

Robust Trust Region for Weakly Supervised Segmentation

Dmitrii Marin
University of Waterloo, Canada
dmitrii.a.marin@gmail.com

Yuri Boykov
University of Waterloo, Canada
yboykov@uwaterloo.ca

Abstract

Acquisition of training data for the standard semantic segmentation is expensive if requiring that each pixel is labeled. Yet, current methods significantly deteriorate in weakly supervised settings, e.g. where a fraction of pixels is labeled or when only image-level tags are available. It has been shown that regularized losses—originally developed for unsupervised low-level segmentation and representing geometric priors on pixel labels—can considerably improve the quality of weakly supervised training. However, many common priors require optimization stronger than gradient descent. Thus, such regularizers have limited applicability in deep learning. We propose a new robust trust region approach¹ for regularized losses improving the state-of-the-art results. Our approach can be seen as a higher-order generalization of the classic chain rule. It allows neural network optimization to use strong low-level solvers for the corresponding regularizers, including discrete ones.

1. Introduction

Our paper proposes a higher-order optimization technique for neural network training. While focused on semantic image segmentation, our main algorithmic idea is simple and general - integrate the standard *trust region* principle into the context of *backpropagation*, i.e. the chain rule. We reinterpret the classic chain rule: instead of the chain of gradients/derivatives for a composition of functions, we formulate the corresponding chain of hidden optimization subproblems. Then, inspired by the *trust region* principle, we can substitute a standard linear approximation solver (gradient descent) at any chain with a better higher-order solver. In short, we replace the classic differentiation chain rule by the trust region chain rule in the context of backpropagation.

Our work is motivated by the well-known challenges presented to the gradient descent by typical regularization losses or geometric priors/energies ubiquitous in the context of weakly-supervised or unsupervised segmentation.

To validate our approach, we present semantic segmentation results improving the state-of-the-art in the challenging setting where the training data has only a fraction of pixels labeled. The generality of our main principle (trust region chain rule) and our promising results for a difficult problem encourage further research. In fact, this work applies trust region principle only to the last “chain” in the network. We discuss several promising extensions for future work.

The rest of the introduction is organized as follows. To create a specific context for our general approach to network training, we review loss functions relevant for weakly-supervised or unsupervised segmentation. First, Sec. 1.1 discusses several standard geometric priors, regularization energies, clustering criteria, and their powerful solvers originally developed for low-level segmentation or general machine learning. Then, Sec. 1.2 outlines the use of such regularization objectives as losses for network training in the context of weakly supervised semantic (high-level) segmentation. We also review the standard *trust region* principle (Sec. 1.4) and highlight our main contributions (Sec. 1.5) based on the general idea of applying trust region (with powerful solvers) to network training.

1.1. Regularized energies in low-level segmentation

Assuming discrete segmentation $s \in \{1, 2, \dots, K\}^N$ where K is the number of categories and N is the number of image pixels, one common low-level segmentation energy can be represented as

$$E(s) = - \sum_i \log P(I_i | s_i) + \sum_{\{i,j\} \in \mathcal{N}} w_{ij} [s_i \neq s_j] \quad (1)$$

where I_i is a low-level feature (e.g. intensity, color, texture) at pixel i with distribution functions $P(\cdot | k)$ for each category k , neighborhood system \mathcal{N} describes any pairwise connectivity (typically 4-, 8-grid [5] or denser [34]), weights w_{ij} represent given pairwise affinities (typically Gaussian kernel for low-level features I_i and I_j [7, 5, 54, 34]), and $[\cdot]$ is the Iverson bracket operator returning 1 if the argument is true and 0 otherwise. The energy above combines the log-likelihoods term enforcing consistency with given (low-level) feature distributions and a pairwise reg-

¹https://github.com/dmitrii-marin/robust_trust_region

ularizer (Potts model) term enforcing geometric prior on shape smoothness with alignment to image intensity edges.

The Potts model has several efficient combinatorial [7] and LP-relaxation solvers [32, 35]. Besides, there are many regularization objectives that are closely related to the first-order shape regularization in (1), but derived from a different discrete or continuous formulation of the low-level segmentation and equipped with their own efficient solvers, *e.g.* geodesic active contours [10], snakes [26], power watersheds [18], to name a few. Moreover, there are many other regularization terms going beyond the basic first-order smoothness (boundary length) enforced by the Potts term in (1). The extensions include curvature [57, 47, 46], Pn-Potts [30], convexity [24, 25], etc.

Common continuous formulations of the low-level segmentation use *relaxed* variable $s \in \Delta_K^N$ combining pixel-specific distributions $s_i = (s_i^1, \dots, s_i^K) \in \Delta_K$ over K categories, where Δ_K is the *probability simplex*. In this case the segmentation objective/energy should also be relaxed, *i.e.*, defined over real-values arguments. For example, one basic relaxation of the Potts segmentation energy in (1) is

$$-\sum_i \sum_k s_i^k \log P(I_i|k) + \sum_{\{i,j\} \in \mathcal{N}} w_{ij} \|s_i - s_j\|^2 \quad (2)$$

using a linear relaxation of the likelihood term and a quadratic relaxation of the Potts model. Note that there could be infinitely many alternative relaxations. Any specific choice affects the properties of the relaxed solution, as well as the design of the corresponding optimization algorithm. For example, simple *quadratic* relaxation in (2) is convex suggesting simpler optimization, but its known to be a non-tight relaxation of the Potts model [53] leading to weaker regularization properties unrelated to geometry or shape. There are many better alternatives, *e.g.* using different norms [18] or other convex formulations [13, 12, 11]. The *bilinear* relaxation of the Potts term below

$$-\sum_i \sum_k s_i^k \log P(I_i|k) + \sum_k (1 - s^k)^\top W s^k \quad (3)$$

is tight [53], but it is non-convex and, therefore, more difficult to optimize. In the formula above, vector $s^k := (s_i^k)$ combines segmentation variables for soft-segment k , and $N \times N$ affinity matrix $W_{ij} = w_{ij} [\{i, j\} \in \mathcal{N}]$ represents the neighborhood system \mathcal{N} and all pairwise (*e.g.* Gaussian) affinities w_{ij} between image pixels. Note that Potts regularization is closely related to the *Normalized cut* objective $\sum_k \frac{(1-s^k)^\top W s^k}{1^\top W s^k}$ for unsupervised segmentation [58].

It is common to combine energies like (1),(2),(3) with constraints based on user interactions (weak supervision). While there are different forms of such supervision, the most basic one is based on adding the seed loss [5] defined over pixels in subset Ω_{seeds} with user-specified category labels y_i . Assuming $s_i \in \Delta_K$, it can be written as a partial

cross entropy (PCE) for pixels $i \in \Omega_{\text{seeds}}$

$$E_{\text{seeds}}(s) = - \sum_{i \in \Omega_{\text{seeds}}} \log s_i^{y_i} \quad (4)$$

and, when restricted to *one-hot* s_i representing hard segmentation, it reduces to the hard constraints over seeds [5]. That is, for integer-valued $s_i \in \{1, \dots, K\}$ the seed loss is equivalent to $\sum_{i \in \Omega_{\text{seeds}}} \lambda [s_i = y_i]$ for infinitely large λ .

The log-likelihood loss, *e.g.* the first term in (1) or (3), is common in low-level segmentation and its importance cannot be underestimated. In basic formulations, the distributions of (low-level) features $P(\cdot|k)$ can be assumed given for each category k . However, if such distributions are not known *a priori*, their representation $P(\cdot|\theta_k)$ can explicitly include unknown distribution parameters θ_k for each category k . Then, the overall loss $E(s, \theta)$ adds $\theta = \{\theta_k\}$ as an extra variable. Optimization of $E(s, \theta)$ over both s and θ corresponds to joint estimation of segmentation and *maximum likelihood* (ML) estimation of distribution parameters, as in well-known unsupervised low-level segmentation formulations by Zhu & Yuille [66] and Chan & Vese [14]. Similar ideas are also used in box-interaction methods [54].

1.2. Regularized losses in DNN segmentation

Unlike low-level segmentation methods based on readily available low-dimensional features (like color, texture, contrast edges), deep neural network (DNN) approaches to segmentation learn complex high-dimensional “deep” features that can discriminate semantic categories. Thus, one can refer to such methods as *high-level* segmentation, and to such learned features as *high-level* features.

The most standard way to train segmentation networks is based on *full supervision* requiring a large collection of images where all pixels are accurately labeled. Such training data is expensive to get. The training is based on minimizing the cross-entropy (CE) loss similar to the seed loss in low-level segmentation. For simplicity focusing on a single training image, CE loss is

$$E_{\text{CE}}(s(\theta)) = - \sum_i \log s_i^{y_i}(\theta) \quad (5)$$

where $s(\theta) = \mathbf{f}(\theta) \in \Delta_K^N$ is the (relaxed) segmentation output of the network $\mathbf{f}(\theta)$ with parameters θ . For brevity, here and later in this paper we omit the actual test image from the arguments of the network function \mathbf{f} .

The fundamental difference with low-level segmentation reviewed above is that instead of minimizing losses E directly over segmentation variable s , now the optimization arguments are parameters θ of the network producing such segmentation. Estimating parameters θ can be interpreted as learning deep features. Note that this task is much more complex than ML estimation of distribution parameters for

$P(I|\theta)$ in low-level segmentation with fixed low-level features I , as reviewed above. This explains why network optimization requires a large set of fully labeled training images, rather than a single image (unlabeled or partially-labeled), as in low-level segmentation.

The goal of weakly supervised segmentation is to train the network with as little supervision as possible. First of all, it is possible to train using only a subset of labeled pixels (seeds) in each image [31, 61] in exact analogy with (4)

$$E_{\text{PCE}}(s(\theta)) = - \sum_{i \in \Omega_{\text{seeds}}} \log s_i^{y_i}(\theta) \quad (6)$$

In particular, as shown in [61], this simple, but principled approach can outperform more complex heuristic-based techniques. To improve weakly-supervised training, it is also possible to use standard low-level regularizers, as in Sec. 1.1, that leverage a large number of unlabeled pixels [65, 31, 61, 62, 42]. For example, [62] achieves the state-of-the-art using bilinear relaxation of the Potts model in (3)

$$E_{\text{Potts}}^{\text{bl}}(s(\theta)) = \sum_k (1 - s^k(\theta))^{\top} W s^k(\theta) \quad (7)$$

as an additional regularization loss over all (including unlabeled) pixels. For some $\nu > 0$, their continuous total loss

$$E = E_{\text{PCE}} + \nu E_{\text{Potts}}^{\text{bl}}. \quad (8)$$

More generally, standard regularization losses from low-level segmentation are commonly used in the context of segmentation networks. Such losses and their solvers are ubiquitous in weak-supervision techniques using seeds or boxes to generate fully-labeled *proposals* [28, 38]. Optimization of low-level regularizers is also common for network’s output post-processing, typically improving performance during testing [16]. Also, the corresponding low-level solvers can be directly integrated as solution-improving layers [65].

1.3. Weakly supervised semantic segmentation

Weak supervision for deep neural network semantic segmentation comes in many different forms, *e.g.* image-level tags [50, 49, 31], scribbles/clicks [38, 61, 62, 42], and bounding boxes [49, 28, 27]. These works employ a large variety of strategies to compensate for the lack of labels. The concept of *multiple instance learning* (MIL) naturally fits the weakly supervised setting. Since generic MIL methods produce small unsatisfactory segments, more specialized methods are needed. For example, methods [50, 27] impose constraints on the output of the neural network during learning. There are several segmentation-specific constraints, such as size bias, constraints on present labels, tightness [37], *etc.* [31, 62, 42] incorporate edge alignment constraints. Proposal generation methods [28, 38]

aim to generate/complete the ground truth to use fully-supervised learning. However, DNNs are vulnerable to errors in proposals. More robust approaches use EM [49] or ADMM [42] to iteratively correct errors in “proposals”.

1.4. Classic *trust region* optimization

Trust region is a general approximate iterative local optimization method [4] allowing to use approximations with good solvers when optimizing arbitrarily complex functions. To optimize $g(x)$, it solves sub-problem $\min_{\|x-x_t\| \leq \epsilon} \tilde{g}(x)$ where function $\tilde{g} \approx g$ is an approximation that can be “trusted” in some region $\|x - x_t\| \leq \epsilon$ around the current solution. If \tilde{g} is a linear expansion of g , this reduced to the gradient descent. More accurate higher-order approximations can be trusted over larger regions allowing larger steps. The sub-problem is often formulated as unconstrained Lagrangian optimization $\min_x \tilde{g}(x) + \lambda \|x - x_t\|$ where λ indirectly controls the step size.

1.5. Related optimization work and contributions

The first-order methods based on stochastic gradient descent dominate deep learning due to their simplicity, efficiency, and scalability. However, they often struggle to escape challenging features of the loss profile, *e.g.* “valleys”, as the gradients lack information on the curvature of the loss surface. Adam [29] combines gradients from many iterations to gather such curvature information. On the other hand, the second-order methods compute parameters update in the form $\Delta\theta = H^{-1} \nabla_{\theta} E(\mathbf{f}(\theta))$, *c.f.* (10), where H is the Hessian or its approximation. In neural networks, computing the Hessian is infeasible, so various approximations are used, *e.g.* diagonal or low-rank [2]. The efficient computation of Hessian-vector products is possible [52, 56]; while solving linear systems with Hessian is still challenging [60]. Another group of methods is based on employing Gaussian-Newton matrix and K-FAC approximations [43, 1, 3, 48].

Our approach is related to the proximal methods [44], in particular to the *proximal backpropagation* [22] and *penalty method* [9]. In these works, the “separation” of the gradient update into implicit layer-wise optimization problems is formulated as a gradient update of a certain energy function. Taylor *et al.* [63] use *ADMM* splitting approach to separate optimization over different layers in distributed fashion. These works focus on neural network parameter optimization replacing backpropagation altogether. In contrast to [9, 63, 22], we are primarily focused on optimization for complex loss functions in the context of the weakly supervised semantic segmentation, see Sec. 1.2, while others focus on replacing the backpropagation in the intermediate layers. Also, unlike us, these methods use the squared Euclidean norm in their proximal formulations. Chen and Teboulle [15] generalize the proximal methods to *Bregman divergences*, a more general class of functions which

includes both the Euclidean distance and KL-divergence. Nesterov in [45] uses the Euclidean norm with a higher power improving the convergence of the proximal method.

Our contribution are as follows:

- New trust region optimization for DNN segmentation integrating higher-order low-level solvers into training. Differentiability of the loss is not required as long as there is a good solver, discrete or continuous. The classic differentiation chain rule is replaced by the trust region chain rule in the context of backpropagation.
- The local optimization in trust region framework allows to use arbitrary metrics, instead of Euclidean distance implicit for the standard gradient descent. We discuss different metrics for the space of segmentations and motivate a robust version of KL-divergence.
- We show benefits of our optimization for regularization losses in weakly supervised DNN segmentation, compared to the gradient descent. We set new state-of-the-art results for weakly supervised segmentation with scribbles achieving consistently the best performance at all levels of supervision, *i.e.* from point-clicks to full-length scribbles.

2. Trust region for loss optimization

Backpropagation is the dominant method for optimizing network losses during training. It represents the gradient descent with respect to model parameters θ where the gradient's components are gradually accumulated using the classic *chain rule* while traversing the network layers starting from the output directly evaluated by the loss function.

Motivated by the use of hard-to-optimize regularization losses (Sec. 1.1) in the context of weakly-supervised segmentation (Sec. 1.2), we propose higher-order *trust region* approach to network training. While this general optimization approach can be developed for any steps of the backpropagation (*i.e.* chain rule) between internal layers, we focus on the very first step where the loss function is composed with the network output

$$\min_{\theta \in \mathbb{R}^m} E(\mathbf{f}(\theta)) \quad (9)$$

where some scalar loss function

$$E : \mathbb{R}^n \rightarrow \mathbb{R}^1$$

is defined over n -dimensional output of a network/model

$$\mathbf{f} : \mathbb{R}^m \rightarrow \mathbb{R}^n.$$

Since during training the network's input is limited to fixed examples, for simplicity we restrict the arguments of network function \mathbf{f} to its training parameters $\theta \in \mathbb{R}^m$. Also

note that, as a convention, this paper reserves the boldface font for vector functions (*e.g.* network model \mathbf{f}) and for matrix functions (*e.g.* model's Jacobian $\mathbf{J}_{\mathbf{f}}$).

The main technical ideas of the trust region approach to network optimization (9) in this section are fairly general. However, to be specific and without any loss of generality, this and (particularly) later sections can refer to the output of the network as *segmentation* so that

$$\mathbb{R}^n = \mathbb{R}^{N \times K}$$

where N is the number of image pixels and K is the number of distinct semantic classes. This is not essential.

Our general trust region approach to (9) can be seen as a higher-order extension of the classic chain rule for the composition $E \circ \mathbf{f}$ of the loss functions E and model \mathbf{f} . For the classic chain rule in the standard backpropagation procedure, it is critical that both E and \mathbf{f} are differentiable. In this case, the classic chain rule for the objective in (9) gives the following gradient descent update for parameters θ

$$\Delta\theta = -\alpha \nabla E^\top \mathbf{J}_{\mathbf{f}} \quad (10)$$

where $\Delta\theta \equiv \theta - \theta_t$ is an update of the model parameters from the current solution, α is the learning rate, ∇ is the gradient operator, and $\mathbf{J}_{\mathbf{f}}$ is the model's *Jacobian*

$$\mathbf{J}_{\mathbf{f}} := \begin{bmatrix} \frac{\partial \mathbf{f}_i}{\partial \theta^j} \end{bmatrix}.$$

We would like to rewrite the classic chain rule (10) in an equivalent form explicitly using a variable for segmentation $s \in \mathbb{R}^n$, which is an implicit (hidden) argument of the loss function E in (9). Obviously, equation (10) is equivalent to two separate updates for the segmentation $\Delta s \equiv s - s_t$ and for the model parameters $\Delta\theta \equiv \theta - \theta_t$

$$\Delta s = -\alpha \nabla E^\top \quad (11)$$

$$\Delta\theta = \Delta s \mathbf{J}_{\mathbf{f}} \quad (12)$$

where the gradient ∇E is computed at the current segmentation $s_t := \mathbf{f}(\theta_t)$. Note that $s \in \mathbb{R}^n$ represents points (*e.g.* segmentations) in the same space as the network output $\mathbf{f}(\theta) \in \mathbb{R}^n$, the two should be clearly distinguished in the discourse. We will refer to s as (explicit) *segmentation variable*, while $\mathbf{f}(\theta)$ is referred to as *segmentation output*.

The updates in (11) and (12) correspond to two distinct optimization sub-problems. Clearly, (11) is the gradient descent step for the loss $E(s)$ locally optimizing its linear Taylor approximation $\tilde{E}_{\text{linear}}(s) = E(s_t) + \nabla E^\top \Delta s$ over (explicit) segmentation variable $s \in B(s_t) \subset \mathbb{R}^n$ in a neighborhood (ball) around s_t

$$s_{t+1} = \arg \min_{s \in B(s_t)} \tilde{E}_{\text{linear}}(s). \quad (13)$$

While less obvious, it is easy to verify that θ -update in (12) is exactly the gradient descent step

$$\Delta\theta = -\frac{1}{2}\nabla_{\theta} \|s_{t+1} - \mathbf{f}(\theta)\|^2 \quad (14)$$

corresponding to optimization of the least-squares objective

$$\min_{\theta} \|s_{t+1} - \mathbf{f}(\theta)\|^2 \quad (15)$$

based on the solution $s_{t+1} \equiv \Delta s + \mathbf{f}(\theta_t)$ for problem (13).

Our trust region approach to network training (9) is motivated by the principled separation of the chain rule (10) into two sub-problems (13) and (15). Instead of the gradient descent, low-level optimization of the loss in (13) can leverage powerful higher-order solvers available for many popular loss functions, see Sec. 1.1. In particular, the majority of common robust loss functions for unsupervised or weakly-supervised computer vision problems are well-known to be problematic for the gradient descent. For example, their robustness (boundedness) leads to *vanishing gradients* and sensitivity to local minima. At the same time, the gradient descent can be left responsible for the least-squares optimization in (15). While it is still a hard problem due to size and non-convexity of the typical models $\mathbf{f}(\theta)$, at least the extra difficulties introduced by complex losses E can be removed into a different sub-problem.

Formally, our trust-region approach to training (9) generalizes our interpretation of the classic chain rule in sub-problems (13) and (15) as shown in iterative stages A, B:

STAGE A (low-level optimization)

$$s_{t+1} = \arg \min_s \tilde{E}(s) + \lambda d_A(s, \mathbf{f}(\theta_t)) \quad (16)$$

STAGE B (network parameters update)

$$\min_{\theta} \underbrace{d_B(s_{t+1}, \mathbf{f}(\theta))}_{\downarrow} \quad (17)$$

$$\Delta\theta = -\gamma \nabla_{\theta} d_B(s_{t+1}, \mathbf{f}(\theta)) \quad (18)$$

where \tilde{E} is some loss approximation, d_A and d_B are some distance/divergence measures. Instead of α in (11) and fixed weight $\frac{1}{2}$ in (14), the overall learning speed of our training procedure is controlled by two parameters: (A) scalar λ indirectly determining the step size from the current solution $s_t = \mathbf{f}(\theta_t)$ in (16), and (B) scalar γ defining the step size for the gradient descent in (18). While both λ and γ are important for the learning speed, we mostly refer to λ as a trust region parameter, while the term *learning rate* is reserved primarily for parameter γ in (18), as customary for the gradient descent step size in network optimization. Note that similarly to the gradient descent (10), stages A/B are iterated until convergence. While it is sensible to make several

B-steps (18) in a row, in general, it is not necessary to wait for convergence in sub-problem (17) before the next A-step.

Our formulation offers several significant generalizations of the classic chain rule. **First**, instead of the linear approximation (13) implied by the *gradient descent* (11), we target higher-order approximations of the loss \tilde{E} in (16). In some cases, the exact loss E could be used². The corresponding powerful low-level solvers for (16) are readily available for many types of useful robust losses, see Sec. 1.1. Note that for exact solvers when $\tilde{E} = E$, one may argue for $\lambda = 0$ allowing the network to learn from the best solutions for regularized loss E implying global optima in (9). However, such fixed proposals (Sec. 1.2) may result in overfitting to mistakes due to well-known biases/weaknesses in common regularizers. Constraining loss optimization (9) to the network output manifold in \mathbb{R}^n motivates $\lambda > 0$ in (16). More discussion is in [40, Sec.5].

Second, besides continuous/differentiable losses required by the standard backpropagation (chain rule), our trust region approach (stages A/B) allows training based on losses defined over discrete domains. There are several reasons why this extension is significant. For example, besides continuous solvers, optimization in (16) now can use a significantly larger pool of solvers including many powerful discrete/combinatorial methods. Moreover, this approach enables training of models with discrete decision functions, e.g. *step function* instead of *sigmoid*, or *hard-max* instead of the *soft-max*. This is further discussed in [40, Sec.5].

Third, the standard gradient descent (10) is implicitly defined over Euclidean metric, that manifests itself in our equations (13) and (15) via the local neighborhood topology (Euclidean ball B) and the least-squares objective (squared Euclidean distance). In contrast, when replacing ball $B(s_t)$ in (13) by the *trust region* term in (16), we explicitly define the trust region “shape” using function d_A . It could be any application-specific distance metric, quasi- or pseudo-metric, divergence, *etc.* Similarly, any appropriately motivated distance, distortion, or divergence function d_B in (17) can replace the least squares objective in (15).

On the negative side, our trust region formulation could be more expensive due to the computational costs of the low-level solvers in stage A. In practice, it is possible to amortize stage A over multiple iterations of stage B.

3. Robust metric for trust region

The choice of metrics d_A and d_B defining the shape of the trust region above is application dependent. In the case of segmentation, the output of a neural network is typically ob-

²Note that parameter λ in (16) controls two properties: the size of the trust region for approximation \tilde{E} , as well as the network’s training speed. While using exact loss $\tilde{E} = E$ implies that the trust region for such “approximation” should be the whole domain (*i.e.* $\lambda = 0$), the competing interest of limiting the training speed in (17) may require $\lambda > 0$.




Figure 1. The unknown true labeling Z corresponds to observed image I . The observed labeling Y is assumed to be generated from the true Z by a simple corruption model (20).

tained via the soft-max function. Hence, the space, in which the trust region operates, is the space of multiple categorical distributions over K categories: Δ_K^N .

Below, we generally discuss (robust) metrics over pairs of arbitrary probability distributions p, q in Δ_K^N . The goal of this section is to motivate our choice of metrics d_A and d_B in problems (16), (17) so that distribution p can be associated with the segmentation variable s , and distribution q can be associated with the network output $f(\theta)$. Besides this connection, the following discussion of metrics over probability distributions is independent of the context of networks.

Note, metrics d_A or d_B do not have to be proper distances for the purposes of trust region optimization. Instead, one may use any divergence measure defined on space Δ_K^N . Let us consider the Kullback–Leibler divergence:

$$\text{KL}(p||q) = \sum_{i=1}^N \sum_{l=1}^K p_i^l \log \frac{p_i^l}{q_i^l} = - \sum_{i=1}^N \sum_{l=1}^K p_i^l \log q_i^l - H(p)$$

where $p, q \in \Delta_K^N$, and p_i^l is the probability of pixel i to have label l , and $H(p)$ is the entropy of distribution p .

A practically important case is when the distribution p is degenerate or one-hot, *i.e.* for each pixel i there exists label y_i such that $p_i^{y_i} = 1$ and for any label $k \neq y_i$ probability $p_i^k = 0$. In that case $H(p) = 0$ and

$$\text{KL}(p||q) = \sum_i -\log q_i^{y_i}, \quad (19)$$

which is the cross-entropy or negative log-likelihood, a standard loss when q is the probability estimate outputted by a neural network. In the following we assume (19).

During the trust region procedure, intermediate solutions generated by a solver in (16) may have a noticeable amount of misclassified pixels. It is known that many standard losses for neural networks, including cross-entropy (19), can result in training sensitive to idiosyncrasies in the datasets including mistakes in the ground truth [23, 39, 21]. Therefore, a robust distance measure may be needed. Our experiments show that robustness is critical. We propose a simple error model depicted in graphical model in Fig. 1. Let random variable Y_i be the observed noisy label of pixel




Figure 2. Robust loss as function of logits x_i . There are $K = 2$ classes; the ground truth label is $y_i = 1$. If the current prediction q_i^1 is confident and does not coincide with y_i , see $x_i \ll 0$ or $q_i^1 \approx 0$ on the plot, robust loss (b) becomes flatter avoiding the over-penalize in case of mistakes in the ground truth. In contrast, standard cross-entropy (a) behaves linearly, which may be detrimental to learning if the ground truth is mistaken.




Figure 3. Classification accuracy on Fashion-MNIST dataset [64] with noisy labels using a network with two convolutional, two fully-connected layers and *robust loss* (22). The original labels were uniformly corrupted with probability $\frac{1}{2}$. The best accuracy is achieved at $\varepsilon = 0.4$, which is close to the actual noise level.

i and Z_i be its hidden true label. We assume that the probability of observing label l given true label k is

$$\Pr(Y_i = l | Z_i = k) = \begin{cases} 1 - \varepsilon, & l = k, \\ \frac{\varepsilon}{K-1}, & l \neq k, \end{cases} \quad (20)$$

where ε is called the *outlier probability* [36]. The probability of pixel i having label l given image I is

$$\begin{aligned} \Pr(Y_i = l | I) &= \sum_{z=1}^K \Pr(Y_i = l | Z_i = z) \Pr(Z_i = z | I) = \\ &= a + b \Pr(Z_i = l | I) \end{aligned} \quad (21)$$

where $a = \frac{\varepsilon}{K-1}$ and $b = 1 - K a$. The probability $\Pr(Z_i = z | I)$ is unknown and is replaced by probability estimate q_i^z yielding a robust version of divergence (19):

$$\sum_i -\log (a + b q_i^{y_i}). \quad (22)$$

Figure 2 compares cross-entropy (19) with robust loss (22).

Our robust cross-entropy (22) is related to a more general approach for classification [51, 59]. In [51], the corresponding robust cross-entropy (*forward correction*) is

$$\sum_i -\log \tilde{q}_i^{y_i} \quad (23)$$

where $\tilde{q}_i = T^\top q_i$, and q_i is the vector of probability estimates at pixel i , and $T = [T_{lk}]$ is the *noise transition matrix*: $T_{lk} = \Pr(Y = k | Z = l)$. The effect of different ε is shown in example in Fig. 3.

In practice, different pixels require different values of ε in (20). For example, in the scribble-based weakly supervised segmentation, the labels of seed pixels Ω_{seeds} are known for sure. So, $\varepsilon = 0$ for such pixels, and $\varepsilon > 0$ for all other pixels. Thus, the robust “metric” is

$$\text{KL}_{\varepsilon, \Omega_{\text{seeds}}}(p||q) = \sum_{i \notin \Omega_{\text{seeds}}} -\log(a + b q_i^{y_i}) + \sum_{i \in \Omega_{\text{seeds}}} -\log q_i^{y_i}. \quad (24)$$

In sum, we propose the following robust metrics for the trust region iterations (16) and (18):

$$\begin{cases} d_A(p, q) = \text{KL}(p||q), \\ d_B(p, q) = \text{KL}_{\varepsilon, \Omega_{\text{seeds}}}(p||q). \end{cases} \quad (25)$$

4. Results in weakly supervised segmentation

To validate our approach (16-18) we use standard efficient discrete solvers [7] for the loss

$$\tilde{E} = E_{\text{PCE}} + E_{\text{Potts}} \quad (26)$$

where $E_{\text{Potts}}(s) = \sum_{\{i,j\} \in \mathcal{N}} w_{ij} [s_i \neq s_j]$ is the second (regularization) term in standard low-level energy (1). In this case, optimization in (16) is limited to the corners of the simplex where E_{PCE} reduces to the hard constraints over the seeds. In (16-18) we use robust metrics (25). The overall method is summarized in Alg. 1.

One natural baseline for Alg. 1 is a standard method based on stochastic gradient descent (SGD) for regularized loss (8) proposed in [62], see Sec. 1.2. Indeed, $E_{\text{Potts}}^{\text{bl}}$ is a relaxation of E_{Potts} , as discussed in Sec. 1.1. Thus, (8) is a relaxation of (26). Alg. 1 with combinatorial solver for \tilde{E} in (26) can be seen as a *discrete* trust region approximation for (8). In general, our approach (16-18) allows other discrete or continuous solvers and/or other approximations \tilde{E} .

First, PCE-GD baseline is the standard SGD optimizing partial cross-entropy (6). It has been shown in [62, 61] that such approach outperforms more complex proposal (fake ground truth) generation methods such as [38]. Second, Grid-GD is the SGD over regularized loss (8) where the CRF neighbourhood is the standard 8-grid. Third, Dense-GD is the approach of [62] that uses the common fully-connected (dense) Potts CRF of [34].

We use the ScribbleSup [38] annotations for Pascal VOC 2012 [20] dataset. ScribbleSup supplies scribbles, *i.e.* a small subset of image pixels ($\approx 3\%$) is labeled while the vast majority of pixels is left unlabeled.

4.1. Implementation details

In all our experiments we used DeeplabV3+ [17] with MobileNetV2 [55] as a backbone model.

Pretraining: We use the standard ImageNet [19] pre-training of the backbone models. In addition, before the optimization via Grid-GD (7) and Grid-TR (16-18) starts, the DeeplabV3+ models are pretrained by the PCE loss (6).

Meta-parameters: We train 60 epochs. We tuned the learning rates for all methods on the val set. Other meta-parameters for competitive methods were set as in the corresponding papers/code. The learning rate is polynomial with power 0.9, momentum is 0.9, batch size is 12.

Grid-TR STAGE A (16): The low-level solver³ of the grid CRF is the α -expansion [8, 33, 6] with 8-grid neighbourhood system. The max number of α -expansion iterations is 5 achieving convergence in most cases. We restrict the set of labels to those present in the image. We amortize the STAGE A compute time by integrating it with data loading. The training is 1.3 times slower than Dense-GD.

Grid-TR STAGE B (18): To amortize the time consumed by the graph cuts, we perform $M = 5$ epochs of neural network weights updates (17) for each update of the segmentation variables (18). We use a global learning rate schedule spanning throughout iterations. See Alg. 1.

³GCOv3.0: <https://vision.cs.uwaterloo.ca/code/>

Algorithm 1: Robust Trust Region for Potts model

- 1 Initialize model \mathbf{f} using ImageNet pretraining ;
 - 2 Tune parameters θ of model \mathbf{f} by optimizing PCE-GD loss (6) ;
 - 3 Initialize γ with the base learning rate ;
 - 4 **repeat**
 - 5 **for each image in dataset do**
 - 6 compute segmentation variable s via (16) using metric d_A in (25) and loss (26);
 - 7 **end**
 - 8 **for M epochs do**
 - 9 **for each image (batch) in dataset do**
 - 10 update the network parameters θ using stochastic gradient descent for loss (17) with robust metric d_B in (25) ;
 - 11 update rate γ in accord with schedule;
 - 12 **end**
 - 13 **end**
 - 14 **until** required number of epochs is reached;
-

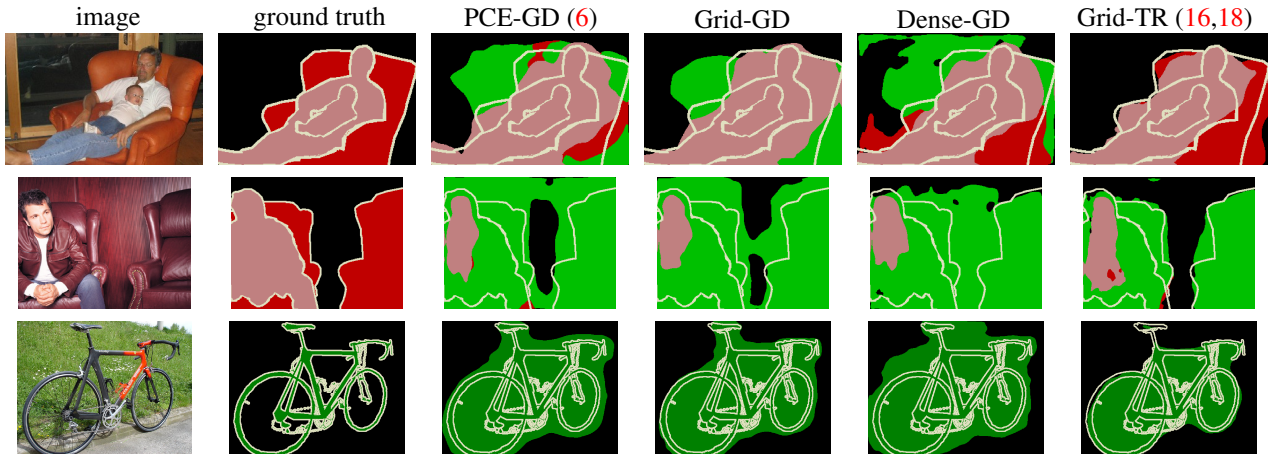


Figure 4. Examples of the full-scribble training results, see Tab. 1 and Figure 5. Note the better edge alignment of our Grid-TR.

scribble length	0	0.3	0.5	0.8	1
full supervision	0.70				
PCE-GD	0.50	0.57	0.59	0.61	0.61
Dense-GD	0.55	0.61	0.62	0.63	0.64
Grid-GD	0.54	0.60	0.62	0.64	0.64
Grid-TR (our)	0.57	0.63	0.64	0.66	0.67

Table 1. Results for ScribbleSup, see description in Figure 5.




Figure 6. The quality of segment boundary alignment. The networks were trained on the full-length scribbles.




Figure 5. Segmentation performance on the val set of ScribbleSup [38, 20] using DeeplabV3+ [17] with MobileNetV2 [55] backbone. The supervision level varies horizontally, with 1 corresponding to the full scribbles. Our “Grid-TR” outperforms other competitors for all scribble lengths and provides a new state-of-the-art.

4.2. Segmentation quality

The quantitative results of the weakly supervised training for semantic segmentation are presented in Figure 5 and Tab. 1. The results are presented with different levels of supervision varying from the clicks (denoted as length 0) to the full-length scribbles (denoted as length 1). Decreasing supervision results in degraded performance for all methods. We are interested to compare how different approaches

perform at different levels of supervision. Our Grid-TR outperforms all the competitors at each level of supervision.

The examples of images and results shown in Fig. 4 demonstrate the advantages of our method, particularly w.r.t. edge alignment. Quantitatively, we evaluate the accuracy of semantic boundaries using standard trimaps [30, 34, 16, 41]. A trimap corresponds to a narrow band around the ground truth segment boundaries of varying width. An accuracy measure, *e.g.* mIoU, is computed for pixels within each band. The results are shown in Fig. 6 where our approach demonstrates superior performance.

Acknowledgements

We thank Yaoliang Yu for the insightful discussion on related proximal methods and pointing out related literature. We also thank Vladimir Kolmogorov for suggesting prior studies of the tightness of the Potts model relaxations.

References

- [1] Jimmy Ba, Roger Grosse, and James Martens. Distributed second-order optimization using kronecker-factored approximations. 2017. [3](#)
- [2] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, August 2006. [3](#)
- [3] Aleksandar Botev, Hippolyt Ritter, and David Barber. Practical Gauss-Newton optimisation for deep learning. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 557–565, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. [3](#)
- [4] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. [3](#)
- [5] Yuri Boykov and Marie-Pierre Jolly. Interactive graph cuts for optimal boundary & region segmentation of objects in N-D images. In *ICCV*, volume I, pages 105–112, July 2001. [1](#), [2](#)
- [6] Yuri Boykov and Vladimir Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE transactions on pattern analysis and machine intelligence*, 26(9):1124–1137, 2004. [7](#)
- [7] Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 23(11):1222–1239, November 2001. [1](#), [2](#), [7](#)
- [8] Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 23(11):1222–1239, 2001. [7](#)
- [9] Miguel Carreira-Perpinan and Weiran Wang. Distributed optimization of deeply nested systems. In *Artificial Intelligence and Statistics*, pages 10–19. PMLR, 2014. [3](#)
- [10] Vicent Caselles, Ron Kimmel, and Guillermo Sapiro. Geodesic active contours. *International journal of computer vision*, 22(1):61–79, 1997. [2](#)
- [11] Antonin Chambolle, Daniel Cremers, and Thomas Pock. A convex approach to minimal partitions. *SIAM Journal on Imaging Sciences*, 5(4):1113–1158, 2012. [2](#)
- [12] Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011. [2](#)
- [13] Tony Chan, S Esedoglu, and M Nikolova. Algorithms for finding global minimizers of image segmentation and denoising models. *SIAM journal on applied mathematics*, 66(5):1632–1648, 2006. [2](#)
- [14] Tony F Chan and Luminita A Vese. Active contours without edges. *IEEE Transactions on image processing*, 10(2):266–277, 2001. [2](#)
- [15] Gong Chen and Marc Teboulle. Convergence analysis of a proximal-like minimization algorithm using bregman functions. *SIAM Journal on Optimization*, 3(3):538–543, 1993. [3](#)
- [16] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. [3](#), [8](#)
- [17] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. [7](#), [8](#)
- [18] Camille Couprie, Leo Grady, Laurent Najman, and Hugues Talbot. Power watershed: A unifying graph-based optimization framework. *IEEE transactions on pattern analysis and machine intelligence*, 33(7):1384–1399, 2010. [2](#)
- [19] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [7](#)
- [20] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *International journal of computer vision*, 111(1):98–136, 2015. [7](#), [8](#)
- [21] Benoît Fréney and Michel Verleysen. Classification in the presence of label noise: a survey. *IEEE transactions on neural networks and learning systems*, 25(5):845–869, 2013. [6](#)
- [22] Thomas Frerix, Thomas Möllenhoff, Michael Moeller, and Daniel Cremers. Proximal backpropagation. In *International Conference on Learning Representations*, 2018. [3](#)
- [23] Stuart Geman, Elie Bienenstock, and René Doursat. Neural networks and the bias/variance dilemma. *Neural computation*, 4(1):1–58, 1992. [6](#)
- [24] Lena Gorelick, Olga Veksler, Yuri Boykov, and Claudia Nieuwenhuis. Convexity shape prior for binary segmentation. *IEEE transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 39(2):258–271, February 2017. [2](#)
- [25] Hossam Isack, Lena Gorelick, Karin Ng, Olga Veksler, and Yuri Boykov. K-convexity shape priors for segmentation. In *European Conference on Computer Vision (ECCV)*, Munich, Germany, September 2018. [2](#)
- [26] Michael Kass, Andrew Witkin, and Demetri Terzopoulos. Snakes: Active contour models. *International journal of computer vision*, 1(4):321–331, 1988. [2](#)
- [27] Hoel Kervadec, Jose Dolz, Shanshan Wang, Eric Granger, and Ismail Ben Ayed. Bounding boxes for weakly supervised segmentation: Global constraints get close to full supervision. In *Proceedings of the Third Conference on Medical Imaging with Deep Learning*, volume 121 of *Proceedings of Machine Learning Research*, pages 365–381, Montreal, QC, Canada, 06–08 Jul 2020. PMLR. [3](#)
- [28] Anna Khoreva, Rodrigo Benenson, Jan Hosang, Matthias Hein, and Bernt Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 876–885, 2017. [3](#)
- [29] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [3](#)

- [30] Pushmeet Kohli, Philip HS Torr, et al. Robust higher order potentials for enforcing label consistency. *International Journal of Computer Vision*, 82(3):302–324, 2009. 2, 8
- [31] Alexander Kolesnikov and Christoph H Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *European Conference on Computer Vision*, pages 695–711. Springer, 2016. 3
- [32] Vladimir Kolmogorov. Convergent Tree-Reweighted Message Passing for Energy Minimization. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 28(10):1568–1583, October 2006. 2
- [33] Vladimir Kolmogorov and Ramin Zabih. What energy functions can be minimized via graph cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(2):147–159, 2004. 7
- [34] Philipp Krahenbuhl and Vladlen Koltun. Efficient inference in fully connected CRFs with Gaussian edge potentials. In *NIPS*, 2011. 1, 7, 8
- [35] M Pawan Kumar, Vladimir Kolmogorov, and Philip Torr. An analysis of convex relaxations for map estimation of discrete mrfs. *JMLR*, 2009. 2
- [36] Jan Larsen, L Nonboe, Mads Hintz-Madsen, and Lars Kai Hansen. Design of robust neural network classifiers. In *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP'98 (Cat. No. 98CH36181)*, volume 2, pages 1205–1208. IEEE, 1998. 6
- [37] Victor Lempitsky, Pushmeet Kohli, Carsten Rother, and Toby Sharp. Image segmentation with a bounding box prior. In *2009 IEEE 12th international conference on computer vision*, pages 277–284. IEEE, 2009. 3
- [38] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3159–3167, 2016. 3, 7, 8
- [39] Naresh Manwani and PS Sastry. Noise tolerance under risk minimization. *IEEE transactions on cybernetics*, 43(3):1146–1151, 2013. 6
- [40] Dmitrii Marin and Yuri Boykov. Robust trust region for weakly supervised segmentation (technical report). *CoRR (arXiv)*, abs/2104.01948, 2021. 5
- [41] Dmitrii Marin, Zijian He, Peter Vajda, Priyam Chatterjee, Sam Tsai, Fei Yang, and Yuri Boykov. Efficient segmentation: Learning downsampling near semantic boundaries. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2131–2141, 2019. 8
- [42] Dmitrii Marin, Meng Tang, Ismail Ben Ayed, and Yuri Boykov. Beyond gradient descent for regularized segmentation losses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10187–10196, 2019. 3
- [43] James Martens and Roger Grosse. Optimizing neural networks with kronecker-factored approximate curvature. In *International conference on machine learning*, pages 2408–2417. PMLR, 2015. 3
- [44] Bernard Martinet. Brief communication. Regularization of variational inequalities by successive approximations. *French journal of informatics and operational research. Red Series*, 4(R3):154–158, 1970. 3
- [45] Yurii E Nesterov. Inexact accelerated high-order proximal-point methods. Technical report, CORE, 2020. 4
- [46] Claudia Nieuwenhuis, Eno Toeppe, Lena Gorelick, Olga Veksler, and Yuri Boykov. Efficient squared curvature. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014. 2
- [47] Carl Olsson, Johannes Ulen, Yuri Boykov, and Vladimir Kolmogorov. Partial enumeration and curvature regularization. In *International Conference on Computer Vision (ICCV)*, Sydney, Australia, December 2013. 2
- [48] Kazuki Osawa, Yohei Tsuji, Yuichiro Ueno, Akira Naruse, Rio Yokota, and Satoshi Matsuoka. Large-scale distributed second-order optimization using kronecker-factored approximate curvature for deep convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12359–12367, 2019. 3
- [49] George Papandreou, Liang-Chieh Chen, Kevin P Murphy, and Alan L Yuille. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1742–1750, 2015. 3
- [50] Deepak Pathak, Philipp Krahenbuhl, and Trevor Darrell. Constrained convolutional neural networks for weakly supervised segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1796–1804, 2015. 3
- [51] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1944–1952, 2017. 7
- [52] Barak A Pearlmutter. Fast exact multiplication by the hessian. *Neural computation*, 6(1):147–160, 1994. 3
- [53] Pradeep Ravikumar and John Lafferty. Quadratic programming relaxations for metric labeling and markov random field map estimation. In *Proceedings of the 23rd international conference on Machine learning*, pages 737–744, 2006. 2
- [54] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. Grabcut - interactive foreground extraction using iterated graph cuts. In *ACM trans. on Graphics (SIGGRAPH)*, 2004. 1, 2
- [55] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. 7, 8
- [56] Nicol N Schraudolph. Fast curvature matrix-vector products for second-order gradient descent. *Neural computation*, 14(7):1723–1738, 2002. 3
- [57] Alexander Shekhovtsov, Pushmeet Kohli, and Carsten Rother. Curvature prior for mrf-based segmentation and shape inpainting. In *Joint DAGM (German Association for Pattern Recognition) and OAGM Symposium*, pages 41–51. Springer, 2012. 2

- [58] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22:888–905, 2000. 2
- [59] Sainbayar Sukhbaatar, Joan Bruna, Manohar Paluri, Lubomir Bourdev, and Rob Fergus. Training convolutional networks with noisy labels. In *3rd International Conference on Learning Representations, ICLR 2015*, 2015. 7
- [60] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *International conference on machine learning*, pages 1139–1147. PMLR, 2013. 3
- [61] Meng Tang, Abdelaziz Djelouah, Federico Perazzi, Yuri Boykov, and Christopher Schroers. Normalized Cut Loss for Weakly-supervised CNN Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 3, 7
- [62] Meng Tang, Federico Perazzi, Abdelaziz Djelouah, Ismail Ben Ayed, Christopher Schroers, and Yuri Boykov. On Regularized Losses for Weakly-supervised CNN Segmentation. In *European Conference on Computer Vision (ECCV)*, 2018. 3, 7
- [63] Gavin Taylor, Ryan Burmeister, Zheng Xu, Bharat Singh, Ankit Patel, and Tom Goldstein. Training neural networks without gradients: A scalable ADMM approach. In *International conference on machine learning*, pages 2722–2731, 2016. 3
- [64] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017. 6
- [65] Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip HS Torr. Conditional random fields as recurrent neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1529–1537, 2015. 3
- [66] Song Chun Zhu and Alan Yuille. Region competition: Unifying snakes, region growing, and Bayes/MDL for multiband image segmentation. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 18(9):884–900, Sept. 1996. 2