees, our approach is powerful in modelling complex data likelihoods and label relationships. The resulting optimization problem is very challenging because it can have exponentially many variables and constraints. We show that this challenging optimization can be efficiently solved by combining a modified column generation and cutting-planes techniques. Experimental results on both binary (Graz-02, Weizmann horse, Oxford flower) and multi-class (MSRC-21, PASCAL VOC 2012) segmentation datasets demonstrate the power of the learned nonlinear nonparametric potentials.

Index Terms—Conditional random fields, Decision trees, Structured support vector machines, Image segmentation.

CONTENTS

I	Introduction	2
II	Learning tree potentials in CRFs	3
II-A	Segmentation using CRFs models . . .	3
II-B	Energy Formulation	3
II-C	Learning CRFs in the large-margin framework	3
II-D	Learning tree potentials using column generation	4
II-E	Speeding up optimization using cutting-plane	5
III	Experiments	6
III-A	Experimental setup	6
III-B	Comparing with baseline methods . . .	7
III-C	Comparing with state-of-the-art methods	7
IV	Conclusion	9

F. Liu, R. Qiao, C. Shen are with The University of Adelaide, Australia. G. Lin is with Nanyang Technological University, Singapore. This work was done when G. Lin was with The University of Adelaide. Email: {fayao.liu, ruizhi.qiao, chunhua.shen}@adelaide.edu.au, guosheng.lin@gmail.com

Appearing in *IEEE Transactions on Neural Networks and Learning Systems*, 26 March 2017.

I. INTRODUCTION

The goal of object segmentation is to produce a pixel level segmentation of different object categories. It is challenging as the objects may appear in various backgrounds and in different visual conditions. CRFs [19] model the conditional distribution of labels given observations, and represents the state-of-the-art in image/object segmentation [10], [15], [24], [27], [32], [34]. The max-margin principle has also been applied to predict structured outputs, including SSVMs [36], and max-margin Markov networks [35]. These three methods share similarities when viewed as optimization problems using different loss functions. Szummer *et al.* [34] proposed to learn linear coefficients of CRFs potentials using SSVMs and graph cuts. To date, most of these methods assume a pre-defined parametric model for the potential functions, and typically only the linear coefficients of the parametric model are learned. This can greatly limit the flexibility of the model capability of CRFs, and thus calls for effective methods to incorporate nonlinear nonparametric models for learning the potential functions in CRFs.

As similar in standard support vector machines (SVMs), nonlinearity can be achieved by introducing nonlinear kernels for SSVMs. However, the time complexity of nonlinear SVMs is roughly $O(n^{3.5})$ with n being the number of training data examples. This time complexity is problematic for SSVMs, where the number of constraints grows exponentially in the description length of the label y . Moreover, nonlinear functions can significantly slow down the test time in most cases. Because of these reasons, currently most SSVMs applications use linear kernels (or linear parametric potential functions in CRFs), despite the fact that nonlinear functions usually deliver more promising prediction accuracy. In this work, we address this issue by combining CRFs with nonparametric decision trees. Both CRFs and decision trees have gained tremendous success in computer vision. Decision trees are capable of modelling complex relations and generalize well on test data. Unlike kernel methods, decision trees are fast to evaluate and can be used to select informative features.

In this work, we propose to use ensembles of decision trees to map the image content to both the unary terms and the pairwise interaction values in CRFs. The proposed method is termed as CRFTree. Specifically, we formulate both the unary and pairwise potentials as *nonparametric* forests—ensembles of decision trees, and learn the ensemble parameters and the trees in a single optimization framework. In this way, the nonlinearity is easily introduced into CRFs learning without confronting the kernel dilemma. Furthermore, we learn class-wise decision trees for each object. Due to the rich structure and flexibility of decision trees, our approach is powerful in modelling complex data likelihoods and label relationships. The resulting optimization problem is very challenging in the sense that it can involve exponentially or even infinitely many variables and constraints. We summarize our main contributions as follows.

1. We formulate the unary and pairwise potentials as ensembles of decision trees, and show how to jointly learn the ensemble parameters and the trees as a unified

optimization problem within the large-margin framework. In this fashion, we achieve nonlinear potential learning on both the unary and pairwise terms.

2. We learn class-wise decision trees (potentials) for each object that appears in the image.
3. We show how to train the proposed CRFTree model efficiently. In particular, we combine the column generation and cutting-planes techniques to approximately solve the resulting optimization problem, which can involve exponentially many variables and constraints.
4. We empirically demonstrate that CRFTree outperforms existing methods for image segmentation. On both binary and multi-class segmentation datasets we show the advantages of the learned nonlinear nonparametric potentials of decision trees.

Related work We briefly review the recent works that are relevant to ours. A few attempts have been made to apply nonlinear kernels in SSVMs. Yu *et al.* [37] and Severyn *et al.* [29] developed sampled cuts based methods for training SSVMs with kernels. Sampled cuts methods were originally proposed for standard kernel SVMs. When applied to SSVMs, the performance is compromised [24]. In [5], the image-mask pair kernels are designed to exploit image-level structural information for object segmentation. However, these kernels are restricted to the unary term. Although not in the large margin framework, the kernel CRFs proposed in [20] incorporates kernels into the CRFs learning. The authors only demonstrated the efficacy of their method on a synthetic and a small scale protein dataset. To sum up, these approaches are hampered by the heavy computation complexity. Furthermore, it is not a trivial task to design appropriate kernels for structured problems. Recently, Lucchi *et al.* [24] proposed a two-step solution to tackle this problem. Specifically, they train linear SSVMs by using kernelized feature vectors that are obtained from training a standard non-linear kernel SVMs model. They experimentally demonstrate that the kernel transferred linear SVMs model achieves similar performance as the Gaussian SVMs. However, this approach is heuristic and it cannot be shown theoretically that their formulation approximates a nonlinear SSVMs model. Besides, their method consumes extra usage of memory and training time since the dimension of the transformed features equals to the number of support vectors, while the latter is linearly proportional to the size of the training data [33]. Moreover, compared to the above mentioned works of [5] and [24], we achieve nonlinear learning on both the unary and the pairwise terms while theirs are limited to nonlinear unary potential learning. The recent work of Shen *et al.* [31] generalizes standard boosting methods to structured learning, which shares similarities to our work here. However, our method bears critical differences from theirs: 1) We design a column generation method for non-linear tree potentials learning in CRFs directly from the SSVMs formulation. Different from the case in [31], which can directly derive column generation method analogous to LPBoost [8], our derivation here is more challenging. This is because we can not obtain the most violated constraint from the constraints of the dual problem, on which the column generation technique relies. We instead inspect the KKT condition to seek for

the most violated constraint. This is an important difference compared to existing column generation techniques. 2) We develop a CRFs learning method for multi-class semantic segmentation, while [31] only shows CRFs learning for binary foreground/background segmentation. Our experiments on the MSRC-21 dataset shows that our method achieves state-of-the-art results. 3) We learn class-wise decision trees (potentials) for each object that appears in the image. This is different from [31]. The work of decision tree fields [28] is close to ours in that they also use decision trees to model the pairwise potentials. The major difference is that in [28] potential functions are constructed by directly summing the energy tables associated with the set of nodes taken during evaluating the decision trees. Their trees are generally deep, with depth 15 for the unary potential and 6 for the pairwise potential in their experiment. By contrast, we model the potential functions as an ensemble of decision trees and learn them in the large margin framework. In our method, the decision trees are shallow and simple with binary outputs.

II. LEARNING TREE POTENTIALS IN CRFS

We present the details of our method in this section by first introducing the CRFs models for segmentation, then formulating the energy functions and showing how to learn decision tree potentials in the large-margin framework.

A. Segmentation using CRFs models

Before presenting our method, we first revisit how to use CRFs models to perform image segmentation. Given an image instance $\mathbf{x}$ and its corresponding labelling $\mathbf{y}$, CRFs [19] models the conditional distribution of the form

$$P(\mathbf{y}|\mathbf{x}; \mathbf{w}) = \frac{1}{Z} \exp(-E(\mathbf{y}, \mathbf{x}; \mathbf{w})). \quad (1)$$

where $\mathbf{w}$ are parameters and Z is the normalization term. The energy E of an image $\mathbf{x}$ with segmentation labels $\mathbf{y}$ over the nodes (superpixels) $\mathcal{N}$ and edges $\mathcal{S}$, takes the following form:

$$E(\mathbf{y}, \mathbf{x}; \mathbf{w}) = \sum_{p \in \mathcal{N}} \Phi^{(1)}(y^p, \mathbf{x}; \mathbf{w}) + \sum_{(p,q) \in \mathcal{S}} \Phi^{(2)}(y^p, y^q, \mathbf{x}; \mathbf{w}). \quad (2)$$

Here $\mathbf{x} \in \mathcal{X}$, $\mathbf{y} \in \mathcal{Y}$; $\Phi^{(1)}$ and $\Phi^{(2)}$ are the unary and pairwise potentials, both of which depend on the observations as well as the parameter $\mathbf{w}$. CRFs seeks an optimal labeling that achieves maximum a posterior (MAP), which mainly involves a two-step process [34]: 1) Learning the model parameters from the training data; 2) Inferring a most likely label for the test data given the learned parameters. The segmentation problem thus reduces to minimizing the energy (or cost) over $\mathbf{y}$ by the learned parameters $\mathbf{w}$, which is $\mathbf{y}^* = \operatorname{argmin}_{\mathbf{y} \in \mathcal{Y}} E(\mathbf{y}, \mathbf{x}; \mathbf{w})$. When the energy function is submodular, this inference problem can be efficiently solved via graph cuts [34].

B. Energy Formulation

Given the energy function in Eqn. (2), we show how to construct the unary and pairwise potentials using decision trees. We denote $\mathbf{x}^p$ as the features of superpixel p ($p = 1, \dots, n$),

with its label $y^p \in \{1, \dots, K\}$, where K is the number of classes. Let $\mathcal{H}$ be a set of decision trees, which can be infinite. Each $h_j^{(1)}(\cdot) \in \mathcal{H}$ takes $\mathbf{x}^p$ as the input, and $h_j^{(2)}(\cdot, \cdot) \in \mathcal{H}$ takes a pair $(\mathbf{x}^p, \mathbf{x}^q)$ as the input to output $\{0, 1\}$. We introduce $(K+1)$ groups of decision trees, in which K groups are for the unary potential and one group for the pairwise potential. For the unary potential, the K groups of decision trees are denoted by $\mathbf{H}_c^{(1)} (c = 1, \dots, K)$, which correspond to K categories. Each $\mathbf{H}_c^{(1)}$ is associated with the c -th class. In other words, for each class, we maintain its own unary feature mappings. Each group of decision trees for the unary potential can be written as: $\mathbf{H}_c^{(1)} = [\bar{h}_{c1}^{(1)}, \bar{h}_{c2}^{(1)}, \dots]^\top$, which are the output of decision trees: $\bar{h}_{cj}^{(1)}$. All decision trees of the unary potential are denoted by $\mathbf{H}^{(1)} = [\mathbf{H}_1^{(1)}, \mathbf{H}_2^{(1)}, \dots, \mathbf{H}_K^{(1)}]$. Accordingly, for the pairwise potential, the group of decision trees is denoted by $\mathbf{H}^{(2)}$, and $\mathbf{H}^{(2)} = [\bar{h}_1^{(2)}, \bar{h}_2^{(2)}, \dots]^\top$ being the output of all $\bar{h}_j^{(2)}$. The whole set of decision trees is denoted by $\mathbf{H} = [\mathbf{H}^{(1)}, \mathbf{H}^{(2)}]$. We then construct the unary and pairwise potentials as

$$\Phi^{(1)}(y^{(p)}, \mathbf{x}) = \mathbf{w}_{y^p}^{(1)\top} \mathbf{H}_{y^p}^{(1)}(\mathbf{x}^p). \quad (3)$$

$$\Phi^{(2)}(y^{(p)}, y^{(q)}, \mathbf{x}) = \mathbf{w}^{(2)\top} \mathbf{H}^{(2)}(\mathbf{x}^p, \mathbf{x}^q) I(y^p \neq y^q). \quad (4)$$

where $I(\cdot)$ is an indicator function which equals 1 if the input is true and 0 otherwise. Then the energy function in Eqn. (2) can be written as:

$$E(\mathbf{y}, \mathbf{x}; \mathbf{w}, \mathbf{H}) = \sum_{p \in \mathcal{N}} \mathbf{w}_{y^p}^{(1)\top} \mathbf{H}_{y^p}^{(1)}(\mathbf{x}^p) + \sum_{(p,q) \in \mathcal{S}} \mathbf{w}^{(2)\top} \mathbf{H}^{(2)}(\mathbf{x}^p, \mathbf{x}^q) I(y^p \neq y^q). \quad (5)$$

Next we show how to learn these decision tree potentials in the large-margin framework.

C. Learning CRFs in the large-margin framework

Instead of directly minimizing the negative log-likelihood loss, we here learn the CRFs parameters in the large margin framework, similar to [34]. Given a set of training examples $\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^m$, the large-margin based CRFs learning solves the following optimization:

$$\begin{aligned} \min_{\mathbf{w}, \xi \geq 0} \quad & \frac{1}{2} \|\mathbf{w}\|_2^2 + \frac{C}{m} \sum_i \xi_i \\ \text{s.t.:} \quad & E(\mathbf{y}, \mathbf{x}_i; \mathbf{w}, \mathbf{H}) - E(\mathbf{y}_i, \mathbf{x}_i; \mathbf{w}, \mathbf{H}) \geq \Delta(\mathbf{y}_i, \mathbf{y}) - \xi_i, \\ & \forall i = 1, \dots, m, \text{ and } \forall \mathbf{y} \in \mathcal{Y}; \end{aligned} \quad (6)$$

where $\Delta : \mathcal{Y} \times \mathcal{Y} \mapsto \mathbb{R}$ is a loss function associated with the prediction and the true label mask. In general, we have $\Delta(\mathbf{y}, \mathbf{y}) = 0$ and $\Delta(\mathbf{y}, \mathbf{y}') > 0$ for any $\mathbf{y}' \neq \mathbf{y}$. Intuitively, the optimization in Eqn. (6) is to encourage the energy of the ground truth label $E(\mathbf{y}_i, \mathbf{x}_i; \mathbf{w})$ to be lower than that of any other *incorrect* labels $E(\mathbf{y}, \mathbf{x}_i; \mathbf{w})$ by at least a margin $\Delta(\mathbf{y}_i, \mathbf{y})$.

To learn the potential functions we proposed in §II-B in the large-margin framework, we introduce the following

definitions. For the unary part, we define $\mathbf{w}^{(1)} = \mathbf{w}_1^{(1)} \odot \mathbf{w}_2^{(1)} \odot \dots \odot \mathbf{w}_K^{(1)}$, where $\odot$ stacks two vectors, and

$$\Psi^{(1)}(\mathbf{y}, \mathbf{x}; \mathbf{H}^{(1)}) = \sum_{p \in \mathcal{N}} \mathbf{H}_{y^p}^{(1)}(\mathbf{x}^p) \otimes y^p. \quad (7)$$

where $\otimes$ denotes the tensor operation (e.g., $\mathbf{x}^p \otimes y^p = [I(y^p = 1)\mathbf{x}^{p^\top}, \dots, I(y^p = K)\mathbf{x}^{p^\top}]^\top$). Recall that $\mathbf{x}^p$ denotes the p -th superpixel of the image $\mathbf{x}$. Here, $\Psi^{(1)}$ acts as the unary feature mapping. Clearly we have:

$$\mathbf{w}^{(1)\top} \Psi^{(1)}(\mathbf{y}, \mathbf{x}; \mathbf{H}^{(1)}) = \sum_{p \in \mathcal{N}} \Phi^{(1)}(y^p, \mathbf{x}). \quad (8)$$

For the pairwise part, we define the pairwise feature mapping as:

$$\Psi^{(2)}(\mathbf{y}, \mathbf{x}; \mathbf{H}^{(2)}) = \sum_{(p,q) \in \mathcal{S}} \mathbf{H}^{(2)}(\mathbf{x}^p, \mathbf{x}^q) I(y^p \neq y^q). \quad (9)$$

Then we have the following relation:

$$\mathbf{w}^{(2)\top} \Psi^{(2)}(\mathbf{y}, \mathbf{x}; \mathbf{H}^{(2)}) = \sum_{(p,q) \in \mathcal{S}} \Phi^{(2)}(y^p, y^q, \mathbf{x}). \quad (10)$$

We further define $\mathbf{w} = \mathbf{w}^{(1)} \odot \mathbf{w}^{(2)}$, and the joint feature mapping as

$$\Psi(\mathbf{y}, \mathbf{x}; \mathbf{H}) = \Psi^{(1)}(\mathbf{y}, \mathbf{x}; \mathbf{H}^{(1)}) \odot \Psi^{(2)}(\mathbf{y}, \mathbf{x}; \mathbf{H}^{(2)}). \quad (11)$$

With the definitions of $\mathbf{w}$ and Ψ , the energy function can then be written as:

$$\begin{aligned} E(\mathbf{y}, \mathbf{x}; \mathbf{w}, \mathbf{H}) &= \sum_{p \in \mathcal{N}} \Phi^{(1)}(y^p, \mathbf{x}; \mathbf{w}, \mathbf{H}^{(1)}) \\ &\quad + \sum_{(p,q) \in \mathcal{S}} \Phi^{(2)}(y^p, y^q, \mathbf{x}; \mathbf{w}, \mathbf{H}^{(2)}) \\ &= \mathbf{w}^\top \Psi(\mathbf{y}, \mathbf{x}; \mathbf{H}). \end{aligned} \quad (12)$$

Now we can apply the large-margin framework to learn CRFs using the proposed energy functions by rewriting the optimization problem in Eqn. (6) as:

$$\begin{aligned} \min_{\mathbf{w}, \xi} \quad & \frac{1}{2} \|\mathbf{w}\|_2^2 + \frac{C}{m} \sum_i \xi_i \\ \text{s.t.:} \quad & \mathbf{w}^\top [\Psi(\mathbf{y}, \mathbf{x}_i; \mathbf{H}) - \Psi(\mathbf{y}_i, \mathbf{x}_i; \mathbf{H})] \geq \Delta(\mathbf{y}_i, \mathbf{y}) - \xi_i, \\ & \forall i = 1, \dots, m, \text{ and } \forall \mathbf{y} \in \mathcal{Y}; \\ & \mathbf{w} \geq 0, \xi \geq 0. \end{aligned} \quad (13)$$

Note that we add the $\mathbf{w} \geq 0$ constraint to ensure submodular property of our energy functions, which we will discuss the details later in §II-E. Up until now, we are ready to learn $\mathbf{w}$ and Ψ (or $\mathbf{H}$) in a single optimization problem formulated in Eqn. (13), but it is not clear how. Next we demonstrate how to solve the optimization problem in Eqn. (13) by using column generation and cutting-plane.

D. Learning tree potentials using column generation

We aim to learn a set of decision trees $\mathbf{H}$ and the potential parameter $\mathbf{w}$ by solving the optimization problem in Eqn. (13). However, jointly learning $\mathbf{H}$ and $\mathbf{w}$ is generally difficult. Here we propose to apply column generation techniques [8], [30] to alternatively construct the set of decision trees and

solve for $\mathbf{w}$. From the point of view of column generation techniques, the dimension of the primal variable $\mathbf{w}$ is infinitely large; the column generation is to iteratively select (generate) variables for solving the optimization. In our case, infinitely many dimension of $\mathbf{w}$ corresponds to infinitely many decision trees, thus we iteratively generate decision trees to solve the optimization.

Basically, we construct a working set of decision trees (denoted as $\mathcal{W}_{\mathbf{H}}$). During each column generation iteration we perform two steps. In the first step, we generate new decision trees and add them to $\mathcal{W}_{\mathbf{H}}$. In the second step, we solve a restricted optimization problem in Eqn. (13) on the current working set $\mathcal{W}_{\mathbf{H}}$ to obtain the solution of $\mathbf{w}$. We repeat these two steps until convergence. Next we describe how to generate decision trees in a principal way by using the dual solution of the optimization in Eqn. (13), which is similar to the conventional column generation technique. First we derive the Lagrange dual problem of Eqn. (13), which can be written as

$$\begin{aligned} \max_{\lambda, \theta} \quad & \sum_{i, \mathbf{y}} \lambda_{(i, \mathbf{y})} \Delta(\mathbf{y}_i, \mathbf{y}) \\ & - \frac{1}{2} \left\{ \sum_{i, \mathbf{y}} \lambda_{(i, \mathbf{y})} [\Psi(\mathbf{y}, \mathbf{x}_i; \mathbf{H}) - \Psi(\mathbf{y}_i, \mathbf{x}_i; \mathbf{H})] + \theta \right\}^2 \quad (14) \\ \text{s.t.:} \quad & 0 \leq \sum_{\mathbf{y}} \lambda_{(i, \mathbf{y})} \leq \frac{C}{m}, \forall i = 1, \dots, m; \theta \geq 0, \lambda \geq 0. \end{aligned}$$

Here θ, λ are the dual variables. When using column generation technique, one need to find the most violated constraint in the dual. However, the constraints of the dual problem do not involve decision trees $\mathbf{H}$. Instead of examining the dual constraint, we inspect the KKT condition, which is an important difference compared to existing column generation techniques. According to the KKT condition, when at optimal, the following condition holds for the primal solution $\mathbf{w}$ and the current working set $\mathcal{W}_{\mathbf{H}}$:

$$\mathbf{w} \geq \sum_{i, \mathbf{y}} \lambda_{(i, \mathbf{y})} [\Psi(\mathbf{y}, \mathbf{x}_i; \mathbf{H}) - \Psi(\mathbf{y}_i, \mathbf{x}_i; \mathbf{H})]. \quad (15)$$

All of those generated $\mathbf{H} \in \mathcal{W}_{\mathbf{H}}$ satisfy the above condition. Obviously, generating new decision trees which most violate the above condition would contribute the most to the optimization of Eqn. (13). Hence the strategy of generating new decision trees is to solve the following problem:

$$\mathbf{H}^* = \operatorname{argmax}_{\mathbf{H}} \sum_{i, \mathbf{y}} \lambda_{(i, \mathbf{y})} [\Psi(\mathbf{y}, \mathbf{x}_i; \mathbf{H}) - \Psi(\mathbf{y}_i, \mathbf{x}_i; \mathbf{H})]. \quad (16)$$

Then $\mathbf{H}^*$ is added to the current working set $\mathcal{W}_{\mathbf{H}}$. If $\mathbf{H}^*$ still satisfies the condition in Eqn. (15), the current solution of $\mathbf{H}$ and $\mathbf{w}$ is already the globally optimal one.

The optimization in Eqn. (16) for generating new decision trees can be independently decomposed into solving the unary part and the pairwise part. Hence $\mathbf{H}^*$ can be written as: $\mathbf{H}^* = [\mathbf{H}^{(1)*}, \mathbf{H}^{(2)*}]$. For the unary part, we learn class-wise decision trees, namely, we generate K decision trees corresponding to K categories at each column generation iteration. Hence $\mathbf{H}^{(1)*}$ is composed of K decision trees: $\mathbf{H}^{(1)*} = [\mathbf{h}_1^{(1)*}, \dots, \mathbf{h}_K^{(1)*}]$. More specifically, according to the definition of $\Psi(\mathbf{y}, \mathbf{x})$ in

Eqn. (11), we solve the following K problems:

$$\forall c=1, \dots, K :$$

$$\begin{aligned} \hat{h}_c^{(1)*}(\cdot) &= \operatorname{argmax}_{\hat{h} \in \mathcal{H}} \sum_{i, \mathbf{y}} \lambda_{(i, \mathbf{y})} \left[\sum_{\substack{p \in \mathcal{N}, \\ y^p = c}} \hat{h}_{y^p}^{(1)}(\mathbf{x}_i^p) - \sum_{\substack{p \in \mathcal{N}, \\ y^p \neq c}} \hat{h}_{y_i^p}^{(1)}(\mathbf{x}_i^p) \right] \\ &= \operatorname{argmax}_{\hat{h} \in \mathcal{H}} \sum_{i, \mathbf{y}} \left[\sum_{\substack{p \in \mathcal{N}, \\ y^p = c}} \underbrace{\lambda_{(i, \mathbf{y})} \hat{h}_{y^p}^{(1)}(\mathbf{x}_i^p)}_{\text{positive}} \right. \\ &\quad \left. - \sum_{\substack{p \in \mathcal{N}, \\ y_i^p \neq c}} \underbrace{\lambda_{(i, \mathbf{y})} \hat{h}_{y_i^p}^{(1)}(\mathbf{x}_i^p)}_{\text{negative}} \right]. \end{aligned} \quad (17)$$

To solve the above optimization problems, we here train K weighted decision tree classifiers. Specifically, when training decision trees for the c -th class, the training data is composed of those superpixels whose ground truth label or predicted label is equal to the category label c . Since the output of the decision tree is in $\{0, 1\}$ and $\lambda_{(i, \mathbf{y})} \geq 0$, the maximization in Eqn. (17) is achieved if $\hat{h}_c^{(1)}$ outputs 1 for each of the superpixel p with $y^p = c$, and outputs 0 for each of the superpixel p with $y_i^p \neq c$. Therefore, as indicated by the horizontal curly braces in Eqn. (17), superpixels with the predicted labels of category c are used as positive training examples, while superpixels with ground truth labels of category c are used as negative training examples. The dual solution λ serves as weightings of the training data.

For the pairwise part, we generate one decision tree in each column generation iteration, hence $\mathbf{H}^{(2)*}$ can be written as $\mathbf{H}^{(2)*} = [\hat{h}^{(2)*}]$, the new decision tree for the pairwise part is generated as:

$$\begin{aligned} \hat{h}^{(2)*}(\cdot, \cdot) &= \operatorname{argmax}_{\hat{h} \in \mathcal{H}} \sum_{i, \mathbf{y}} \lambda_{(i, \mathbf{y})} \left[\sum_{(p, q) \in \mathcal{S}} \underbrace{\hat{h}^{(2)}(\mathbf{x}^p, \mathbf{x}^q) I(y^p \neq y^q)}_{\text{positive}} \right. \\ &\quad \left. - \sum_{(p, q) \in \mathcal{S}} \underbrace{\hat{h}^{(2)}(\mathbf{x}^p, \mathbf{x}^q) I(y_i^p \neq y_i^q)}_{\text{negative}} \right]. \end{aligned} \quad (18)$$

Similar to the unary case, we train a weighted decision tree classifier with λ as training example weightings. The positive and negative training data are indicated by the horizontal curly braces in Eqn. (18). $\hat{h}^{(2)}$ is the response of a decision tree applied on the pairwise features constructed by two neighbouring superpixels ($\mathbf{x}^p, \mathbf{x}^q$), *e.g.*, color differences or shared boundary lengths.

With the above analysis, we can now apply column generation to jointly learn the decision trees $\mathbf{H}^{(1)}, \mathbf{H}^{(2)}$ and $\mathbf{w}$. The column generation (CG) procedure iterates the following two steps:

1) Solve Eqn. (17), Eqn. (18) to generate decision trees $\mathbf{H}^{(1)*}, \mathbf{H}^{(2)*}$;

2) Add $\mathbf{H}^{(1)*}$ and $\mathbf{H}^{(2)*}$ to working set $\mathcal{W}_H$ and resolve for the primal solution $\mathbf{w}$ and dual solution λ .

We show two segmentation examples on the Oxford flower dataset produced by our method with different CG iterations in Fig. 1. As can be seen, our method refines the segmentation with the increase of CG iterations. Since this dataset is relatively simple, a few CG iterations are enough to get satisfactory results.

Algorithm 1 CRFTree using column generation

1. **Input:** training examples $(\mathbf{x}_1; \mathbf{y}_1), (\mathbf{x}_2; \mathbf{y}_2), \dots$; maximum iteration number.
 2. **Initialize** $(\lambda, \mathbf{y})$, and decision tree working set $\mathcal{W}_H \leftarrow \emptyset$
 3. **Repeat**
 4. Find decision trees $\mathbf{H}^*$ by solving Eqn. (17), Eqn. (18). Add $\mathbf{H}^*$ to working set $\mathcal{W}_H$.
 5. Call Alg. 2 using working set $\mathcal{W}_H$ to solve for $\mathbf{w}$ and λ .
 6. **Until** the maximum iteration is reached.
 7. **Output:** $\mathbf{w}, \mathbf{H} \in \mathcal{W}_H$.
-

Algorithm 2 Cutting-planes for solving the 1-slack primal

1. **Input:** cutting-plane termination threshold ϵ_{cp} , and inputs from Alg. 1.
 2. **Initialize:** working set $\mathcal{W} \leftarrow \emptyset$; $\mathbf{r}$.
 3. **Repeat**
 4. $\mathcal{W} \leftarrow \mathcal{W} \cup \{(r_1, \dots, r_m, \mathbf{y}_1^*, \dots, \mathbf{y}_m^*)\}$.
 5. Obtain primal and dual solutions $\mathbf{w}, \xi; \lambda$ by solving (19) on $\mathcal{W}$.
 6. **For** $i = 1, \dots, m$
 7. Solve the inference problem in Eqn. (20) using Graph-Cut to find the most violated $\mathbf{y}_i^*$.
 8. **End for**
 9. **Until**
 10. $\frac{1}{m} \mathbf{w}^\top \left[\sum_{i=1}^m r_i [\Psi(\mathbf{y}_i^*, \mathbf{x}_i) - \Psi(\mathbf{y}_i, \mathbf{x}_i)] \right] \geq \frac{1}{m} \sum_{i=1}^m r_i \Delta(\mathbf{y}_i, \mathbf{y}_i') - \xi - \epsilon_{cp}$.
 10. **Output:** $\mathbf{w}, \lambda$.
-

For solving the primal problem in the second step, it involves a large number of constraints due to the large output space $\{\mathbf{y} \in \mathcal{Y}\}$. We next show how to apply the cutting-plane technique [13] to efficiently solve this problem.

E. Speeding up optimization using cutting-plane

To apply cutting-plane for solving the optimization in Eqn. (13), we first derive its 1-slack formulation. The 1-slack SSVMS formulation was first introduced by [13]. The 1-slack formulation of our method can be written as:

$$\begin{aligned} \min_{\mathbf{w} \geq 0, \xi \geq 0} \quad & \frac{1}{2} \|\mathbf{w}\|_2^2 + C\xi \\ \text{s.t.:} \quad & \frac{1}{m} \mathbf{w}^\top \left[\sum_{i=1}^m r_i \cdot [\Psi(\mathbf{y}, \mathbf{x}_i; \mathbf{H}) - \Psi(\mathbf{y}_i, \mathbf{x}_i; \mathbf{H})] \right] \\ & \geq \frac{1}{m} \sum_{i=1}^m r_i \Delta(\mathbf{y}_i, \mathbf{y}) - \xi, \forall \mathbf{r} \in \{0, 1\}^m; \forall \mathbf{y} \in \mathcal{Y}. \end{aligned} \quad (19)$$

Cutting-plane methods work by finding the most violated constraint for each example i

$$\mathbf{y}_i^* = \operatorname{argmin} \mathbf{w}^\top \Psi(\mathbf{y}, \mathbf{x}; \mathbf{H}) - \Delta(\mathbf{y}_i, \mathbf{y}) \quad (20)$$

at every iteration and add it to the constraint working set. The sketch of our method is summarized in Algorithm 1, which calls Algorithm 2 to solve the 1-slack optimization.

Implementation details To deal with the unbalanced appearance of different categories in the dataset, we define $\Delta(\mathbf{y}_i, \mathbf{y})$ as weighted Hamming loss, which weighs errors for a given class inversely proportional to the frequency it appears in the training data, as similar in [24]. In the inference problem of Eqn. (20), when using the hamming loss as the label cost Δ , the label cost term can be absorbed into the unary part. We therefore can apply Graph-cut to efficiently solve Eqn. (20). As for more complicated label cost functions, an efficient inference algorithm is proposed in [4]. During each

Fig. 1: Segmentation examples (each row being an example) produced by our model on images from the Oxford 17 Flower dataset with different column generation iterations. From left to right: Test images, 2nd, 4th, 6th, and 10th iteration.

CG iteration, our method first solves Eqn. (17), (18) given the current $\mathbf{x}$ and ξ , and then solves a quadratic programming (QP) problem given $\mathbf{H}$. When solving Eqn. (17), (18), we train weighted decision tree classifiers using the highly optimized decision tree training method of [3].

Discussions on the submodularity It is known that if graph cuts are to be applied to achieve globally optimum labelling in segmentation, the energy function must be submodular. For foreground/background segmentation in which a (super-)pixel label takes value in $\{0, 1\}$, we show that our method keeps this submodular property. It is commonly known that an energy function is submodular if its pairwise term satisfies: $\eta_{pq}(0, 0) + \eta_{pq}(1, 1) \leq \eta_{pq}(0, 1) + \eta_{pq}(1, 0)$. Recall that our pairwise energy is written as $\eta_{pq}(y^p, y^q) = \mathbf{w}^{(2)\top} \mathbf{H}^{(2)}(\mathbf{x}^p, \mathbf{x}^q) I(y^p \neq y^q)$. Clearly we have $(\eta_{pq}(0, 0) = \eta_{pq}(1, 1) = 0)$ because of the indicator function $I(y^p \neq y^q)$. The second thing is to ensure $\eta_{pq}(1, 0) + \eta_{pq}(0, 1) \geq 0$. Given the non-negativeness constraint we impose on $\mathbf{w}$ in our model, and the output of decision trees in our method taking values from $\{0, 1\}$, we have $\eta_{pq}(1, 0) \geq 0$ and $\eta_{pq}(0, 1) \geq 0$. We thus accomplish the proof of the submodularity of our model. In the case of multi-object segmentation, the inference is done by the α -expansion of graph cuts.

Discussions on the non-negative constraint on $\mathbf{w}$ Our learning framework aligns with boosting methods, where we learn a non-negative weighted ensemble of weak structured learners (constructed by decision trees), which is analogous to weak learners in boosting methods. This is similar to boosting methods, such as AdaBoost, LPBoost [8], where the non-negative weighting is commonly used. Further, a weak structured learner generated by our column generation method is expected to make positive contribution to the learning objective. If it is of no use to the objective, the weight will approach zero. Therefore it is reasonable to enforce the non-negative constraint on $\mathbf{w}$.

III. EXPERIMENTS

To demonstrate the effectiveness of the proposed method, we first compare our model with some most related baseline methods, which are SVMs, AdaBoost and SSVMs. In section III-C, we show that our method achieves state-of-the-art results by exploiting recent advances in feature learning [6], [16].

A. Experimental setup

The datasets evaluated here include three binary datasets (Weizmann horse, Oxford flower and Graz-02) and two multi-

class datasets (MSRC-21 and PASCAL VOC 2012). The Weizmann horse dataset¹ consists of 328 horse images from various backgrounds, with groundtruth masks available for each image. We use the same data split as in [5] and [17]. The Oxford 17 category flower dataset [26] is composed of 849 flower images. Those with too small foreground are removed, which leaves 753 for segmentation purpose [26]. The data split stated in [26] is used to perform the evaluation. During our experiment, images of the Weizmann horse and the Oxford flower datasets are resized to 256×256 . The Graz-02 dataset² contains 3 categories (bike, car and people). This dataset is considered challenging as the objects appear at various background and with different poses. We follow the evaluation protocol in [25] to use 150 for training and 150 for testing for each category. The MSRC-21 dataset [32] is a popular multi-class segmentation benchmark with 591 images containing objects from 21 categories. We follow the standard split to divide the dataset into training/validation/test subsets. The PASCAL VOC 2012 dataset³ is a widely used benchmark for semantic segmentation, which contains 2913 images from the trainval set and 1456 images from the test set, making up 21 categories. Unlike many state-of-the-arts methods such as [38], we do not use any additional training data for this dataset.

We start with over-segmenting the images into superpixels using SLIC [1], with ~ 700 superpixels generated per image. We extract dense SIFT descriptors and color histograms around each superpixel centroid with different block sizes (12×12 , 24×24 , 36×36). The dense SIFT descriptors are then quantized into bag-of-words features using nearest neighbour search with a codebook size of 400. We construct four types of pairwise features also using different block sizes to enforce spatial smoothness, which are color difference in LUV space, color histogram difference, texture difference in terms of LBP operators as well as shared boundary length [10]. The column generation iteration number of our CRFTree is set to 50 based on a validation set. We learn tree potentials with the tree depth being 2. Training on the MSRC-21 dataset on a standard PC machine takes around 16 hours.

¹<http://www.msri.org/people/members/eranb/>

²<http://www.emt.tugraz.at/~pinz/>

³<http://host.robots.ox.ac.uk/pascal/VOC/voc2012/>

Category	bike	car	people
	intersection/union (foreground, background)(%)		
SVMs	67.8 (51.9, 83.8)	69.7 (46.8, 92.6)	65.0 (44.5, 85.5)
AdaBoost	71.2 (57.6, 84.9)	71.0 (49.4, 92.6)	67.7 (48.7, 86.7)
SSVMs	72.2 (58.6, 85.8)	76.9 (60.0, 94.2)	70.9 (53.8, 87.9)
CRFTree	76.4 (65.0, 87.8)	79.5 (64.0, 95.0)	74.2 (58.7, 89.7)
CRFTree (FL)	78.3 (67.7, 88.9)	83.0 (70.1, 95.9)	75.7 (61.0, 90.5)
	pixel accuracy (foreground, background)(%)		
SVMs	79.5 (67.4, 91.5)	77.3 (57.2, 97.3)	77.7 (63.8, 91.6)
AdaBoost	83.8 (77.3, 90.3)	80.1 (63.5, 96.6)	80.5 (69.0, 91.9)
SSVMs	83.8 (76.1, 91.6)	85.5 (73.8, 97.2)	83.9 (75.8, 92.1)
CRFTree	87.8 (83.9, 91.8)	87.0 (76.4, 97.7)	85.9 (78.4, 93.4)
CRFTree (FL)	89.1 (85.8, 92.4)	90.0 (82.1, 98.0)	86.9 (80.0, 94.0)

TABLE I: The average intersection-over-union score and average pixel accuracy comparison on the Graz-02 dataset. We include the foreground and background results in the brackets. Our method CRFTree with nonlinear and class-wise potentials learning performs better than all the baseline methods.

B. Comparing with baseline methods

We first compare CRFTree with some conventional methods, which are linear SVMs, AdaBoost and SSVMs to demonstrate the superiority of our method. For SVMs and AdaBoost, each superpixel is classified independently without CRFs. We mainly evaluate on the more challenging Graz-02 and MSRC-21 dataset in this part. The regularization parameter C of SVMs, SSVMs and our CRFTree are selected from $\{1, 10, 100, 1000\}$ based on a validation set. We use depth-2 decision trees for training AdaBoost and our CRFTree. The maximum iteration number of AdaBoost is chosen from $\{50, 100, 200\}$. For our method, we treat the foreground and background as two categories in the binary case to learn class-wise potentials.

Graz-02 For a comprehensive evaluation, we use two measurements to quantify the performance on the Graz-02 dataset, which are intersection over union score and the pixel accuracy (including foreground and background). We report the results in Table I. As can be observed, AdaBoost based on a depth-2 decision tree performs better than the linear SVMs. On the other hand, structured methods which jointly consider local information and spatial consistency are able to significantly outperform the simple binary models. By introducing nonlinear and class-wise potential learning, our method is able to gain further improvement over SSVMs.

MSRC-21 We learn class-wise potentials using our CRFTree for each of the 21 classes on the MSRC dataset. The compared results are summarized in Table II (upper part). Similar conclusions can be drawn as on the Graz-02 dataset and our CRFTree again outperforms all its baseline competitors.

C. Comparing with state-of-the-art methods

Since features play a pivotal role in the performance of vision algorithms, we exploit recent advances in feature learning to pursue state-of-the-art results, *i.e.*, unsupervised feature learning [6] and convolutional neural networks (CNN) [16]. Specifically, for the unsupervised feature learning, we first learn a dictionary $\mathbf{B}$ of size 400 and patch size 6×6 based on the evaluated image dataset using Kmeans, and then use the soft threshold coding [6] to encode patches extracted from each superpixel block. The final feature vectors (we call it

encoding feature here) are obtained by performing a three-level max pooling over the superpixel block. For the CNN features, we use the Alex model [16] trained on the ImageNet⁴ to generate CNN features. These two versions of our method are denoted as CRFTree (FL) and CRFTree (CNN) respectively. We only report the results of CRFTree (CNN) on the MSRC-21 and PASCAL VOC 2012 datasets since our method already performs very well by using the encoding features on the three binary datasets.

Weizmann horse We quantify the performance by the global pixel-wise accuracy S_a and the foreground intersection over union score S_o , as did in [5]. S_a measures the percentage of pixels correctly classified while S_o directly reflects the segmentation quality of the foreground. The results are reported in Table III. Our method performs better than the kernel structural learning method of [5], which may result from the fact that they only introduced nonlinearity into the unary part while our method achieves nonlinearity on both unary and pairwise terms. The best S_a score is obtained by [21]. However their method relies on an assumption that a perfect bounding box of the horse is available for each test image, which is not practically applicable. On the contrary, we provide a principal and general way of nonlinearly learning CRFs parameters. We show some segmentation examples of our method in Fig. 2.

Oxford flower As in [5], we also use S_a and S_o to measure the performance on the Oxford flower dataset, and report the results in Table IV. Our method performs comparable to the original work of [26] on this dataset in terms of S_o while again obtains better results than the closely related state-of-the-art work of [5]. It is also worth noting that the method in [26] is very domain specific, which relies on modelling the flower's shape (center and petal), while ours is generally applicable.

Graz-02 As in the work of [25], [9], [2], [17], we also evaluate the F-score on the Graz-02 dataset besides the above mentioned intersection over union score and pixel accuracy. The F-score is defined as $F = 2pr/(p + r)$, where p is the precision and r is the recall. We summarize the results in Table V and Table I. From Table V, it can be seen that our method significantly outperforms all the compared methods, which fully demonstrate the power of nonlinear and class-wise potential learning. Furthermore, we can observe from Table I that compared with the previous results, adding more features help to improve the performance.

MSRC-21 The compared results with state-of-the-art works are reported in the lower part of Table II. As we can see, by incorporating more advanced features, our CRFTree gains significant improvements over the previous results which only use bag-of-words and color histogram features. It is worth noting that our method performs better than the closely related work of Lucchi *et al.* [24] which claims exploiting non-linear kernels. It has to be pointed out that we did not employ any global potentials (while in [24], they improve the global and average per-category accuracy from 70, 73 to 82 and 76 by adding global information). If global or higher potentials are incorporated into our model, further performance promotion can be expected. We show some qualitative evaluation exam-

⁴<http://image-net.org>

	building	grass	tree	cow	sheep	sky	aeroplane	water	face	car	bicycle	flower	sign	bird	book	chair	road	cat	dog	body	boat	Average	Global
SVMs	54	92	73	41	54	80	51	67	51	41	59	41	28	8	64	17	75	41	23	20	7	47.0	63.7
AdaBoost	68	92	83	48	58	87	43	69	58	43	64	41	32	14	70	28	79	47	22	41	6	52.0	68.6
SSVMs	65	92	81	42	76	84	65	70	75	54	87	62	31	14	76	31	78	61	30	25	2	57.2	70.8
CRFTree	53	87	85	59	84	90	77	82	81	54	90	57	62	22	81	59	80	71	26	49	15	64.9	73.9
CRFTree (FL)	66	95	89	83	89	90	90	83	76	74	83	71	69	46	87	73	87	84	53	68	20	75.1	82.2
CRFTree (CNN)	73	96	89	82	92	96	89	86	93	78	86	91	71	75	85	76	86	91	63	83	41	82.0	86.2
Shotton <i>et al.</i> [32]	49	88	79	97	97	78	82	54	87	74	72	74	36	24	93	51	78	75	35	66	18	67	72
Ladicky <i>et al.</i> [18]	80	96	86	74	87	99	74	87	86	87	82	97	95	30	86	31	95	51	69	66	9	75	86
Gonfaus <i>et al.</i> [11]	60	78	77	91	68	88	87	76	73	77	93	97	73	57	95	81	76	81	46	56	46	75	77
Lucchi <i>et al.</i> [24]	59	90	92	82	83	94	91	80	85	88	96	89	73	48	96	62	81	87	33	44	30	76	82
Lucchi <i>et al.</i> [23]	67	89	85	93	79	93	84	75	79	87	89	92	71	46	96	79	86	76	64	77	50	78.9	83.7

TABLE II: Segmentation results on the MSRC dataset. We report the pixel-wise accuracy for each category as well as the average per-category scores and the global pixel-wise accuracy. (1) The upper part presents the comparison with baseline methods, which all use bag-of-words and color histogram features. Our method CRFTree gains impressive improvements over SSVMs while far better than simple linear models. (2) The lower part shows the results of our method using unsupervised feature learning and CNN features (denoted as CRFTree (FL) and CRFTree (CNN) respectively) compared with state-of-the-art methods on this dataset.

Fig. 2: Segmentation examples on Weizmann horse. 1st row: Test images; 2nd row: Ground truth; 3rd row: segmentation produced by CRFTree.

Fig. 3: Segmentation examples on MSRC. 1st row: Test images; 2nd row: Ground truth; 3rd row: CRFTree with unsupervised feature learning.

Fig. 4: Segmentation examples on PASCAL VOC 2012. 1st row: Images; 2nd row: Ground truth; 3rd row: CRFTree prediction results.

ples in Fig. 3.

PASCAL VOC 2012 We generate deep features of each superpixel by averaging the pixel-wise feature map scores within the superpixel obtained from a pretrained FCN model [22]. We then train our CRFTree model on the standard PASCAL VOC 2012 training dataset with the generated deep

features. Following the standard evaluation procedure for the Pascal VOC challenge, we upload our segmentation results to the test server and use the average intersection over union as the evaluation metric. We compare against several state-of-the-art methods ([12], [7], [22], [38]) on the test set of the PASCAL VOC 2012 dataset. The results are reported in

Method	Sa	So
Levin & Weiss [21]	95.5	-
Cosegmentation [14]	80.1	-
Bertelli <i>et al.</i> [5]	94.6	80.1
Kuttel <i>et al.</i> [17]	94.7	-
CRFTree (FL)	94.6	80.4

TABLE III: Performance of different methods on the Weizmann Horse dataset.

Method	Sa	So
Nilsback <i>et al.</i> [26]	-	94.0
Bertelli <i>et al.</i> [5]	97.7	92.3
CRFTree (FL)	98.0	94.2

TABLE IV: Performance of different methods on the Oxford FLower dataset. Our method CRFTree performs better than the compared methods.

Table VI. As seen from the table, our CRFTree beats the Hypercolumn [12] and the CFM [7] and outperforms the FCN [22] by a notable margin. Although our method is triumphed by [38], it should be noted that their result is obtained by using extra training data (11,685 images vs 1456 images used for training our CRFTree). Some qualitative evaluation examples of our method are illustrated in Fig. 4.

IV. CONCLUSION

Nonlinear structured learning has been a promising yet challenging topic in the community. In this work, we have proposed a nonlinear structured learning method of tree potentials for image segmentation. The unary and pairwise potentials are ensembles of class-wise trees, with the ensemble parameters and the trees jointly learned in a unified large-margin framework. In this way, nonlinearity is easily introduced into the CRFs learning. The resulted model involves exponential number of variables and constraints. We therefore derive a novel algorithm combining a modified column generation method and the cutting-plane technique for efficient model training. We have exemplified the superiority of the proposed nonlinear potential learning method by comparing against state-of-the-art methods on both binary and multi-class object

Method	bike	car	people	average
Marszalek & Schmid [25]	61.8	53.8	44.1	53.2
Fulkerson <i>et al.</i> [9]	66.4	54.7	51.4	57.5
Aldavert <i>et al.</i> [2]	71.9	62.9	58.6	64.5
Kuettel <i>et al.</i> [17]	63.2	74.8	66.4	68.1
CRFTree (FL)	80.7	82.4	75.8	79.5

TABLE V: Comparing with state-of-the-art methods on the Graz-02 dataset. We report the F-score (%) for each class and the average over classes. Our method CRFTree outperforms all the compared methods with a large margin.

Method	intersection/union
CFM [7]	61.8
Hypercolumn [12]	62.6
FCN-8s [22]	62.2
Zheng <i>et al.</i> [38]	72.0
CRFTree	65.4

TABLE VI: The average intersection-over-union scores of different methods on the PACAL VOC 2012 test dataset. Our method CRFTree achieves comparable performance with the state-of-the-arts. Note that [38] used extra training data.

segmentation datasets. A potential disadvantage of our method is that it is prone to overfitting due to the outstanding non-linear learning capacity. This can be alleviated by using more training data. On the other hand, as we show in Table II, our method using pre-trained CNN features has shown the best performance. Therefore it is worth exploiting to further combine our method with deep learning techniques in the future work.

REFERENCES

- [1] R. Achanta, K. Smith, A. Lucchi, P. Fua, and S. Ssstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2012.
- [2] D. Aldavert, A. Ramisa, R. L. de Mntaras, and R. Toledo. Fast and robust object segmentation with the integral linear classifier. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2010.
- [3] R. Appel, T. J. Fuchs, P. Dollár, and P. Perona. Quickly boosting decision trees - pruning underachieving features early. In *Proceedings of International Conference on Machine Learning*, 2013.
- [4] A. Bauer, S. Nakajima, and K.-R. Müller. Efficient exact inference with loss augmented objective in structured learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2016.
- [5] L. Bertelli, T. Yu, D. Vu, and B. Gokturk. Kernelized structural SVM learning for supervised object segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2011.
- [6] A. Coates and A. Y. Ng. The importance of encoding versus training with sparse coding and vector quantization. In *Proceedings of International Conference on Machine Learning*, 2011.
- [7] J. Dai, K. He, and J. Sun. Convolutional feature masking for joint object and stuff segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [8] A. Demiriz, K. P. Bennett, and J. Shawe-Taylor. Linear programming boosting via column generation. *Machine Learning*, 2002.
- [9] B. Fulkerson, A. Vedaldi, and S. Soatto. Localizing objects with smart dictionaries. In *Proceedings of European Conference on Computer Vision*, 2008.
- [10] B. Fulkerson, A. Vedaldi, and S. Soatto. Class segmentation and object localization with superpixel neighborhoods. In *Proceedings of IEEE International Conference on Computer Vision*, 2009.
- [11] J. M. Gonfaus, X. B. Bosch, J. van de Weijer, A. D. Bagdanov, J. S. Gual, and J. G. Sabaté. Harmony potentials for joint classification and segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2010.
- [12] B. Hariharan, P. A. Arbeláez, R. B. Girshick, and J. Malik. Hypercolumns for object segmentation and fine-grained localization. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [13] T. Joachims. Training linear SVMs in linear time. In *Proceedings of ACM SIGKDD ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2006.
- [14] A. Joulin, F. R. Bach, and J. Ponce. Discriminative clustering for image co-segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2010.
- [15] M. Kim. Mixtures of conditional random fields for improved structured output prediction. *IEEE Transactions on Neural Networks and Learning Systems*, 2016.
- [16] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Proceedings of Advances in Neural Information Processing Systems*, 2012.
- [17] D. Kuettel and V. Ferrari. Figure-ground segmentation by transferring window masks. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2012.
- [18] L. Ladicky, C. Russell, P. Kohli, and P. H. S. Torr. Associative hierarchical crfs for object class image segmentation. In *Proceedings of IEEE International Conference on Computer Vision*, 2009.
- [19] J. Lafferty, A. McCallum, and F. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of International Conference on Machine Learning*, 2001.
- [20] J. Lafferty, X. Zhu, and Y. Liu. Kernel conditional random fields: Representation and clique selection. In *Proceedings of International Conference on Machine Learning*, 2004.
- [21] A. Levin and Y. Weiss. Learning to combine bottom-up and top-down segmentation. In *Proceedings of European Conference on Computer Vision*, 2006.
- [22] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [23] A. Lucchi, Y. Li, and P. Fua. Learning for structured prediction using approximate subgradient descent with working sets. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2013.
- [24] A. Lucchi, Y. Li, K. Smith, and P. Fua. Structured image segmentation using kernelized features. In *Proceedings of European Conference on Computer Vision*, 2012.
- [25] M. Marszalek and C. Schmid. Accurate object localization with shape masks. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2007.
- [26] M.-E. Nilsback and A. Zisserman. Delving deeper into the whorl of flower segmentation. *Image Vision Computing*, 2010.
- [27] S. Nowozin, P. Gehler, and C. H. Lampert. On parameter learning in CRF-based approaches to object class image segmentation. In *Proceedings of European Conference on Computer Vision*, 2010.
- [28] S. Nowozin, C. Rother, S. Bagon, T. Sharp, B. Yao, and P. Kohli. Decision tree fields. In *Proceedings of IEEE International Conference on Computer Vision*, 2011.
- [29] A. Severyn and A. Moschitti. Fast support vector machines for structural kernels. In *Proceeding of European Conference on Machine Learning & Knowledge Discovery in Databases*, 2011.
- [30] C. Shen and Z. Hao. A direct formulation for totally-corrective multi-class boosting. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2011.
- [31] C. Shen, G. Lin, and A. van den Hengel. Structboost: Boosting methods for predicting structured output variables. *CoRR*, 2013.
- [32] J. Shotton, M. Johnson, and R. Cipolla. Semantic texton forests for image categorization and segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2008.
- [33] I. Steinwart. Sparseness of support vector machines. *Journal of Machine Learning Research*, 2003.
- [34] M. Szummer, P. Kohli, and D. Hoiem. Learning CRFs using graph cuts. In *Proceedings of European Conference on Computer Vision*, 2008.
- [35] B. Taskar, C. Guestrin, and D. Koller. Max-margin Markov networks. In *Proceedings of Advances in Neural Information Processing Systems*, 2004.
- [36] I. Tsochantaris, T. Joachims, T. Hofmann, and Y. Altun. Large margin methods for structured and interdependent output variables. *Journal of Machine Learning Research*, 2005.
- [37] C.-N. Yu and T. Joachims. Training structural SVMs with kernels using sampled cuts. In *Proceedings of ACM SIGKDD ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2008.
- [38] 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 *Proceedings of IEEE International Conference on Computer Vision*, 2015.