BoundaryNet: An Attentive Deep Network with Fast Marching Distance Maps for Semi-automatic Layout Annotation

by

Abhishek Trivedi, Ravi Kiran Sarvadevabhatla

in

International Conference on Document Analysis and Recognition (ICDAR) : 1 -16

Report No: IIIT/TR/2021/-1



Centre for Visual Information Technology International Institute of Information Technology Hyderabad - 500 032, INDIA September 2021

BoundaryNet: An Attentive Deep Network with Fast Marching Distance Maps for Semi-automatic Layout Annotation

Abhishek Trivedi^[0000-0002-6763-4716] and Ravi Kiran Sarvadevabhatla $(\boxtimes)^{[0000-0003-4134-1154]}$

Centre for Visual Information Technology (CVIT) International Institute of Information Technology, Hyderabad - 500032, INDIA http://ihdia.iiit.ac.in/BoundaryNet/ {abhishek.trivedi@research.,ravi.kiran@}iiit.ac.in

Abstract. Precise boundary annotations of image regions can be crucial for downstream applications which rely on region-class semantics. Some document collections contain densely laid out, highly irregular and overlapping multi-class region instances with large range in aspect ratio. Fully automatic boundary estimation approaches tend to be data intensive, cannot handle variable-sized images and produce sub-optimal results for aforementioned images. To address these issues, we propose BoundaryNet, a novel resizing-free approach for high-precision semi-automatic layout annotation. The variable-sized user selected region of interest is first processed by an attention-guided skip network. The network optimization is guided via Fast Marching distance maps to obtain a good quality initial boundary estimate and an associated feature representation. These outputs are processed by a Residual Graph Convolution Network optimized using Hausdorff loss to obtain the final region boundary. Results on a challenging image manuscript dataset demonstrate that BoundaryNet outperforms strong baselines and produces high-quality semantic region boundaries. Qualitatively, our approach generalizes across multiple document image datasets containing different script systems and layouts, all without additional fine-tuning. We integrate BoundaryNet into a document annotation system and show that it provides high annotation throughput compared to manual and fully automatic alternatives.

Keywords: document layout analysis · interactive · deep learning

1 Introduction

Document images exhibit incredible diversity in terms of language [40,36,6], content modality (printed [12,34], handwritten [35,32,16,4]), writing surfaces (paper, parchment [29], palm-leaf [30,17]), semantic elements such as text, tables, photos, graphics [6,34,43] and other such attributes. Within this variety, handwritten and historical documents pose the toughest challenges for tasks such as Optical Character Recognition (OCR) and document layout parsing.



Fig. 1: Compare the contours of semantic region instances for printed documents (top) [41] and historical document images (bottom). The latter are very diverse, often found damaged, contain densely laid out overlapping region instances (lines, holes) with large range in aspect ratios and high local curvature. These factors pose a challenge for region annotation.

In this work, we focus on historical documents. These documents form an important part of world's literary and cultural heritage. The mechanised process of machine printing imparts structure to modern-era paper documents. In contrast, historical document images are typically handwritten, unstructured and often contain information in dense, non-standard layouts (Fig. 1). Given the large diversity in language, script and non-textual elements in these documents, accurate spatial layout parsing can assist performance for other document-based tasks such as word-spotting [21], optical character recognition (OCR), style or content-based retrieval [39,44]. Despite the challenges posed by such images, a number of deep-learning based approaches have been proposed for fully automatic layout parsing [35,30,3,26]. However, a fundamental trade off exists between global processing and localized, compact nature of semantic document regions. For this reason, fully automatic approaches for documents with densely spaced, highly warped regions often exhibit false negatives or imprecise region boundaries. In practice, correction of predicted boundaries can be more burdensome than manual annotation itself.

Therefore, we propose an efficient semi-automatic approach for parsing images with dense, highly irregular layouts. The user selected bounding-box enclosing the region of interest serves as a weak supervisory input. Our proposed deep neural architecture, BoundaryNet, processes this input to generate precise region contours which require minimal to no manual post-processing.

Numerous approaches exist for weakly supervised bounding-box based semantic parsing of scene objects [2,37]. However, the spatial dimensions and aspect ratios of semantic regions in these datasets are less extreme compared to ones found in handwritten documents (Fig. 1). More recently, a number of approaches model the annotation task as an active contour problem by regressing boundary points on the region's contour [1,24,27,7,10]. However, the degree of curvature for document region contours tends to be larger compared to regular object datasets. The image content and associated boundaries are also distorted by the standard practice of resizing the image to a common height and width. For these reasons, existing approaches empirically tend to produce imprecise contours, especially for regions with high warp, extreme aspect ratio and multiple curvature points (as we shall see).

To address these shortcomings, we propose a two-stage approach (Sec. 3). In the first stage, the variable-sized input image is processed by an attentionbased fully convolutional network to obtain a region mask (Sec. 3.1). The region mask is morphologically processed to obtain an initial set of boundary points (Sec. 3.2). In the second stage, these boundary points are iteratively refined using a Residual Graph Convolutional Network to generate the final semantic region contour (Sec. 3.3). As we shall show, our design choices result in a highperforming system for accurate document region annotation.

Qualitatively and quantitatively, BoundaryNet outperforms a number of strong baselines for the task of accurate boundary generation (Sec. 4). BoundaryNet handles variable-sized images without resizing, in real-time, and generalizes across document image datasets with diverse languages, script systems and dense, overlapping region layouts (Sec. 4). Via end-to-end timing analysis, we showcase BoundaryNet's superior annotation throughput compared to manual and fully-automatic approaches (Sec. 5.2).

Source code, pre-trained models and associated documentation are available at http://ihdia.iiit.ac.in/BoundaryNet/.

2 Related Work

Annotating spatial regions is typically conducted in three major modes – manual, fully automatic and semi-automatic. The manual mode is obviously laborintensive and motivates the existence of the other two modes. Fully automatic approaches fall under the task categories of semantic segmentation [8] and instance segmentation [13]. These approaches work reasonably well for printed [6,43] and structured handwritten documents [3], but have been relatively less successful for historical manuscripts and other highly unstructured documents containing distorted, high-curvature regions [16,30].

Given the challenges with fully automatic approaches, semi-automatic variants operate on the so-called 'weak supervision' provided by human annotators. The weak supervision is typically provided as class label [38,11], scribbles [5,9,15] or bounding box [2] for the region of interest with the objective of predicting the underlying region's spatial support. This process is repeated for all image regions relevant to the annotation task. In our case, we assume box-based weak

supervision. Among box-based weakly supervised approaches, spatial support is typically predicted as a 2-D mask [2] or a boundary contour [1,24,27,7,10].

Contour-based approaches generally outperform mask-based counterparts and provide the flexibility of semi-automatic contour editing [1,7,27,10,24]. We employ a contour-based approach. However, unlike existing approaches, (i) BoundaryNet efficiently processes variable-sized images without need for resizing (ii) Boundary points are adaptively initialized from an inferred estimate of region mask instead of a fixed geometrical shape (iii) BoundaryNet utilizes skip connection based attentional guidance and boundary-aware distance maps to semantically guide region mask production (iv) BoundaryNet also produces region class label reducing post-processing annotation efforts. Broadly, our choices help deal with extreme aspect ratios and highly distorted region boundaries typically encountered in irregularly structured images.

3 BoundaryNet

Overview: Given the input bounding box, our objective is to obtain a precise contour of the enclosed semantic region (e.g. text line, picture, binding hole). BoundaryNet's processing pipeline consists of three stages – see Fig. 2. In the first stage, the bounding box image is processed by a Mask-CNN (MCNN) to obtain a good quality estimate of the underlying region's spatial mask (Sec. 3.1). Morphological and computational geometric procedures are used to sample contour points along the mask boundary (Sec. 3.2). A graph is constructed with contour points as nodes and edge connectivity defined by local neighborhoods of each contour point. The intermediate skip attention features from MCNN and contour point location are used to construct feature representations for each graph node. Finally, the feature-augmented contour graph is processed by a Graph Convolutional Network (Anchor GCN - Sec. 3.3) iteratively to obtain final set of contour points which define the predicted region boundary.

Semantic regions in documents are often characterized by extreme aspect ratio variations across region classes and uneven spatial distortion. In this context, it is important to note that BoundaryNet processes the input as-is without any resizing to arbitrarily fixed dimensions. This helps preserve crucial appearance detail. Next, we describe the components of BoundaryNet.

3.1 Mask-CNN (MCNN)

As the first step, the input image is processed by a backbone network ('Skip Attention Backbone' in Fig. 2). The backbone has U-Net style long-range skip connections with the important distinction that no spatial downsampling or upsampling is involved. This is done to preserve crucial boundary information. In the first part of the backbone, a series of residual blocks are used to obtain progressively refined feature representations (orange blocks). The second part of the backbone contains another series of blocks we refer to as Skip Attentional



Fig. 2: The architecture of BoundaryNet (top) and various sub-components (bottom). The variable-sized $H \times W$ input image is processed by Mask-CNN (MCNN) which predicts a region mask estimate and an associated region class (Sec. 3.1). The mask's boundary is determined using a contourization procedure (light brown) applied on the estimate from MCNN. M boundary points are sampled on the boundary (Sec. 3.2). A graph is constructed with the points as nodes and edge connectivity defined by $\leq k$ -hop neighborhoods of each point. The spatial coordinates of a boundary point location p = (x, y) and corresponding backbone skip attention features from MCNN f^r are used as node features for the boundary point. The feature-augmented contour graph $\mathbb{G} = (F, A)$ is iteratively processed by Anchor GCN (Sec. 3.3) to obtain the final output contour points defining the region boundary. Note that all filters in MCNN have a 3×3 spatial extent. The orange lock symbol on region classifier branch indicates that it is trained standalone, i.e. using pre-trained MCNN features.

Guidance (SAG). Each SAG block produces increasingly compressed (channelwise) feature representations of its input. To accomplish this feat without losing crucial low-level feature information, the output from immediate earlier SAG block is fused with skip features originating from a lower-level residual block layer (refer to 'Skip Attention Backbone' and its internal module diagrams in Fig. 2). This fusion is modulated via an attention mechanism (gray 'Attn Gating' block) [28].

The final set of features generated by skip-connection based attentional guidance (magenta) are provided to the 'Mask Decoder' network which outputs a region mask binary map. In addition, features from the last residual block (Res-128) are fed to 'Region Classifier' sub-network which predicts the associated region class. Since input regions have varying spatial dimensions, we use



Fig. 3: Contourization (Sec. 3.2): a - input image, b - thresholded initial estimate from MCNN, c - after area-based thresholding and joining centroids of largest connected components by an adaptive $m = \frac{H}{7}$ pixel-thick line where H is the height of the input image, d - after morphological closing, contour extraction, b-spline fitting and uniform point sampling, e - estimated contour (red) and ground-truth (blue) overlaid on input image.

adaptive average pooling [14] to ensure a fixed-dimensional fully connected layer output (see 'Region Classifier' in Fig. 2).

The input image is processed by an initial convolutional block with stride 2 filters before the resulting features are relayed to the backbone residual blocks. The spatial dimensions are restored via a transpose convolution upsampling within 'Mask Decoder' sub-network. These choices help keep the feature representations compact while minimizing the effect of downsampling.

3.2 Contourization

The pixel predictions in the output from 'Mask Decoder' branch are thresholded to obtain an initial estimate of the region mask. The result is morphologically processed, followed by the extraction of mask contour. A b-spline representation of mask contour is further computed to obtain a smoother representation. Mlocations are uniformly sampled along the b-spline contour curve to obtain the initial set of region boundary points. Figure 3 illustrates the various steps.

An advantage of the above procedure is that the set of mask-based boundary points serves as a reasonably accurate estimate of the target boundary. Therefore, it lowers the workload for the subsequent GCN stage which can focus on refining the boundary estimate.

3.3 Anchor GCN

The positional and appearance-based features of boundary points from the contourization stage (Sec. 3.2) are used to progressively refine the region's boundary estimate. For this, the boundary points are first assembled into a contour graph. The graph's connectivity is defined by $\leq k$ -hop neighbors for each contour point node. The node's s-dimensional feature representation is comprised of (i) the contour point 2-D coordinates p = (x, y) (ii) corresponding skip attention features from MCNN f^r - refer to 'GCN node feature' in Fig. 2 for a visual illustration.

The contour graph is represented in terms of two matrices - feature matrix Fand adjacency matrix A [18,42]. F is a $M \times s$ matrix where each row corresponds to the *s*-dimensional boundary point feature representation described previously. The $M \times M$ binary matrix A encodes the $\leq k$ -hop connectivity for each boundary point. Thus, we obtain the contour graph representation $\mathbb{G} = (F, A)$ (denoted 'Residual Graph' at bottom-right of Fig. 2). We briefly summarize GCNs next. **Graph Convolutional Network (GCN):** A GCN takes a graph \mathbb{G} as input and computes hierarchical feature representations at each node in the graph while retaining the original connectivity structure. The feature representation at the (i + 1)-th layer of the GCN is defined as $H_{i+1} = f(H_i, A)$ where H_i is a $p \times F_i$ matrix whose *j*-th row contains the *i*-th layer's feature representation for node indexed by j ($1 \leq j \leq N$). f (the so-called propagation rule) determines the manner in which node features of previous layer are aggregated to obtain current layer's feature representation. We use the following propagation rule [19]:

$$f(H_i, A) = \sigma(D^{\frac{-1}{2}} \widetilde{A} D^{\frac{-1}{2}} H_i W_i)$$
(1)

where $\widetilde{A} = A + I$ represents the adjacency matrix modified to include selfloops, D is a diagonal node-degree matrix (i.e. $D_{jj} = \sum_m \widetilde{A}_{jm}$) and W_i are the trainable weights for *i*-th layer. σ represents a non-linear activation function (ReLU in our case). Also, $H_0 = F$ (input feature matrix).

Res-GCN: The residual variant of GCN operates via an appropriate 'residual' modification to the GCN layer's feature representation and is defined as $H_{i+1} = f(H_i, A) + H_i$.

The input contour graph features are processed by a series of Res-GCN blocks [22] sandwiched between two GCN blocks. The Anchor GCN module culminates in a 2-dimensional fully connected layer whose output constitutes per-point displacements of the input boundary locations. To obtain the final boundary, we perform iterative refinement of predicted contour until the net displacements are negligibly small by re-using GCN's prediction for the starting estimate at each iteration [24].

3.4 Training and Inference

We train BoundaryNet in three phases.

First Phase: In this phase, we aim to obtain a good quality estimate of the boundary contour from MCNN. For this, the binary prediction from MCNN is optimized using per-pixel class-weighted binary focal loss [23]:

$$l_{BFL} = \alpha_c y (1-p)^{\gamma} \cdot \log p + (1-y)p^{\gamma} \cdot \log(1-p)$$
⁽²⁾

Table 1: Train, Validation and Test split for different region types.



Fig. 4: Some examples of region image (top) and color-coded versions of fast marching distance map (middle) and the attention map computed by final SAG block of BoundaryNet (bottom). The relatively larger values at the crucial boundary portions can be clearly seen in the attention map.

where $y \in \{0,1\}$ is ground-truth label, p is the corresponding pixel-level prediction, $\alpha_c = N_b/N_f$ is the ratio of background to foreground (region mask) pixel counts and γ is the so-called focusing hyperparameter in focal loss. The class-weighting ensures balanced optimization for background and foreground, indirectly aiding contour estimation. The focal loss encourages the optimization to focus on the harder-to-classify pixels.

To boost the precision of estimate in a more boundary-aware manner, we first construct a distance map using a variant of the Fast Marching method [33]. The procedure assigns a distance of 0 to boundary pixels and progressively higher values to pixels based on contours generated by iterative erosion and dilation of ground-truth region mask (see Figure 4. The distance map is inverted by subtracting each entry from the global maximum within the map. Thus, the highest weights are assigned to boundary pixels, with the next highest set of values for pixels immediately adjacent to the boundary. The inverted distance map is then normalized ([0, 1]) to obtain the final map Ψ . The class-weighted binary focal loss spatial map \mathcal{L}_{BFL} is constituted from per-pixel losses l_{BFL} (Eq. 2) and further weighted by Ψ as follows:

$$\mathcal{L}_{FM} = (1 + \Psi) \odot \mathcal{L}_{BFL} \tag{3}$$

where \odot stands for the Hadamard product. The above formulation is preferred to solely weighting \mathcal{L}_{BFL} with Ψ to mitigate the vanishing gradient issue.

Second Phase: In this phase of training, MCNN's weights are frozen and the estimate of region mask is obtained as described previously (Sec. 3.1). The con-

tour graph constructed from points on region mask boundary (Sec. 3.2) is fed to Anchor GCN. The output nodes from Anchor GCN are interpolated $10 \times$ through grid sampling. This ensures maximum optimal shifts towards ground-truth contour for original graph nodes and avoids graph distortion.

Let \mathcal{G} be the set of points (x-y locations) in ground-truth contour and \mathcal{B} , the point set predicted by Anchor GCN. Let E_1 be the list of minimum Euclidean distances calculated per ground-truth point $g_i \in \mathcal{G}$ to a point in \mathcal{B} , i.e. $e_i = \min_j || g_i - b_j ||, e_i \in E_1, b_j \in \mathcal{B}$. Let E_2 be a similar list obtained by flipping the roles of ground-truth and predicted point sets. The Hausdorff Distance loss [31] for optimizing Anchor GCN is defined as:

$$L_C(E_1, E_2) = 0.5(\sum_i e_i + \sum_j e_j)$$
(4)

where $e_j \in E_2$.

Third Phase: In this phase, we jointly fine-tune the parameters of both MCNN and Anchor GCN in an end-to-end manner. The final optimization is performed by minimizing L_{FT} loss defined as: $L_{FT} = L_C + \lambda \mathcal{L}_{FM}$. As we shall see, the end-to-end optimization is crucial for improved performance (Table 3).

The region classification sub-branch is optimized using categorical crossentropy loss (\mathcal{L}_{CE}^r) after all the phases mentioned above. During this process, the backbone is considered as a pre-trained feature extractor, i.e. backpropagation is not performed on MCNN backbone's weights.

3.5 Implementation Details

MCNN: The implementation details of MCNN can be found in Fig. 2. The input $H \times W \times 3$ RGB image is processed by MCNN to generate a corresponding $H \times W$ region mask representation (magenta branch in Fig. 2) and a region class prediction (orange branch) determined from the final 8-way softmax layer of the branch. In addition, the outputs from the SAG blocks are concatenated and the resulting $\frac{H}{2} \times \frac{W}{2} \times 120$ output (shown at the end of dotted green line in Fig. 2) is used to determine the node feature representations f^r used in the downstream Anchor GCN module.

For MCNN training, the focal loss (Eq. 2) is disabled at the beginning, i.e. $\gamma = 0$. The batch size is set to 1 with an initial learning rate of $3e^{-5}$. A customized variant of Stochastic Gradient Descent with Restart [25] is conducted. Two fresh restarts are performed by increasing learning rate $5 \times$ for 3 epochs and dropping it back to counter potential loss saturation. The focal loss is invoked with $\gamma = 2$ when \mathcal{L}_{FM} (Eq. 3) starts to plateau. At this stage, the learning rate is set to decay by 0.5 every 7 epochs.

Contourization: The region formed by pixels labelled as region interior in MCNN's output is morphologically closed using a 3×3 disk structuring element. Major regions are extracted using area-based thresholding. The final region interior mask is obtained by connecting all the major sub-regions through their

Table 2: Region-wise average and overall Hausdorff Distance (HD) for different baselines and BoundaryNet on Indiscapes dataset.

	$ $ HD \downarrow	Hole	Line Segment	Degradation	Character	Picture	Decorator	Library Marker	Boundary Line
BoundaryNet	17.33	6.95	20.37	10.15	7.58	51.58	20.17	16.42	5.45
Polygon-RNN++[1]	30.06	5.59	66.03	7.74	5.11	105.99	25.11	9.97	15.01
Curve-GCN[24]	39.87	8.62	142.46	14.55	10.25	68.64	32.11	19.51	22.85
DACN[7]	41.21	8.48	105.61	14.10	11.42	91.18	26.55	22.24	50.16
DSAC[27]	54.06	14.34	237.46	10.40	8.27	65.81	39.36	23.34	33.53

centroids. A b-spline curve is fit to the boundary of the resulting region and M = 200 boundary points are uniformly sampled along the curve - this process is depicted in Fig. 3.

Anchor GCN: Each boundary point's 122-dimensional node feature is obtained by concatenating the 120-dimensional feature column (f^r in Fig. 2) and the point's 2-D coordinates p = (x, y) (normalized to a $[0, 1] \times [0, 1]$ grid). Each contour point is connected to its 20 nearest sequential neighbors in the contour graph, ten on each side along the contour (see 'Mask Contour' in Fig. 2), i.e. maximum hop factor k = 10. The graph representation is processed by two GCN and six residual GCN layers (see 'Residual GCN' in Fig. 2 for architectural details). The resulting features are processed by a fully connected layer to produce 2-D residuals for each of the boundary points. The iterative refinement of boundary points is performed two times. During training, the batch size is set to 1 with a learning rate of $1e^{-3}$.

End-to-end framework: For joint optimization, the batch size set to 1 with a relatively lower learning rate of $1e^{-5}$. Weighting coefficient λ (in Eq. 4) is set to 200.

Throughout the training phases and for loss computation, the predicted points and ground-truth are scaled to a unit normalized $([0, 1] \times [0, 1])$ canvas. Also, to ensure uniform coverage of all region classes, we perform class-frequency based mini-batch resampling and utilize the resultant sequences for all phases of training.

4 Experimental Setup

Performance Measure: As performance measure, we use the Hausdorff Distance (HD) [20] between the predicted contour and its ground-truth counterpart (Sec. 3.5). Note that smaller the HD, the better is the boundary prediction. The per-region HD is obtained as the average over the HD of associated region instances.

For all the models, we use performance on the validation set to determine the optimal hyperparameters and determine architectural choices. Subsequently, we optimize the models on the combined training and validation splits and conduct a one-time evaluation on the test split.

Baselines: To perform a comparative evaluation of BoundaryNet, we include multiple state-of-the-art semi-automatic annotation approaches - DSAC [27],



Fig. 5: Qualitative comparison of BoundaryNet with baselines on sample test images from Indiscapes dataset. For each region, the ground-truth contour is outlined in white. The IoU score is also mentioned for reference (see Sec. 5).

Polygon-RNN++ [1], Curve-GCN [24] and DACN [7]. These approaches exhibit impressive performance for annotating semantic regions in street-view dataset and for overhead satellite imagery. However, directly fine-tuning the baselines resulted in bad performance due to the relatively fewer annotation nodes regression and domain gap between document images and imagery (street-view, satellite) for which the baselines were designed. Therefore, we use the original approaches as a guideline and train modified versions of the baseline deep networks.

Evaluation dataset: For training and primary evaluation, we use Indiscapes [30], a challenging historical document dataset of handwritten manuscript images. It contains 526 diverse document images containing 9507 regions spanning the following categories: Holes, Line Segments, Physical Degradation, Character Component, Picture, Decorator, Library Marker, Boundary Line, Page Boundary (omitted for our evaluation). Details of the training, validation and test splits can be viewed in Table 1.

5 Results

5.1 Indiscapes

Quantitative Baseline Comparison: As Table 2 shows, BoundaryNet outperforms other baselines by a significant margin in terms of overall Hausdorff Distance (HD). Considering that images in the test set have widths as large as 6800 pixels, the results indicate a high degree of precision for obtained contours. The performance of BoundaryNet is slightly lower than the best on regions such



Fig. 6: Semantic region boundaries predicted by BoundaryNet. The colors distinguish instances – they are *not* region labels (written in shorthand alongside the regions). The dotted line separates Indiscapes dataset images (top) and those from other document collections (bottom). Note: BoundaryNet has been trained only on Indiscapes.

as 'Holes', 'Library Marker' and 'Degradation' due to the filtering induced by the GCN. However, notice that the performance for region present most frequently - 'Line Segment' - is markedly better than other baselines.

Qualitative Baseline Comparison: The performance of BoundaryNet and top three baseline performers for sample test images can be viewed in Figure 5. In addition to HD, we also mention the IoU score. As the results demonstrate, HD is more suited than IoU for standalone and comparative performance assessment of boundary precision. The reason is that IoU is an area-centric measure, suited for annotating objects in terms of their rigid edges (e.g. objects found in real-world scenery). As example B in Fig. 5 shows, a boundary estimate which fails to enclose the semantic content of the region properly can still have a high IoU. In contrast, semantic regions found in documents, especially character lines, typically elicit annotations which aim to minimally enclose the region's semantic content in a less rigid manner. Therefore, a contour-centric measure which penalizes boundary deviations is more suitable.

Table 3: Performance for ablative variants of BoundaryNet. The + refers to MCNN's output being fed to mentioned ablative variants of AGCN.

Component	Ablation type	Default Configuration in BoundaryNet	$\mathtt{HD}\downarrow$
MCNN	Max Residual Channels=64	Max Residual Channels=128	21.77
MCNN	No Focal Loss	Focal Loss	22.98
MCNN	No Fast Marching weights Penalization	Fast Marching weights Penalization	23.96
MCNN	Normal skip connection, no attention gating	Skip connection with attention gating	28.27
MCNN	No AGCN	AGCN	19.17
+AGCN	\leq 5-hop neighborhood	≤ 10 -hop neighborhood	19.26
+AGCN	≤ 15 -hop neighborhood	≤ 10 -hop neighborhood	20.82
+AGCN	$1 \times$ spline interpolation	$10 \times$ interpolation	20.48
+AGCN	1 iteration	2 iterations	19.31
+AGCN	100 graph nodes	200 graph nodes	20.37
+AGCN	300 graph nodes	200 graph nodes	19.98
+AGCN	Node features: backbone only $f^r(x, y)$	$f^r(x,y), (x,y)$	20.50
+Fine-Tuning	g No end-to-end finetuning	End to end finetuning	18.79
BoundaryNet	_	original	17.33

Qualitative Results (Image-level): Examples of document images with BoundaryNet predictions overlaid can be seen in Figure 6. The images above the dotted line are from the Indiscapes dataset. The documents are characterized by dense layouts, degraded quality (first image), ultra wide character lines (second image). Despite this, BoundaryNet provides accurate annotation boundaries. Note that BoundaryNet also outputs region labels. This results in amortized time and labor savings for the annotator since region label need not be provided separately. Region Classifier performance can be seen in Figure 7 (left).

Performance on other document collections: To determine its general utility, we used BoundaryNet for semi-automatic annotation of documents from other historical manuscript datasets (South-East Asian palm leaf manuscripts, Arabic and Hebrew documents). The results can be viewed in the images below the dotted line in Figure 6. Despite not being trained on images from these collections, it can be seen that BoundaryNet provides accurate region annotations.

Ablations: To determine the contribution of various architectural components, we examined lesioned variants of BoundaryNet for ablation analysis. The results can be viewed in Table 3. As can be seen, the choices related to the MCNN's loss function, presence of error penalizing distance maps, number of points sampled on mask boundary estimate, spline interpolation, all impact performance in a significant manner.

5.2 Timing Analysis

To determine BoundaryNet utility in a practical setting, we obtained documentlevel annotations for test set images from Indiscapes dataset. The annotations for each image were sourced using an in-house document annotation system in three distinct modes: Manual Mode (hand-drawn contour generation and region labelling), Fully Automatic Mode (using an existing instance segmenta-



Fig. 7: (left) Confusion Matrix from Region classifier branch, (right) Documentlevel end-to-end annotation duration distribution for various approaches depicted as a violin plot(the white dot represents mean duration - see Sec. 5.2).

tion approach [30] with post-correction using the annotation system) and Semi-Automatic Mode (manual input of region bounding boxes which are subsequently sent to BoundaryNet, followed by post-correction). For each mode, we recorded the end-to-end annotation time at per-document level, including manual correction time. The distribution of annotation times for the three modes can be seen in Figure 7 (right). As can be seen, the annotation durations for the BoundaryNetbased approach are much smaller compared to the other approaches, despite BoundaryNet being a semi-automatic approach. This is due to the superior quality contours generated by BoundaryNet which minimize post-inference manual correction burden.

6 Conclusion

In this paper, we propose BoundaryNet, a novel architecture for semi-automatic layout annotation. The advantages of our method include (i) the ability to process variable dimension input images (ii) accommodating large variation in aspect ratio without affecting performance (iii) adaptive boundary estimate refinement. We demonstrate the efficacy of BoundaryNet on a diverse and challenging document image dataset where it outperforms competitive baselines. Finally, we show that BoundaryNet readily generalizes to a variety of other historical document datasets containing dense and uneven layouts. Going ahead, we plan to explore the possibility of incorporating BoundaryNet into popular instance segmentation frameworks in an end-to-end manner.

References

1. Acuna, D., et al.: Efficient interactive annotation of segmentation datasets with polygon-rnn++. In: CVPR. pp. 859–868 (2018) 3, 4, 10, 11

- Bonechi, S., Andreini, P., et al.: Coco_ts dataset: Pixel–level annotations based on weak supervision for scene text segmentation. In: ICANN. pp. 238–250 (2019) 2, 3, 4
- 3. Breuel, T.M.: Robust, simple page segmentation using hybrid convolutional mdlstm networks. ICDAR 01, 733–740 (2017) 2, 3
- Buss, J.F., Rosenberg, A.L., Knott, J.D.: Vertex types in book-embeddings. Tech. rep., Amherst, MA, USA (1987) 1
- Can, Y.B., et al.: Learning to segment medical images with scribble-supervision alone. In: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support, pp. 236–244. Springer (2018) 3
- Clausner, C., Antonacopoulos, A., Derrick, T., Pletschacher, S.: Icdar2019 competition on recognition of early indian printed documents-reid2019. In: ICDAR. pp. 1527–1532 (2019) 1, 3
- 7. Dong, Z., Zhang, R., Shao, X.: Automatic annotation and segmentation of object instances with deep active curve network. IEEE Access (2019) 3, 4, 10, 11
- Fu, J., Liu, J., Wang, Y., Zhou, J., Wang, C., Lu, H.: Stacked deconvolutional network for semantic segmentation. IEEE TIP (2019) 3
- Garz, A., Seuret, M., Simistira, F., Fischer, A., Ingold, R.: Creating ground truth for historical manuscripts with document graphs and scribbling interaction. In: DAS. pp. 126–131 (2016) 3
- 10. Gur, S., Shaharabany, T., Wolf, L.: End to end trainable active contours via differentiable rendering. In: ICLR (2020) 3, 4
- Gurjar, N., Sudholt, S., Fink, G.A.: Learning deep representations for word spotting under weak supervision. In: DAS. pp. 7–12 (2018) 3
- 12. Harley, A.W., Ufkes, A., Derpanis, K.G.: Evaluation of deep convolutional nets for document image classification and retrieval. In: ICDAR (2015) 1
- 13. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: ICCV (2017) 3
- He, K., Zhang, X., Ren, S., Sun, J.: Spatial pyramid pooling in deep convolutional networks for visual recognition. IEEE TPAMI 37(9), 1904–1916 (2015) 6
- Kassis, M., El-Sana, J.: Scribble based interactive page layout segmentation using gabor filter. In: ICFHR. pp. 13–18 (2016) 3
- Kassis, M., Abdalhaleem, A., Droby, A., Alaasam, R., El-Sana, J.: Vml-hd: The historical arabic documents dataset for recognition systems. In: 1st Intl. Wksp on Arabic Script Analysis and Recognition. pp. 11–14. IEEE (2017) 1, 3
- 17. Kesiman, M.W.A., et al.: Benchmarking of document image analysis tasks for palm leaf manuscripts from southeast asia. Journal of Imaging **4**(2), 43 (2018) 1
- Kipf, T.N., Welling, M.: Variational graph auto-encoders. arXiv preprint arXiv:1611.07308 (2016) 7
- Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: ICLR (2017) 7
- Klette, R., Rosenfeld, A. (eds.): Digital Geometry. The Morgan Kaufmann Series in Computer Graphics, Morgan Kaufmann, San Francisco (2004) 10
- Lais Wiggers, K., de Souza Britto Junior, A., Lameiras Koerich, A., Heutte, L., Soares de Oliveira, L.E.: Deep Learning Approaches for Image Retrieval and Pattern Spotting in Ancient Documents. arXiv e-prints (2019) 2
- Li, G., Muller, M., Thabet, A., Ghanem, B.: Deepgcns: Can gcns go as deep as cnns? In: ICCV. pp. 9267–9276 (2019) 7
- Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV. pp. 2980–2988 (2017) 7
- 24. Ling, H., Gao, J., Kar, A., Chen, W., Fidler, S.: Fast interactive object annotation with Curve-GCN. In: CVPR. pp. 5257–5266 (2019) 3, 4, 7, 10, 11

- 16 Trivedi and Sarvadevabhatla
- Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983 (2016) 9
- Ma, L., Long, C., Duan, L., Zhang, X., Li, Y., Zhao, Q.: Segmentation and recognition for historical tibetan document images. IEEE Access (2020) 2
- Marcos, D., Tuia, D., et al.: Learning deep structured active contours end-to-end. In: CVPR. pp. 8877–8885 (2018) 3, 4, 10
- 28. Oktay, O., et al.: Attention u-net: Learning where to look for the pancreas. In: Medical Imaging with Deep Learning (2018) 5
- Pal, K., Terras, M., Weyrich, T.: 3d reconstruction for damaged documents: imaging of the great parchment book. In: Frinken, V., Barrett, B., Manmatha, R., Märgner, V. (eds.) HIP@ICDAR 2013. pp. 14–21. ACM (2013) 1
- Prusty, A., Aitha, S., Trivedi, A., Sarvadevabhatla, R.K.: Indiscapes: Instance segmentation networks for layout parsing of historical indic manuscripts. In: ICDAR. pp. 999–1006 (2019) 1, 2, 3, 11, 14
- Ribera, J., Güera, D., Chen, Y., Delp, E.J.: Locating objects without bounding boxes. CVPR (June 2019) 9
- 32. Saini, R., Dobson, D., et al.: Icdar 2019 historical document reading challenge on large structured chinese family records. In: ICDAR. pp. 1499–1504 (2019) 1
- Sethian, J.A.: A fast marching level set method for monotonically advancing fronts. PNAS 93(4), 1591–1595 (1996) 8
- Shahab, A., Shafait, F., et al.: An open approach towards benchmarking of table structure recognition systems. In: DAS. pp. 113–120 (2010) 1
- Simistira, F., Seuret, M., Eichenberger, N., Garz, A., Liwicki, M., Ingold, R.: Divahisdb: A precisely annotated large dataset of challenging medieval manuscripts. In: ICFHR. pp. 471–476 (2016) 1, 2
- Slimane, F., Ingold, R., Kanoun, S., Alimi, A.M., Hennebert, J.: Database and evaluation protocols for arabic printed text recognition. DIUF-University of Fribourg-Switzerland (2009) 1
- Song, C., Huang, Y., Ouyang, W., Wang, L.: Box-driven class-wise region masking and filling rate guided loss for weakly supervised semantic segmentation. In: CVPR. pp. 3136–3145 (2019) 2
- Tang, M., Perazzi, F., et al.: On regularized losses for weakly-supervised cnn segmentation. In: ECCV. pp. 507–522 (2018) 3
- Wiggers, K.L., Junior, A.S.B., Koerich, A.L., Heutte, L., de Oliveira, L.E.S.: Deep learning approaches for image retrieval and pattern spotting in ancient documents. ArXiv abs/1907.09404 (2019) 2
- 40. Yalniz, I.Z., Manmatha, R.: A fast alignment scheme for automatic ocr evaluation of books. In: ICDAR. pp. 754–758 (2011) 1
- Yang, X., et al.: Learning to extract semantic structure from documents using multimodal fully convolutional neural networks. In: CVPR. pp. 5315–5324 (2017)
 2
- 42. Zhang, L., Song, H., Lu, H.: Graph node-feature convolution for representation learning. arXiv preprint arXiv:1812.00086 (2018) 7
- Zhong, X., Tang, J., Jimeno Yepes, A.: Publaynet: Largest dataset ever for document layout analysis. In: ICDAR. pp. 1015–1022 (2019) 1, 3
- 44. Zhou, J., Guo, B., Zheng, Y.: Document image retrieval based on convolutional neural network. In: Advances in Intelligent Information Hiding and Multimedia Signal Processing. pp. 221–229 (2020) 2