

Seq-UPS: Sequential Uncertainty-aware Pseudo-label Selection for Semi-Supervised Text Recognition

Gaurav Patel Jan Allebach Qiang Qiu

School of Electrical and Computer Engineering, Purdue University, USA

{gpatel10, allebach, qqiu}@purdue.edu

Abstract

This paper looks at semi-supervised learning (SSL) for image-based text recognition. One of the most popular SSL approaches is pseudo-labeling (PL). PL approaches assign labels to unlabeled data before re-training the model with a combination of labeled and pseudo-labeled data. However, PL methods are severely degraded by noise and are prone to over-fitting to noisy labels, due to the inclusion of erroneous high confidence pseudo-labels generated from poorly calibrated models, thus, rendering threshold-based selection ineffective. Moreover, the combinatorial complexity of the hypothesis space and the error accumulation due to multiple incorrect autoregressive steps posit pseudo-labeling challenging for sequence models. To this end, we propose a pseudo-label generation and an uncertainty-based data selection framework for semi-supervised text recognition. We first use Beam-Search inference to yield highly probable hypotheses to assign pseudo-labels to the unlabelled examples. Then we adopt an ensemble of models, sampled by applying dropout, to obtain a robust estimate of the uncertainty associated with the prediction, considering both the character-level and word-level predictive distribution to select good quality pseudo-labels. Extensive experiments on several benchmark handwriting and scene-text datasets show that our method outperforms the baseline approaches and the previous state-of-the-art semi-supervised text-recognition methods.

1. Introduction

Text recognition has garnered a great deal of attention in recent times [35], owing primarily to its commercial applications. Since the introduction of deep learning, great strides [6, 9, 14, 31, 33, 34, 49, 50, 54, 55, 62] have been made in recognition accuracy on various publicly available benchmark datasets [27, 28, 29, 36, 41, 42, 45, 46, 56]. These models, however, are heavily reliant on a large amount of labeled data with complete character-sequence as labels, which is laborious to obtain. Aside from fully-supervised text recognition, very few attempts have been made to utilize unlabelled data samples to improve the model’s performance [1, 26, 64].




Figure 1: An overview of Beam-Search inference (beam width = 2) on recognizing an unlabeled text-image and populating the hypotheses set. The framework considers all the accumulated hypotheses to approximate the *total uncertainty* (\mathcal{U}_{total}) by importance-sampling.

Semi-supervised learning paradigms have been developed to address the preceding issues, and primarily pseudo-label-based semi-supervised learning (PL-SSL) methods have caught much attention. In a PL-SSL configuration, a smaller labeled set is used to train an initial seed model and then applied to a larger amount of unlabeled data to generate hypotheses. Furthermore, the unlabeled data along with its most reliable hypothesis as the label is combined with the training data for re-training, this methodology utilizes both the labeled and unlabelled data-points to re-train the complete model, allowing the entire network to exploit the latent knowledge from unlabelled data-points as well. However, on the other hand, PL-SSL is sensitive to the quality of the selected pseudo-labels and suffers due to the inclusion of erroneous highly confident pseudo-labels generated from poorly calibrated models, resulting in noisy training [47]. Moreover, for image-based text recognition, that requires predictions of characters at each time-step for each input image, pseudo-labeling is much more challenging due to the combinatorial vastness of the hypotheses space and the fact that a single incorrect character prediction renders the entire

predicted sequence false. Additionally, in the PL-SSL setup, handling erroneous predictions from the model trained with a relatively small amount of labeled data, and being able to exclude them in the beginning of the training cycle is highly essential. Therefore, correct hypotheses generation and selection are of massive importance in such a framework.

This work proposes a reliable hypotheses generation and selection method for PL-SSL for image-based text recognition. We suggest a way to estimate the uncertainty associated with the prediction of character sequences for an input image that gives a firm estimate of the reliability of the pseudo-label for an unlabelled data sample and then based on the estimated uncertainty, select the examples which have a low likelihood of an incorrectly assigned pseudo-label. Our methodology stems from the two primary observations that suggest (a) for pseudo-labeling based SSL schemes, choosing predictions with low uncertainty reduces the effect of poor calibration, thus improving generalization [47] and (b) a high positive correlation exists between the predictive uncertainty and the token error rate, for a deep neural network-based language model, suggesting if the model produces a high uncertainty for an input image, it's highly likely that the prediction used as the pseudo-label is incorrect [53].

Nevertheless, the majority of the unsupervised uncertainty estimation methods have concentrated on conventional unstructured prediction tasks, such as image classification and segmentation. Uncertainty estimation of an input image for text recognition, inherently a sequence prediction task, is highly non-trivial and poses various challenges [39]: (a) Recognition models do not directly generate a distribution over an infinite set of variable-length sequences, and (b) an auto-regressive sequence prediction task, such as text recognition, does not have a fixed hypotheses set; therefore, debarring expectation computation on the same.

To circumvent these challenges we use Beam-Search inference (Figure 1) to yield high probability hypotheses for each unlabelled data-point on the seed model trained with the given labeled data. Moreover, the hypotheses obtained are used to approximate the predictive distribution and obtain expectations over the set. We term this process as *deterministic inference*, which generates definite and a distinct hypotheses set for each image that aid in approximating the predictive distribution. Furthermore, to compute the uncertainty associated with an input, we take a *Bayesian* approach to ensembles as it produces an elegant, probabilistic and interpretable uncertainty estimates [39].

We use *Monte-Carlo-Dropout* (MC-Dropout) [19], which alleviates the need to train multiple models simultaneously and allows us to utilize Dropout to virtually generate multiple models (with different neurons dropped out of the original model) as Monte-Carlo samples and perform inferences on the sampled models, on each of the sequences in the hypotheses set by *teacher-forcing* [58], terming it as *stochastic inference*. Our motivation to utilize *teacher-forcing* in the pseudo-labeling phase is to enforce prediction consistency across all

the sampled models in the ensemble such that we can estimate the predictive distribution of every hypothesis obtained with *deterministic-inference*. Finally, the predictive posterior for each hypothesis is obtained by taking the expectation over all the sampled models. Furthermore, the obtained predictive posterior is used to compute an information-theoretic estimate of the uncertainty, which estimates the *total uncertainty* [21], considering both the character-level and word-level predictive posteriors and serves as a robust selection criterion for the pseudo-labels. Figure 1 shows an intuitive idea behind Beam-Search inference to generate multiple hypotheses for a normalized uncertainty estimate. Finally, we test our method on several handwriting and scene-text datasets, comparing its performance to state-of-the-art text-recognition methods in semi-supervised setting. Moreover, we demonstrate the robustness of our uncertainty estimation using *prediction rejection curves* [38, 40] based on the Word Error Rate (WER).

In sum the key-points are: (a) We propose an uncertainty-aware pseudo-label-based semi-supervised learning framework, that utilizes Beam-Search inference for pseudo-label assignment, and a character and sequence aware uncertainty estimate for sample selection. (b) We utilize *teacher-forcing* [58], primarily employed to train sequence-models, in the pseudo-labeling phase to enforce prediction consistency across all the sampled models in the ensemble, to estimate the predictive distribution. (c) Finally, the methods are evaluated on several challenging handwriting and scene-text datasets in the SSL setting.

2. Related Work

Text Recognition: Attention-based sequential decoders [5] have become the cutting-edge framework for text recognition in scene-text [33, 54, 60, 63] and handwriting [10, 37, 57, 64]. Besides, various incremental propositions [9, 11, 12] have been made, such as introducing or improving the rectification module [34, 50, 60, 63], designing a multi-directional convolutional feature extractor [17], enhancing the attention mechanism [16, 32], and stacking multiple Bi-LSTM layers for better context modeling [16, 49, 50]. With all the diversity present in text recognition models, the seminal work by Baek *et al.* [3] provides a unifying framework for text recognition that furnishes a modular viewpoint for existing methods by suggesting that the recognition system be divided into four distinct stages of operations, namely: (a) spatial-transformation (Trans.), (b) feature extraction (Feat.), (c) sequence modeling (Seq.), and (d) prediction (Pred.). Furthermore, the framework provides not only the existing methods but also their possible variants [3], and demonstrates that most of the state-of-the-art [14, 31, 34, 49, 50, 55] methods fall under this framework.

Semi-Supervised Learning: When labels are scarce or expensive to obtain, semi-supervised learning provides a powerful framework to exploit unlabeled data. SSL algorithms based on deep learning have demonstrated to be effective on

standard benchmark tasks. Numerous semi-supervised learning works have been published in recent years, including consistency regularization-based methods [30, 44, 43, 52], label propagation [7, 65], self-training [2, 15, 47], data augmentation [8, 59] and entropy regularization [25]. Even though semi-supervised learning is rapidly evolving, it is commonly employed and explored for unstructured predictions tasks *e.g.* classification or semantic segmentation. However, text recognition is a structured prediction task and therefore, off-the-shelf semi-supervised methods are unlikely to be directly applicable for the current use case.

Semi-Supervised Text-Recognition: Notwithstanding the obvious benefits, most text recognition systems do not currently use unlabeled text images. Handwriting recognition in particular, is usually based on fully-supervised training [51, 61], whereas scene-text models are primarily trained on synthetic data [22, 23]. However, Kang *et al.* [26] and Zhang *et al.* [64] recently proposed domain adaptation techniques for using an unlabeled dataset alongside labeled data. Nevertheless, they introduce specific modules to impose domain alignment. Fogel *et al.* [18] proposed a completely unsupervised scheme for handwritten text images, in which a discriminator forces predictions to align with a distribution of a given text corpus. Nonetheless, these methods necessitate restricting the recognizer to use only local predictions. Moreover, Gao *et al.* [20] concoct a reward-based methodology in a reinforcement-learning framework to perform SSL for scene-text recognition, however, they introduce an additional embedding network to compute the embedding distance to formulate the reward function, thus introducing a computational overhead at the time of training. Additionally, Aberdam *et al.* [1] propose a *sequence-to-sequence* contrastive self-supervised learning framework for text recognition, which implicitly shows SSL capabilities due to the generalization power of self-supervised learners [15], the method learns generalized image features with self-supervised pre-training, which is leveraged to initialize the weights at the fine-tuning phase. Moreover, we hypothesize that the model’s performance can further be improved at the fine-tuning phase by utilizing the unlabeled data.

3. Methodology

We start with a preliminary introduction describing the formulation of a generic text recognition model in 3.1. Then, we describe the semi-supervised framework as shown in Figure 2 in subsection 3.2 and describe the proposed sequential uncertainty estimation methodology in 3.3. Lastly in subsection 3.4 we verify the robustness of the introduced uncertainty estimation and show the estimated uncertainty’s correlation with word error rate (WER).

3.1. Preliminary

The text recognition model $f_{\theta}(\cdot)$ attempts to predict the machine readable character sequence Y given an image X . Among the numerous text-recognition models with different

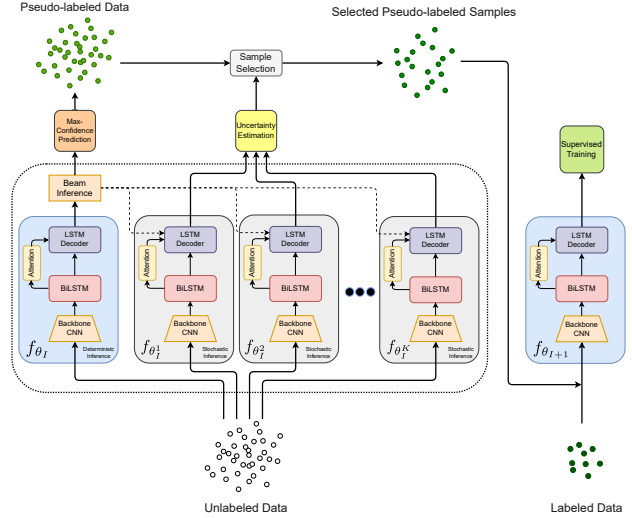


Figure 2: Schematic diagram representing pseudo label generation and uncertainty-based label selection strategy. The figure depicts a snap-shot of the pseudo-labeling process at the end of the I -th training iteration. The unlabeled data points in \mathcal{D}_u are assigned pseudo-labels by the *deterministic inference*. Furthermore, uncertainty associated with the each unlabelled sample is computed by *stochastic-inference* with K ensembles; and then samples with low uncertainty values, based on a threshold parameter, are selected to be included in the training set \mathcal{D}_{train} along with \mathcal{D}_l for the $(I + 1)$ -th training iteration. The dotted arrows represent the *teach-forcing* [58] applied at the time of uncertainty estimation. (The spatial-transformation module is omitted for brevity.)

backbones and sequential decoding schemes, we adopt a generic recognition model with an attention based sequential decoding scheme [3]. Most state-of-the-art text-recognition choices dealing with irregular or handwritten texts are derivatives of our chosen model. Our text recognition model consists of primarily four components: (a) a Spatial Transformer Network [24] to ease downstream stages by image normalization, (b) a backbone convolutional feature extractor, (c) a contextual sequence modelling stage, and (d) a sequence model that autoregressively predicts the characters at each time step. Let the extracted convolutional feature map be $V = \{v_i | v_i \in \mathbb{R}^C\}_{i=1}^W$, where W is the number of columns in the feature map and C be the number of channels. The feature map V passes through a sequence model, denoted as $\text{Seq}(\cdot)$ to generate a contextual sequence H , *i.e.*, $H = \text{Seq}(V) = \{h_i | h_i \in \mathbb{R}^D\}_{i=1}^W$, where D is the hidden-state dimension. Furthermore, at the t -th step, the decoder predicts a probability mass across the character space, $p(y_t) = \text{softmax}(W_0 s_t + b_0)$, where W_0 and b_0 are trainable parameters, and s_t is the decoder LSTM hidden state at step t defined as $s_t = \text{LSTM}(y_{t-1}, c_t, s_{t-1})$. Here, c_t is a context vector computed as a weighted sum of the elements in H from the previous step, defined as

$c_t = \sum_{i=1}^I \alpha_{t,i} h_i$, such that $\alpha_{t,i} = \frac{\exp e_{t,i}}{\sum_k \exp e_{t,k}}$, with $e_{t,i} = w_a^\top \tanh(W_b s_{t-1} + W_c h_i + b)$, where w_a, W_b, W_c , and b are trainable parameters. The complete recognition model is trained end-to-end using cross-entropy loss $\mathcal{CE}(\cdot)$ summed over the ground-truth sequence $Y = \{y_t\}_{t=1}^S$, where y_i is the one-hot encoded vector in \mathbb{R}^E , where E is the character vocabulary size and S denotes the sequence length.

$$\mathcal{L}_{CE} = \sum_{t=1}^S \mathcal{CE}(y_t, p(y_t)) = \sum_{t=1}^S \sum_{i=1}^E y_{t,i} \log(p(y_{t,i})). \quad (1)$$

Moreover, due to the autoregressive nature of the LSTM decoder, we can factorize the distribution over Y as product of conditional distributions, conditioned on previously predicted outputs $Y_{<t} = \{y_i\}_{i=1}^{t-1}$, the input X and all the trainable parameters denoted with θ as,

$$P(Y|X, \theta) = \prod_{t=1}^S P(y_t|Y_{<t}, X, \theta). \quad (2)$$

3.2. Pseudo-label Assignment and Selection for Semi-Supervised Text Recognition

Here we present an overview of our uncertainty-based pseudo-label semi-supervised framework, as shown in Figure 2. Let us denote the labeled dataset containing N_l samples as $\mathcal{D}_l = \{X_i^l, Y_i^l\}_{i=1}^{N_l}$, where X_i^l and Y_i^l denote the i -th training image and its corresponding character-sequence label, and let N_u denote the number of images that are presented in the unlabelled set denoted as $\mathcal{D}_u = \{X_i^u\}_{i=1}^{N_u}$, not associated with any ground-truth. In the proposed pseudo-labeling method, the recognition network undergoes multiple rounds of training. In the initial round, the model is first trained in a supervised manner with the labeled dataset, i.e. $\mathcal{D}_{train} = \mathcal{D}_l$. Furthermore, the trained model is used to estimate labels for unlabeled data as $\tilde{\mathcal{D}}_u = \{X_i^u, \tilde{Y}_i^u\}_{i=1}^{N_u}$, where \tilde{Y}_i^u denotes the pseudo-label for the i -th sample from the unlabeled pool. Concretely, given an unlabeled image X_i^u , Beam-Search inference is performed on the trained recognition network parameterized by θ to obtain a set of top- B high confidence predictions denoted by $\mathcal{B}_i = \{Y_i^{(b)}\}_{b=1}^B$ and is assigned the pseudo-label $\tilde{Y}_i^u = \arg \max_{Y_i^{(b)}} \{P(Y_i^{(b)}|X_i^u, \theta)\}_{b=1}^B$. Then we compute the uncertainty, which in its functional form can be denoted as $\mathcal{U}(X_i^u, \mathcal{B}_i)$, associated with the predictions from the input X_i^u . After this, the pseudo-label selection process is performed. Let $\mathbf{q} = \{q_i\}_{i=1}^{N_u}$, $q_i \in \{0, 1\}$, denote a binary vector, where a value of 1 at the i -th location depicts the selection of that particular unlabeled sample along with its pseudo label for subsequent training.

$$q_i = \begin{cases} 1, & \text{if } \mathcal{U}(X_i^u, \mathcal{B}_i) \leq \tau, \\ 0, & \text{otherwise,} \end{cases} \quad (3)$$

where τ is the selection threshold hyperparameter. Unlabeled samples with low uncertainty are selected and included in the

training dataset $\mathcal{D}_{train} = \mathcal{D}_l \cup \{\{X_i^u, \tilde{Y}_i^u\}_{i=1}^{N_u} | \text{if } q_i = 1\}$ for re-training the recognition model.

3.3. Uncertainty Estimation for Sequential Prediction

Now we elaborate on the adopted uncertainty estimation methodology for pseudo-label selection. We take an ensemble-based approach for uncertainty estimation, also referred to as MC-Dropout [19]. Consider an ensemble of models with the predicted probability distribution $\{P(Y|X, \theta^k)\}_{k=1}^K$, with each ensemble model parameter denoted by θ^k , where K denotes the number of ensembles, sampled from a posterior $\pi_p(\theta)$. Here, $\pi_p(\theta)$ is the set of all possible parameters with dropout applied to it with a given probability p . Moreover, the predictive posterior is defined as the expectation over the ensemble. In our case, to approximate the predictive posterior, we generate an ensemble of models by applying dropout, and virtually generate K versions from the same model, and define the predictive distribution as,

$$P(Y|X) := \int P(Y|X, \theta) \pi(\theta) d\theta \simeq \frac{1}{K} \sum_{k=1}^K P(Y|X, \theta^k), \quad \theta^k \sim \pi_p(\theta). \quad (4)$$

Furthermore, the *total uncertainty* [21] associated with the prediction of Y given an input X , is given by the entropy of the predictive distribution,

$$\begin{aligned} \mathcal{H}[P(Y|X)] &:= \mathbb{E}_{P(Y|X)} [-\ln P(Y|X)] \\ &= - \sum_{Y \in \mathcal{Y}} P(Y|X) \ln P(Y|X), \end{aligned} \quad (5)$$

where \mathcal{Y} denotes set of all possible hypotheses. However, the expression in (5) is not tractable for sequence models due to the ungovernable combinatorial vastness of the hypothesis space. Therefore, to devise a robust estimate of uncertainty for text recognition, we define a character and sequence aware uncertainty estimate that is tractable. For each unlabeled example, we use the hypotheses present in the set \mathcal{B}_i which contains the top- B high-probability predicted sequences obtained with Beam-Search inference at the time of decoding. Note that we had already obtained high-confidence hypotheses using *deterministic inference* at the time of pseudo label assignment, as described previously in 3.2. Furthermore, we perform inference on the different instances of the same model obtained with dropout by *teacher-forcing* to calculate the uncertainty measure as described below,

$$\begin{aligned} \mathcal{U}(X_i, \mathcal{B}_i) &= - \sum_{b=1}^B \sum_{t=1}^{L_b} \frac{\omega_b}{L_b} \mathcal{H}[P(y_t|Y_{<t}^{(b)}, X_i)], \\ \omega_b &= \frac{\exp(\frac{1}{T} \ln P(Y^{(b)}|X_i))}{\sum_{b=1}^B \exp(\frac{1}{T} \ln P(Y^{(b)}|X_i))}, \end{aligned} \quad (6)$$

where L_b denotes the length of the hypothesis and T denotes the temperature hyperparameter. Additionally, for the sequence-level measures of uncertainty, they need to be adjusted such that uncertainty associated with each hypothesis $Y^{(b)}$ within the hypotheses set \mathcal{B}_i is weighted in proportion to its probability (confidence) [48], denoted by ω_b . The defined uncertainty estimate takes into consideration the uncertainty at every character prediction and all the word sequences present in the hypotheses set \mathcal{B}_i . Thus, we approximate the entire hypotheses space with a few highly probable predictions. Moreover, our motivation to use *teacher-forcing* is to impose prediction consistency across all the models in the ensemble. Given an input image to the ensemble, the hypotheses set generated by each model is not guaranteed to be the same, therefore, we *teacher-force* the hypotheses in \mathcal{B}_i to infer from the models sampled for the ensemble.

3.4. Robust Uncertainty Estimation?

Ideally, we would like to identify all of the inputs that the model mis-predicts based on its uncertainty measure, and then we may decide to either include or choose not to use the predictions as pseudo-labels. Malinin *et al.* [38, 39, 40] proposed using the *prediction rejection curves* and the *prediction rejection area ratios* to qualitatively and quantitatively assess an uncertainty measure’s correlation to misclassifications, respectively. Instead of considering misclassifications, we examine our uncertainty measure’s correlation to Word Error Rate (WER) = $1 - \frac{\text{word prediction accuracy}}{100}$ (Figure 3). Given that, if the error metric is a linear function of individual errors, the expected *random* rejection curve should be a straight line from the base error rate on the left to the lower right corner if the uncertainty estimates are uninformative. However, if the estimates of uncertainty are *perfect* and always larger for an incorrect prediction than for a correct prediction, then the *oracle* rejection curve would be produced. When at a percentage of rejected examples equal to the number of false predictions, the *oracle* curve will fall linearly to 0 error. A rejection curve produced by imperfect but informative uncertainty estimates will rest between the *random* and *oracle* curves. The area between the selected measure’s rejection curves and *random* curve AR_r and the area between the *oracle* and *random* curves AR_{orc} can be used to assess the quality of the measure’s rejection curve. *prediction rejection area ratio* (PRR) is defined as $PRR = \frac{AR_r}{AR_{orc}}$.

Figure 3 depicts the rejection curves of the baseline text recognition model trained on different portions of labeled data on the handwriting and the scene-text datasets. We show the rejection curve comparison on two measures: (a) Confidence of the most probable hypothesis (in Blue) and (b) Our proposed sequence uncertainty measure (in Orange). For the confidence curves, the examples are rejected based on their confidence values, low-to-high, and for the uncertainty curves based on their uncertainty values, high-to-low. The rejection curves substantiate the superiority of our uncertainty-

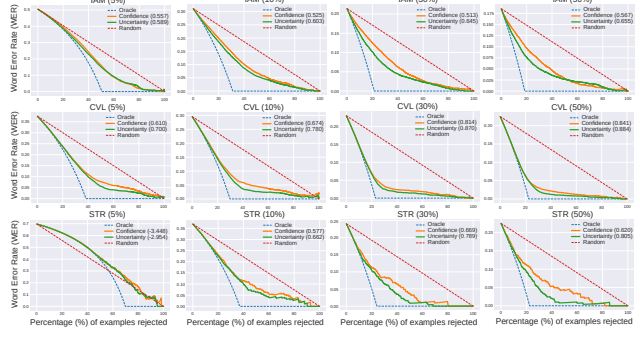


Figure 3: Prediction Rejection Curves. Values in parenthesis in the legend field represent the *PRR* of the corresponding curve.

based pseudo-label selection compared to confidence-based selection. A rejection area ratio of 1 represents optimal rejection, while 0 represents *random* rejection. In addition, the metric is independent of the model’s performance, allowing it to be used to compare models with different base error rates. The higher *PRR* values for uncertainty-based rejection suggest a stronger correlation to WER. Primarily, in Figure 3, we observe that Uncertainty-*PRR* in all the cases is greater than Confidence-*PRR*, validating the uncertainty-measure’s stronger correlation, conveying that the estimated uncertainty is a better estimation of the WER compared to confidence of the prediction. Especially, the gap in the *PRR* values is more significant for models trained with a lower fraction (5% and 10%) of labeled data. Additionally in the supplemental material we provide the prediction rejection curves w.r.t the *character error rate*.

4. Experiments

4.1. Experimental Settings



Figure 4: Examples of (a) handwritten-text images (IAM and CVL) and (b) scene-text images (IIIT5K, SVT, IC13, IC15, SVTP and CUTE)

Datasets: We compare the performances of the methods on a variety of publicly available datasets for handwriting and scene-text recognition. We consider the IAM [41] and CVL [29] word-level datasets for handwriting, and the datasets: IIIT5k [42], SVT [56], IC13 [28], IC15 [27], SVTP [45], and CUTE [46] for scene-texts as shown in Figure 4. To train the handwriting recognition models, we use the training data split of the IAM dataset and the CVL dataset. For scene-text recognition, unlike the previous works [1, 3] that use a large amount of synthetic training data [22, 23] to train scene-text models, we use a combination of multiple real scene-text

Method	IAM [41]				CVL [29]			
	Fraction of labeled data							
	5%		10%		5%		10%	
	Acc \uparrow	CER \downarrow	Acc \uparrow	CER \downarrow	Acc \uparrow	CER \downarrow	Acc \uparrow	CER \downarrow
Supervised Baseline	49.63	28.47	68.00	11.63	62.28	25.20	70.31	23.62
Pseudo-Labeling Baseline	54.57	24.64	71.06	13.23	64.38	29.74	73.18	23.42
SeqCLR (All-to-instance) [1]	64.58	11.63	70.54	9.72	63.10	29.40	70.51	24.15
SeqCLR (Frame-to-instance) [1]	64.81	14.41	71.64	9.55	63.08	27.30	70.63	23.89
SeqCLR (Window-to-instance) [1]	64.16	14.06	71.43	8.68	63.14	27.56	70.58	22.84
Ours	69.37	11.98	74.53	8.85	74.90	23.10	75.37	21.79
Ours w/ SeqCLR (All-to-instance)	73.60	9.03	77.11	7.47	72.41	19.95	75.67	21.26
Ours w/ SeqCLR (Frame-to-instance)	73.87	8.33	76.95	7.47	72.36	22.84	75.75	22.05
Ours w/ SeqCLR (Window-to-instance)	73.37	8.85	76.90	7.29	72.23	22.31	75.76	20.21

Table 1: Semi-Supervised Handwriting Recognition Results: Word level accuracy (Acc) (%) and Character Error Rate (CER) comparisons of models trained with 5% and 10% of the total available labeled data. Best results highlighted in bold.

recognition datasets [4]. More details can be found in the supplemental material.

Training Details: We use the best performing model configuration introduced by Baek *et al.* [3] commonly dubbed as TRBA (TPS-Residual-BiLSTM-Attention) with dropout layers between the convolution layers. We pre-resize all images to 32×100 for training and testing and use 95 symbol classes: 52 case-sensitive letters, 10 digits, and 33 special characters. For the MC-Dropout [19] ensembling we use a dropout probability of 0.1 and $K = 5$ ensembles. Furthermore, we use a beam-width of 5 with the temperature value T set to 0.01 to weigh the hypotheses according to their confidence. Moreover, the sample selection threshold τ for uncertainty is set to a low value of 0.01 for all the experiments in handwriting and scene-text recognition. For all the experiments we perform 5 iterations of pseudo-label based self-training and report the results of the best performing iteration after the first pseudo-labeling step. Furthermore we would like to emphasize that all of the pseudo-labeling and uncertainty estimation is only carried out after the supervised model training, and therefore it does not introduce any computational overhead at the time of training, contrary to the work of Gao *et al.* [20], which simultaneously learn an embedding network with a similar memory footprint as the recognition network, for reward-based gradient estimation. More about the training details can be found in the supplemental material.

Evaluation Metrics: For scene-text datasets we report the *Word Prediction Accuracy* following Baek *et al.* [3] and for handwriting datasets, additionally, we report *Character Error Rate (CER)* which is based on Levenstein distance [13, 51].

4.2. Baseline and State-of-the-art Comparisons

In Table 1 we compare the performance of the methods on handwriting recognition task in the semi-supervised setting, *i.e.*, being trained only with a portion of the labeled training dataset (5% and 10%) on Word Prediction Accuracy

(Acc) and the Character Error Rate (CER) on IAM [41] and CVL [29] word-level datasets. In addition to the supervised baseline we include the pseudo-labeling baseline, in which, we do not employ the threshold based unlabeled sample selection but select all the unlabeled samples along with their pseudo label for model retraining. We make comparisons with the 3 state-of-the-art variants of SeqCLR [1] with the BiLSTM projection head, which primarily differs in the instance mapping strategy used at the time of pre-training. All SeqCLR variants are first pre-trained with all the training examples, *i.e.*, $\mathcal{D}_{train} = \mathcal{D}_l \cup \mathcal{D}_u$, and then the pre-trained model is used as an initialization, and subsequently the entire model is fine-tuned with only a portion of labeled training examples in \mathcal{D}_{train} *i.e.* \mathcal{D}_l , in a fully-supervised manner. For our method, we use the same portions of the labeled split as used in the baseline and the prior arts and train from scratch starting with $\mathcal{D}_{train} = \mathcal{D}_l$. Post the supervised training with \mathcal{D}_{train} , pseudo-labeling is performed to update \mathcal{D}_{train} and the model is re-trained from scratch with the updated \mathcal{D}_{train} . We observe that our method consistently outperforms the baseline and all the variants of SeqCLR in the semi-supervised setting, especially when using 5% of the labeled data, with an average gain of $\approx 4.5\%$ and $\approx 2.9\%$ in accuracy, and around $\approx 2.5\%$ and $\approx 4.5\%$ reduction in CER for the IAM and CVL datasets, respectively. Note that even our vanilla PL-SSL method (without the self-supervised pre-training) performs significantly better than the prior arts.

Similarly, in Table 2, we compare our method’s performance on scene-text recognition task on word prediction accuracy following the same training principles used to train the handwriting recognition models. The last column in Table 2 represents the combined total accuracy *i.e.* all the scene-text datasets combined together for evaluation. The results corroborate the superiority of our methods in the semi-supervised scene-text recognition scenario, with an average gain in the total accuracy of $\approx 6.7\%$ and $\approx 1.2\%$ for model trained with 5% and 10% of the total labeled data on the vanilla PL-SSL strategy, respectively.

Moreover, in Figure 5 we visualize the word prediction accuracy on the handwriting datasets (IAM and CVL) and

Method	IIIT5k [42]		SVT [56]		IC13 [28]		IC15 [27]		SVTP [45]		CUTE [46]		Total	
	Fraction of labeled data													
	5%	10%	5%	10%	5%	10%	5%	10%	5%	10%	5%	10%	5%	10%
Supervised Baseline	41.63	75.60	33.39	69.55	46.11	78.52	26.43	55.51	24.96	54.88	23.26	48.96	35.32	67.30
Pseudo-Labeling Baseline	51.30	76.57	42.50	72.18	54.88	80.89	35.35	57.40	31.16	57.52	30.21	50.00	44.25	69.02
Gao <i>et al.</i> [20]	72.00	74.80	76.20	78.10	81.60	84.00	52.90	54.70	-	-	-	-	-	-
SeqCLR (All-to-instance) [1]	71.03	80.10	65.60	74.50	75.47	82.56	51.37	60.71	50.23	63.88	42.71	54.86	63.06	72.39
SeqCLR (Frame-to-instance) [1]	70.37	78.90	66.77	75.58	74.79	83.25	52.67	62.45	53.33	64.19	43.40	57.29	63.41	72.69
SeqCLR (Window-to-instance) [1]	71.60	80.37	65.38	74.19	76.65	82.37	51.47	61.34	52.87	64.65	41.32	58.33	63.58	72.81
Ours	76.70	81.23	72.64	76.66	80.59	84.63	60.23	62.98	61.40	64.81	53.82	60.07	70.27	73.98
Ours w/ SeqCLR (All-to-instance)	78.73	83.47	75.58	79.60	83.65	87.19	64.37	68.27	66.36	68.22	55.90	64.24	72.99	77.13
Ours w/ SeqCLR (Frame-to-instance)	80.50	83.37	76.20	79.44	84.34	86.90	65.34	67.45	67.75	70.08	58.33	65.97	74.56	77.36
Ours w/ SeqCLR (Window-to-instance)	79.87	83.23	74.34	80.53	85.22	86.70	63.70	67.55	66.36	69.30	57.99	64.93	73.74	77.19

Table 2: Semi-Supervised Scene-Text Recognition Results: Word level accuracy (%) comparisons of models trained with 5% and 10% of the total available labeled data. Best results highlighted in bold.

the scene-text recognition (STR) datasets (all datasets combined) at different fractions of the labels used on the supervised baseline, SeqCLR, and our proposed method with and without SeqCLR pre-training. Note that the SeqCLR trends are shown for its best performing variant.




Figure 5: Word prediction accuracy trends in the semi-supervised setting with 5%, 10%, 30%, and 50% of the total available labeled data used as \mathcal{D}_l , for the IAM, CVL, and the combined scene-text recognition (STR) datasets.

4.3. Ablation Studies

On iteration-wise performance evolution and the τ hyperparameter: Since the introduced methodology is iterative, and consists of multiple self-training steps, in Figure 6 we visualize the vanilla PL-SSL method’s performance on word prediction accuracy and CER at the end of each supervised training iteration, with different portions of labeled training dataset used to train the seed model, for the IAM, CVL, and

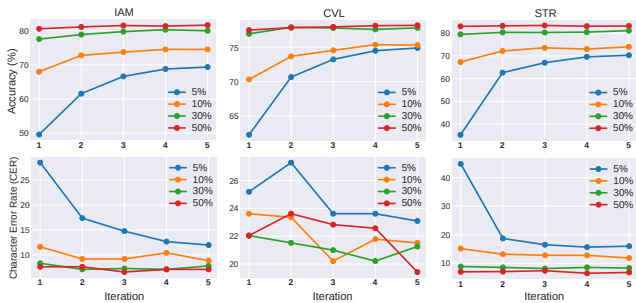


Figure 6: Iteration-wise metric evolution of the pseudo-labeling based semi-supervised learning methodology with randomly initialized weights.

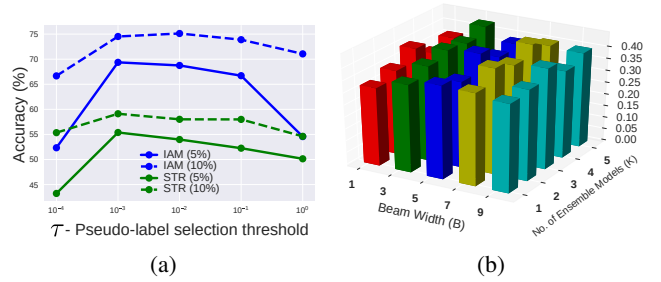


Figure 7: (a) Impact of τ at different values, on IAM and STR validation accuracy with 5% (solid lines) and 10% (dashed lines) of labeled data. (b) *Prediction rejection ratio* (PRR) variation of the employed uncertainty estimation across different values of beam-width (B) and the no. of ensembles (K), on the model trained on IAM dataset, based on the word prediction error metric. Note that values corresponding to the no. of ensembles $K = 1$ depicts *deterministic inference*.

combined STR datasets. We observe improvement in the performance in the initial iteration and then a saturation in the performance, this can be attributed to scenario when the set threshold hyperparameter τ does not allow new samples to be included for the further rounds, thus suggesting a constant τ parameter achieves a state of equilibrium and is not able to include new samples, which motivates future research directions to learn an iteration-wise adaptive threshold parameter or a curriculum of threshold which could break the equilibrium by having variable τ values, for the inclusion of additional samples at later iterations.

As established above, the threshold hyperparameter τ plays a critical role in the effectiveness of the proposed method, to understand its effect on the performance on PL-SSL, we run experiments on the IAM, CVL and STR datasets with different τ values in $\{10^{-i}\}_{i=0}^4$ and note the best performance among the self-training iterations, on their respective validation set. In Figure 7a, we visualize the validation performance at different τ values and observe that when τ is significantly low, the selected examples are less informative,

Dataset	Dropout Probability (p)				
	0.0	0.1	0.2	0.3	0.4
IAM	0.1551	0.1475	0.1409	0.1335	0.1237
CVL	0.1914	0.1819	0.1718	0.1601	0.1513
STR	0.4381	0.4266	0.4169	0.3939	0.3759

Table 3: Expected Calibration Error (ECE) across different dropout probability values (p) with $K = 5$ ensembles, as used in all the experiments, trained with all the samples. Note that values corresponding to $p = 0.0$ denotes *deterministic inference*.

since they are too low in number for improving the model’s performance, especially if large proportion of labeled data is used. If too high, the selected samples become too noisy for training. For all our experiments, we use a fixed τ value of 10^{-2} .

On optimal Beam-width (B) and No. of ensemble models (K) combination: To compute the uncertainty measure defined in (6), we require a finite set of highly probable predictions to approximate the hypotheses space and the predictive distribution of each hypothesis is approximated as described in expression (4) by taking the expectation of the virtually generated ensemble. To identify the optimum beam-width (B) and the no. of ensemble models (K) combination, we visualize the variation of the PRR , described in subsection 3.4, with respect to a finite set of values of B and K as shown in Figure 7b. From the plot we observe no specific trend on the combination of B and K which can generalize the hyperparameter selection, therefore, in all our experiments, we used a fixed value for B and K equal to 5, in practice, one can use the validation set to find the optimal combination by observing the PRR .

Effect of dropout on the prediction distribution: To corroborate our motivation, we conduct experiments to observe the effect of dropout when computing the predictive distribution, we evaluate the quality of the predictive distribution by observing the Expected Calibration Error (ECE) at distinct dropout probability values. In Table 3, we present the ECE values for multiple datasets, at different dropout probability (p) values. As we increase the dropout value, the ECE reduces, hence, demonstrating the utility of dropout in computing the predictive distribution which is eventually used to compute the *total-uncertainty*. Finally, in Figure 8, we show the relationship between the prediction uncertainty and the calibration error (ECE). When pseudo-labels with more certain predictions are chosen, the calibration error for that subset of pseudo-labels is considerably reduced. As a result, the predictions are more likely to result in an accurate pseudo-label for that subset of labels.




Figure 8: The link between prediction uncertainty and expected calibration error (ECE) is depicted here. With different dropout probability p , the ECE of the selected subset diminishes as the uncertainty of the picked pseudo-labels decreases. Note that, compared to *deterministic-inference*, the MC-Dropout-based model ensembling provides a lower calibration error for a fixed number of elements in the uncertainty-based selected subset.

5. Discussion

Limitations: Despite the superior SSL performance, the method discussed is not devoid of certain limitations. Thus, we identify two primary limitations as follows. Firstly, the current framework assumes the same character space for the labeled and the unlabeled examples. Therefore, it may be hard to employ such an SSL framework directly in scenarios when the labeled and the unlabeled examples have character space discrepancy, *e.g.*, cross-lingual text recognition, when labeled examples come from one language and unlabeled samples from the other. The second assumption is based on the presence of dropout layers in the backbone network. In Table 3, we observe improvements in the Expected Calibration Error (ECE) due to the presence of dropout. However, the observation made in [47] and verified as shown in Figure 8 suggests that selecting pseudo-labeled samples with low uncertainty also helps reduce the calibration error. Additionally, directly estimating uncertainty by *deterministic inference* can also be considered, though it may require having a stricter (lower) sub-set selection threshold τ to attenuate the calibration error, however the number of pseudo-labeled examples may be low.

Conclusion: This paper proposes a pseudo-label generation and an uncertainty-based data selection framework for semi-supervised image-based text recognition. The method uses Beam-Search inference to yield highly probable hypotheses for pseudo-label assignment of the unlabelled samples. Furthermore, it adopts an ensemble of models obtained by MC-Dropout to get a robust estimate of the uncertainty related to the input image that considers both the character-level and word-level predictive distributions. Moreover, we show superior SSL performance in handwriting and scene-text recognition tasks. In addition, we also utilize prediction rejection curves to demonstrate the correlation of the uncertainty estimate with the word prediction error rate.

Acknowledgement: Work partially supported by HP and National Science Foundation (1737744).

References

- [1] Aviad Aberdam, Ron Litman, Shahar Tsiper, Oron Anshel, Ron Slossberg, Shai Mazon, R Manmatha, and Pietro Perona. Sequence-to-sequence contrastive learning for text recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [2] Eric Arazo, Diego Ortego, Paul Albert, Noel E O'Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *2020 International Joint Conference on Neural Networks (IJCNN)*, 2020.
- [3] Jeonghun Baek, Geewook Kim, Junyeop Lee, Sungrae Park, Dongyoon Han, Sangdoo Yun, Seong Joon Oh, and Hwal-suk Lee. What is wrong with scene text recognition model comparisons? dataset and model analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [4] Jeonghun Baek, Yusuke Matsui, and Kiyoharu Aizawa. What if we only use real datasets for scene text recognition? toward scene text recognition with fewer labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [5] Dzmitry Bahdanau, Kyung Hyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations (ICLR)*, 2015.
- [6] Fan Bai, Zhanzhan Cheng, Yi Niu, Shiliang Pu, and Shuigeng Zhou. Edit probability for scene text recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [7] Yoshua Bengio, Olivier Delalleau, and Nicolas Le Roux. *Label Propagation and Quadratic Criterion*, pages 193–216. MIT Press, semi-supervised learning edition, January 2006.
- [8] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [9] Ayan Kumar Bhunia, Pinaki Nath Chowdhury, Aneeshan Sain, and Yi-Zhe Song. Towards the unseen: Iterative text recognition by distilling from errors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [10] Ayan Kumar Bhunia, Abhirup Das, Ankan Kumar Bhunia, Perla Sai Raj Kishore, and Partha Pratim Roy. Handwriting recognition in low-resource scripts using adversarial learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [11] Ayan Kumar Bhunia, Shuvojit Ghose, Amandeep Kumar, Pinaki Nath Chowdhury, Aneeshan Sain, and Yi-Zhe Song. Metaht: Towards writer-adaptive handwritten text recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [12] Ayan Kumar Bhunia, Aneeshan Sain, Amandeep Kumar, Shuvojit Ghose, Pinaki Nath Chowdhury, and Yi-Zhe Song. Joint visual semantic reasoning: Multi-stage decoder for text recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [13] Théodore Bluche. Joint line segmentation and transcription for end-to-end handwritten paragraph recognition. *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- [14] Fedor Borisjuk, Albert Gordo, and Viswanath Sivakumar. Rosetta: Large scale system for text detection and recognition in images. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018.
- [15] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big self-supervised models are strong semi-supervised learners. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [16] Zhanzhan Cheng, Fan Bai, Yunlu Xu, Gang Zheng, Shiliang Pu, and Shuigeng Zhou. Focusing attention: Towards accurate text recognition in natural images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017.
- [17] Zhanzhan Cheng, Yangliu Xu, Fan Bai, Yi Niu, Shiliang Pu, and Shuigeng Zhou. Aon: Towards arbitrarily-oriented text recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [18] Sharon Fogel, Hadar Averbuch-Elor, Sarel Cohen, Shai Mazon, and Roei Litman. Scrabblegan: Semi-supervised varying length handwritten text generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [19] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of The 33rd International Conference on Machine Learning (ICML)*, 2016.
- [20] Yunze Gao, Yingying Chen, Jinqiao Wang, and Hanqing Lu. Semi-supervised scene text recognition. *IEEE Transactions on Image Processing*, 30:3005–3016, 2021.
- [21] Purvi Goel and Li Chen. On the robustness of monte carlo dropout trained with noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [22] Ankush Gupta, Andrea Vedaldi, and Andrew Zisserman. Synthetic data for text localisation in natural images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [23] M Jaderberg, K Simonyan, A Vedaldi, and A Zisserman. Synthetic data and artificial neural networks for natural scene text recognition. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2014.
- [24] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.
- [25] Tarun Kalluri, Girish Varma, Manmohan Chandraker, and CV Jawahar. Universal semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [26] Lei Kang, Marçal Rusinol, Alicia Fornés, Pau Riba, and Mauricio Villegas. Unsupervised writer adaptation for synthetic-to-real handwritten word recognition. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2020.
- [27] Dimosthenis Karatzas, Lluís Gomez-Bigorda, Anguelos Nicolaou, Suman Ghosh, Andrew Bagdanov, Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vijay Ramaseshan Chandrasekhar, Shijian Lu, Faisal Shafait, Seiichi Uchida, and Ernest Valveny. Icdar 2015 competition on robust reading. In

2015 13th International Conference on Document Analysis and Recognition (ICDAR), 2015.

- [28] Dimosthenis Karatzas, Faisal Shafait, Seiichi Uchida, Masakazu Iwamura, Lluís Gomez i Bigorda, Sergi Robles Mestre, Joan Mas, David Fernandez Mota, Jon Almazàn Almazàn, and Lluís Pere de las Heras. Icdar 2013 robust reading competition. In *2013 12th International Conference on Document Analysis and Recognition (ICDAR)*, 2013.
- [29] Florian Kleber, Stefan Fiel, Markus Diem, and Robert Sablatnig. Cvl-database: An off-line database for writer retrieval, writer identification and word spotting. In *2013 12th International Conference on Document Analysis and Recognition (ICDAR)*, 2013.
- [30] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *International Conference on Learning Representations (ICLR)*, 2017.
- [31] Chen-Yu Lee and Simon Osindero. Recursive recurrent nets with attention modeling for ocr in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [32] Hui Li, Peng Wang, Chunhua Shen, and Guyu Zhang. Show, attend and read: A simple and strong baseline for irregular text recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- [33] Ron Litman, Oron Anschel, Shahar Tsiper, Roei Litman, Shai Mazor, and R Manmatha. Scatter: selective context attentional scene text recognizer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [34] W. Liu, Chaofeng Chen, Kwan-Yee Kenneth Wong, Zhizhong Su, and Junyu Han. Star-net: A spatial attention residue network for scene text recognition. In *British Machine Vision Conference (BMVC)*, 2016.
- [35] Shangbang Long, Xin He, and Cong Yao. Scene text detection and recognition: The deep learning era. *International Journal of Computer Vision*, 129(1):161–184, 2021.
- [36] S.M. Lucas, A. Panaretos, L. Sosa, A. Tang, S. Wong, and R. Young. Icdar 2003 robust reading competitions. In *2003 7th International Conference on Document Analysis and Recognition (ICDAR)*, 2003.
- [37] Canjie Luo, Yuanzhi Zhu, Lianwen Jin, and Yongpan Wang. Learn to augment: Joint data augmentation and network optimization for text recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [38] Andrey Malinin. Uncertainty estimation in deep learning with application to spoken language assessment (doctoral thesis). In *University of Cambridge*, 2019.
- [39] Andrey Malinin and Mark Gales. Uncertainty estimation in autoregressive structured prediction. In *International Conference on Learning Representations (ICLR)*, 2021.
- [40] Andrey Malinin, Anton Ragni, Kate Knill, and Mark Gales. Incorporating uncertainty into deep learning for spoken language assessment. In *Association for Computational Linguistics (ACL)*, 2017.
- [41] Urs-Viktor Marti and Horst Bunke. The iam-database: an english sentence database for offline handwriting recognition. *International Journal on Document Analysis and Recognition*, 5:39–46, 2002.
- [42] Anand Mishra, Alahari Karteek, and C. V. Jawahar. Scene text recognition using higher order language priors. In *British Machine Vision Conference (BMVC)*, 2012.
- [43] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [44] Yassine Ouali, Céline Hudelot, and Myriam Tami. Semi-supervised semantic segmentation with cross-consistency training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [45] Trung Quy Phan, Palaiahnakote Shivakumara, Shangxuan Tian, and Chew Lim Tan. Recognizing text with perspective distortion in natural scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2013.
- [46] Anhar Risnumawan, Palaiahankote Shivakumara, Chee Seng Chan, and Chew Lim Tan. A robust arbitrary text detection system for natural scene images. *Expert Systems with Applications*, 41(18):8027–8048, 2014.
- [47] Mamshad Nayeem Rizve, Kevin Duarte, Yogesh S Rawat, and Mubarak Shah. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. In *International Conference on Learning Representations (ICLR)*, 2021.
- [48] Verónica Romero, Joan Andreu Sánchez, and Alejandro H. Toselli. Active learning in handwritten text recognition using the derivational entropy. In *2018 16th International Conference on Frontiers in Handwriting Recognition (ICFHR)*, 2018.
- [49] Baoguang Shi, Xiang Bai, and Cong Yao. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2016.
- [50] Baoguang Shi, Xinggang Wang, Pengyuan Lyu, Cong Yao, and Xiang Bai. Robust scene text recognition with automatic rectification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [51] Jorge Sueiras, Victoria Ruiz, Angel Sanchez, and Jose F. Velez. Offline continuous handwriting recognition using sequence to sequence neural networks. *Neurocomputing*, 2018.
- [52] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [53] Apoorv Vyas, Pranay Dighe, Sibong Tong, and Hervé Bourlard. Analyzing uncertainties in speech recognition using dropout. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- [54] Zhaoyi Wan, Jielei Zhang, Liang Zhang, Jiebo Luo, and Cong Yao. On vocabulary reliance in scene text recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [55] Jianfeng Wang and Xiaolin Hu. Convolutional neural networks with gated recurrent connections. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [56] Kai Wang, Boris Babenko, and Serge J. Belongie. End-to-end scene text recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2011.

- [57] Tianwei Wang, Yuanzhi Zhu, Lianwen Jin, Canjie Luo, Xiaoxue Chen, Yaqiang Wu, Qianying Wang, and Mingxiang Cai. Decoupled attention network for text recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
- [58] Ronald J. Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural Computation*, 1989.
- [59] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [60] Mingkun Yang, Yushuo Guan, Minghui Liao, Xin He, Kaigui Bian, Song Bai, Cong Yao, and Xiang Bai. Symmetry-constrained rectification network for scene text recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [61] Mohamed Yousef and Tom E Bishop. Origaminet: weakly-supervised, segmentation-free, one-step, full page text recognition by learning to unfold. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [62] Deli Yu, Xuan Li, Chengquan Zhang, Tao Liu, Junyu Han, Jingtuo Liu, and Errui Ding. Towards accurate scene text recognition with semantic reasoning networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [63] Fangneng Zhan, Shijian Lu, and Chuhui Xue. Verisimilar image synthesis for accurate detection and recognition of texts in scenes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [64] Yaping Zhang, Shuai Nie, Wenju Liu, Xing Xu, D. Zhang, and Heng Tao Shen. Sequence-to-sequence domain adaptation network for robust text image recognition. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [65] Xiaojin Zhu and Zoubin Ghahramani. Learning from labeled and unlabeled data with label propagation. In *CMU-CALD-02-107*, 2002.