

Confidence-Based Discriminative Training for Model Adaptation in Offline Arabic Handwriting Recognition

Philippe Dreuw, Georg Heigold, and Hermann Ney
 RWTH Aachen University
 Human Language Technology and Pattern Recognition
 Ahornstr 55, D-52056 Aachen, Germany
 <surname>@cs.rwth-aachen.de

Abstract

We present a novel confidence-based discriminative training for model adaptation approach for an HMM based Arabic handwriting recognition system to handle different handwriting styles and their variations.

Most current approaches are maximum-likelihood trained HMM systems and try to adapt their models to different writing styles using writer adaptive training, unsupervised clustering, or additional writer specific data.

Discriminative training based on the Maximum Mutual Information criterion is used to train writer independent handwriting models. For model adaptation during decoding, an unsupervised confidence-based discriminative training on a word and frame level within a two-pass decoding process is proposed. Additionally, the training criterion is extended to incorporate a margin term.

The proposed methods are evaluated on the IFN/ENIT Arabic handwriting database, where the proposed novel adaptation approach can decrease the word-error-rate by 33% relative.

1. Introduction

In this paper, we describe our discriminative training and multi-pass decoding system for offline Arabic handwriting, present our novel unsupervised confidence-based discriminative model adaptation approach, and present systematic results on the IFN/ENIT database [15].

Most state-of-the-art single-pass [11] and multi-pass [2, 4, 5] HMM based handwriting recognition systems are trained using the maximum-likelihood criterion.

Similar to the system presented in [13], we apply discriminative training using the Maximum Mutual Information (MMI) which is modified by a margin term. This margin term can be interpreted as an additional observation-

dependent prior weakening the true prior [9], and is identical with the SVM optimization problem of log-linear models [7].

The most common way for unsupervised adaptation is the use of the automatic transcription of a previous recognition pass without the application of confidence scores. Many publications have shown that the application of confidence scores for adaptation can improve recognition results. However, only small improvements are reported for confidence based CMLLR adaptation [1] or MLLR adaptation [6, 14, 16]. In this work, we present a novel unsupervised confidence-based discriminative model adaptation approach using a modified MMI training criterion.

2. System Overview

We are searching for an unknown word sequence $w_1^N := w_1, \dots, w_N$, for which the sequence of features $x_1^T := x_1, \dots, x_T$ best fits to the trained models. We maximize the posterior probability $p(w_1^N | x_1^T)$ over all possible word sequences w_1^N with unknown number of words N . This is modeled by Bayes' decision rule:

$$\hat{w}_1^N = \arg \max_{w_1^N} \{p^\gamma(w_1^N) p(x_1^T | w_1^N)\} \quad (1)$$

with γ being a scaling exponent of the language model.

In this work, we use a writing variant model refinement [4] of our visual model

$$p(x_1^T | w_1^N) \approx \max_{v_1^N | w_1^N} \{p_{\theta_{pm}}^\alpha(v_1^N | w_1^N) p_{\theta_{em,tp}}^\beta(x_1^T | v_1^N, w_1^N)\} \quad (2)$$

with v_1^N a sequence of unknown writing variants, α a scaling exponent of the writing variant probability depending on a parameter set θ_{pm} , and β a scaling exponent of the visual character model depending on a parameter set $\theta_{em,tp}$ for emission and transition model.

2.1. Feature Extraction

Without any preprocessing of the input images, we extract simple appearance-based image slice features X_t at every time step $t = 1, \dots, T$ which are augmented by their spatial derivatives in horizontal direction $\Delta = X_t - X_{t-1}$. In order to incorporate temporal and spatial context into the features, we concatenate 7 consecutive features in a sliding window, which are later reduced by a PCA transformation matrix to a feature vector x_t .

2.2. Discriminative Training

Our baseline hidden Markov model (HMM) based handwriting recognition system is Viterbi trained using the maximum-likelihood training criterion, model length estimation (MLE) for character dependent model lengths, and a lexicon with multiple writing variants as proposed in [3, 4].

In this work, we use a discriminative training approach based on the Modified Maximum Mutual Information criterion as presented in [7]. In the following, we give a brief summary.

Maximum Mutual Information (MMI). In automatic speech recognition (ASR), MMI commonly refers to the maximum likelihood (ML) for the class posteriors.

$$\mathcal{F}^{(\text{MMI})}(\theta) = -\frac{1}{N} \sum_{r=1}^R \log \left(\frac{p_\theta(x_1^{T_r} | w_1^{N_r}) p(w_1^{N_r})}{\sum_{v_1^{M_r}} p_\theta(x_1^{T_r} | v_1^{M_r}) p(v_1^{M_r})} \right). \quad (3)$$

This criterion has proven to perform reasonably as long as the error rate on the training data is not too low, i.e., generalization is not an issue.

Modified Maximum Mutual Information (M-MMI). We define a modified criterion

$$\mathcal{F}_\gamma^{(\text{M-MMI})}(\theta) = \mathcal{R}(\theta, \theta_0) - \frac{J}{R} \sum_{r=1}^R \frac{1}{\gamma} \log \left(\frac{[p_\theta(x_1^{T_r} | w_1^{N_r}) p(w_1^{N_r}) e^{(-\rho\delta(w_1^{N_r}, w_1^{N_r}))}]^\gamma}{\sum_{v_1^{M_r}} [p_\theta(x_1^{T_r} | v_1^{M_r}) p(v_1^{M_r}) e^{(-\rho\delta(w_1^{N_r}, v_1^{M_r}))}]^\gamma} \right). \quad (4)$$

The approximation level γ is an additional parameter to control the smoothness of the criterion. The regularization constant of \mathcal{R} is proportional to $\frac{1}{J}$. Here, I-smoothing is used for regularization [17]. The major difference to the standard MMI formulation is the additional margin term which is non-zero only for the correct w_1^N . This margin term can be interpreted as an additional observation dependent prior, weakening the true prior [9]. Moreover, this training criterion is identical with the SVM optimization problem for $\gamma \rightarrow \infty$ and log-linear models [7]. Keep in mind that GHMMs with globally pooled variances are




Figure 1. Comparison of hinge loss, MMI, and modified MMI

equivalent to a log-linear model with first order features only [8]. See Figure 1 for a comparison of the hinge loss function, MMI, and modified MMI.

Optimization. In [7] it is shown that the objective function $\mathcal{F}_\gamma^{(\text{MMI})}(\Lambda)$ converges pointwise to the SVM optimization problem using the hinge loss function for $\gamma \rightarrow \infty$, similar to [18]. In other words, $\mathcal{F}_\gamma^{(\text{M-MMI})}(\Lambda)$ is a smooth approximation to an SVM with hinge loss function which can be iteratively optimized with standard gradient-based optimization techniques like Rprop [7, 18]. In this work, the approximation level and the margin are chosen beforehand and then kept fixed during the complete optimization using the Rprop algorithm.

State-based confidences and modified MMI. Word confidences can be incorporated into the training criterion by simply weighing the segments with the respective confidence. This is, however, not possible for state-based confidences.

Rprop is a gradient-based optimization algorithm. The gradient of the training criterion under consideration can be represented in terms of the state posteriors $p_{rt}(s|x_1^{T_r})$. These posteriors are obtained by marginalization and normalization of the joint probabilities $p_\theta(x_1^{T_r}, s_1^T, w_1^{N_r})$ over all state sequences through state s at frame t . These quantities can be calculated efficiently by recursion, c.f. forward/backward probabilities. Then, the state-based confidences are incorporated by multiplying the posteriors with the respective confidence before the accumulation. In summary, each frame t contributes $\text{conf}(t)p_{rt}(s|x_1^{T_r})x_t$ to the accumulator of state s .

3. Decoding Architecture

The recognition is performed in two passes, as depicted in Figure 2. System 1 performs the initial and independent recognition pass using the discriminative trained models. The output is required for the text dependent model adaptation in the next step.

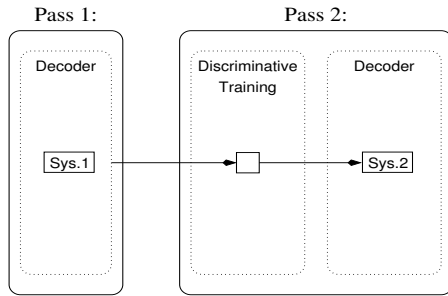


Figure 2. Illustration of the two-pass decoding process using confidence-based discriminative training for model adaptation.

The model adaptation in the second pass is performed by discriminatively training a System 2 on the text output of the first-pass recognition system. Additionally, the confidence-alignments generated during the first-pass decoding can be used on a sentence-, word-, or state-level to exclude the corresponding features from the discriminative training process for model adaptation.

Word Confidences. As we are dealing with isolated word recognition on the IFN/ENIT database, the sentence and word confidences are identical. The segments to be used in the second-pass system are first thresholded on a *word-level* by their word confidences: only complete word *segments* aligned with a high confidence by the first-pass system are used for model adaptation using discriminative training.

State Confidences. Instead of rejecting an entire utterance or word, the system can use state confidence scores to select state-dependent data. State confidence scores are obtained from computing arc posteriors from the lattice output from the decoder. The arc posterior is the fraction of the probability mass of the paths that contain the arc from the mass that is represented by all paths in the lattice. The posterior probabilities can be computed efficiently using the forward-backward algorithm as, for example, described in [10]. The word frames to be used in the second-pass system are first thresholded on a *state-level* by their state confidences: only word *frames* aligned with a high confidence by the first-pass system, are used for model adaptation using discriminative training (see subsection 2.2).

An example for a word-graph and the corresponding 1-best state alignment is given in Figure 3: during the decoding, the ten feature frames (the squares) can be aligned to different words (long arcs) and their states. In this examples, the word-confidence of the 1-best alignment is $c = 0.7$. The corresponding state-confidences are calculated by accumulating state-wise over all other word alignments, i.e. the state-confidence of the 1-best alignment’s fourth state would stay 0.7, all other state-confidences sum up to 1.0.




Figure 3. Example for a word-graph and the corresponding 1-best state alignment: word-confidence of the 1-best alignment is $c = 0.7$. The corresponding state-confidences are calculated by accumulating state-wise over all other word alignments




Figure 4. IFN/ENIT corpora splits used in 2005 and 2007.

4. Experimental Results

The experiments are conducted on the IFN/ENIT database [15]. The database is divided into four training folds with an additional fold for testing [12]. The current database version (v2.0p1e) contains a total of 32492 Arabic words handwritten by about 1000 writers, and has a vocabulary size of 937 Tunisian town names. Additionally, the submitted systems to the ICDAR 2007 competition [11] were trained on all datasets of the IFN/ENIT database and evaluated for known datasets. Here, we follow the same evaluation protocol as for the ICDAR 2005 and 2007 competition (see Figure 4).

4.1. First Pass Decoding

In this section we compare our maximum-likelihood trained baseline system to our discriminative trained systems using the MMI and modified margin-based MMI criterion. The discriminative training is initialized with the respective ML baseline model and iteratively optimized using the Rprop algorithm.

The number of Rprop iterations and the choice of the regularization constant $\mathcal{R}(\theta, \theta_0)$ have to be chosen carefully (c.f. optimization in subsection 2.2), and were empirically

Table 1. Comparison of maximum-likelihood trained baseline system (ML), and discriminative trained systems using MMI criterion and margin-based MMI (M-MMI) criterion after 30 Rprop iterations.

Train	Test	WER[%]		
		ML	MMI	M-MMI
abc	d	10.88	10.59	8.94
abd	c	11.50	10.58	2.66
acd	b	10.97	10.43	8.64
bcd	a	12.19	11.41	9.59
abcd	e	21.86	21.00	19.51
abcde	e	11.14	2.32	2.95

optimized in informal experiments to 30 Rprop iterations.

The results in Table 1 show that the discriminatively trained models clearly outperform the maximum likelihood trained models, especially the models trained with the additional margin term. The strong decrease in word-error-rate (WER) for experiment setup *abd-c* might be due to the training data being separable for the given configurations, whereas the strong improvement for experiment *abcde-e* was expected because of the test set *e* being part of the training data.

4.2. Second Pass Decoding and Unsupervised Model Adaptation

In this section we evaluate our discriminative training for unsupervised model adaptation during a second pass decoding step.

In a first experiment we used the complete first-pass output of the M-MMI system for an unsupervised adaptation. The results in Table 2 show that the M-MMI based unsupervised adaptation cannot improve the system accuracy. With every Rprop iteration, the system is even more biased by the relatively large amount of wrong transcriptions in the adaptation corpus.

Using the word-confidences of our first-pass alignment to reject complete word segments (i.e. feature sequences x_1^T) from the unsupervised adaptation corpus, the results in Table 2 show a slight improvement only in comparison to the M-MMI trained system. Figure 5 shows the resulting WER for different confidence threshold values and the corresponding number of rejected segments. For a confidence threshold of $c = 0.5$, more than 60% of the 6033 segments of set *e* are rejected from the unsupervised adaptation corpus, resulting in a relatively small amount of adaptation data.




Figure 5. Results for word-confidence based discriminative training using different confidence thresholds and their corresponding number of rejected segments.

Table 2. Results for confidence-based model adaptation on the evaluation experiment setup *abcd-e* using a margin-based MMI criterion and 30 Rprop iterations.

Training/Adaptation	WER[%]	CER[%]
ML	21.86	8.11
M-MMI	19.51	7.00
+ unsupervised adaptation	20.11	7.34
+ word-confidences	19.23	7.02
+ state-confidences	17.75	6.49
+ supervised adaptation	2.06	0.77

Using the state-confidences of our first-pass alignment to decrease the contribution of single frames (i.e. features x_t) during the iterative M-MMI optimization process (c.f. optimization in subsection 2.2), the number of features for model adaptation is reduced by approximately 5%: 375 446 frames of 396 416 frames extracted from the 6033 test segments are considered during the optimization, only 20 970 frames are rejected based on confidence thresholding (c.f. also Figure 3). Note that also the character-error-rate (CER) is decreased to 6.49%.

Interestingly, the supervised adaptation on test set *e*, where the correct transcriptions of set *e* are used for an adaptation of the model trained using set *abcd*, can again decrease the word-error-rate of the system down to 2.06%, which is even better than an M-MMI optimization on the full training set *abcde* (c.f. Table 1).

Table 3 shows the final results of our Arabic handwriting

Table 3. Results for confidence-based model adaptation on the IFN/ENIT database using model length estimation (MLE), a margin-based MMI criterion and 30 Rprop iterations.

Train	Test	WER[%]			
		1st pass		2nd pass	
		ML	+MLE	+M-MMI	M-MMI-conf
abc	d	10.88	7.83	6.12	5.95
abd	c	11.50	8.83	6.78	6.38
acd	b	10.97	7.81	6.08	5.84
bcd	a	12.19	8.70	7.02	6.79
abcd	e	21.86	16.82	15.35	14.55

recognition system with additional model length estimation (MLE) as described in [3, 4]. Again, the WER of the MLE based system can be decreased by our proposed modified MMI training during both decoding passes down to 14.55%, which is the currently best known WER in the literature.

5. Conclusions

We presented a novel confidence-based discriminative training using a margin-based Maximum Mutual Information training criterion for model adaptation in offline Arabic handwriting recognition. The advantages of the proposed methods using the HMM based multi-pass decoding system were shown on the IFN/ENIT corpus.

The proposed discriminative training could outperform the maximum-likelihood trained system on all cross folds.

The impact of different writing styles was dealt with a novel confidence-based discriminative training for model adaptation, where the usage of state-confidences during the iterative optimization process based on the modified MMI criterion could decrease the word-error-rate by 33% relative in comparison to a maximum-likelihood trained system.

Acknowledgements. We would like to thank Christian Gollan for his support. This work was partly realized as part of the Quaero Programme, funded by OSEO, French State agency for innovation.

References

- [1] T. Anastasakos and S. Balakrishnan. The use of confidence measures in unsupervised adaptation of speech recognizers. In *International Conference on Spoken Language Processing*, Sydney, Australia, 1998.
- [2] R. Bertolami and H. Bunke. Hmm-based ensemble methods for offline handwritten text line recognition. *Pattern Recognition*, 41:3452–3460, 2008.
- [3] P. Dreuw, S. Jonas, and H. Ney. White-space models for off-line arabic handwriting recognition. In *International Conference on Pattern Recognition*, Tampa, Florida, USA, Dec. 2008.
- [4] P. Dreuw, D. Rybach, C. Gollan, and H. Ney. Writer adaptive training and writing variant model refinement for off-line arabic handwriting recognition. In *ICDAR*, Barcelona, Spain, July 2009.
- [5] G. A. Fink and T. Plötz. Unsupervised estimation of writing style models for improved unconstrained off-line handwriting recognition. In *International Workshop on Frontiers in Handwriting Recognition*, La Baule, France, Oct. 2006.
- [6] C. Gollan and M. Bacchiani. Confidence scores for acoustic model adaptation. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages 4289–4292, Las Vegas, NV, USA, Apr. 2008.
- [7] G. Heigold, T. Deselaers, R. Schlüter, and H. Ney. Modified mmi/mpe: A direct evaluation of the margin in speech recognition. In *International Conference on Machine Learning*, pages 384–391, Helsinki, Finland, July 2008.
- [8] G. Heigold, R. Schlüter, and H. Ney. On the equivalence of Gaussian HMM and Gaussian HMM-like hidden conditional random fields. In *Interspeech*, Antwerp, Belgium, Aug. 2007.
- [9] T. Jebara. *Discriminative, generative, and imitative learning*. PhD thesis, Massachusetts Institute of Technology, 2002.
- [10] T. Kemp and T. Schaaf. Estimating confidence using word lattices. In *European Conference on Speech Communication and Technology*, Rhodes, Greece, 1997.
- [11] V. Märgner and H. E. Abed. ICDAR 2007 Arabic handwriting recognition competition. In *ICDAR*, volume 2, pages 1274–1278, Sept. 2007.
- [12] V. Märgner, M. Pechwitz, and H. Abed. ICDAR 2005 Arabic handwriting recognition competition. In *ICDAR*, volume 1, pages 70–74, Seoul, Korea, Aug. 2005.
- [13] R. Nopsuwanchai and D. Povey. Discriminative training for HMM-based offline handwritten character recognition. In *ICDAR*, pages 114–118, 2003.
- [14] M. Padmanabhan, G. Saon, and G. Zweig. Lattice-based unsupervised mlr for speaker adaptation. In *ISCA ITRW Automatic Speech Recognition: Challenges for the Millenium*, Paris, France, 2000.
- [15] M. Pechwitz, S. S. Maddouri, V. Märgner, N. Ellouze, and H. Amiri. IFN/ENIT-database of handwritten Arabic words. In *Colloque International Francophone sur l’Ecrit et le Document (CIFED)*, Hammamet, Tunis, Oct. 2002.
- [16] M. Pitz, F. Wessel, and H. Ney. Improved mlr speaker adaptation using confidence measures for conversational speech recognition. In *International Conference on Spoken Language Processing*, Beijing, China, 2000.
- [17] D. Povey and P. C. Woodland. Minimum phone error and I-smoothing for improved discriminative training. In *International Conference on Acoustics, Speech, and Signal Processing*, volume 1, Orlando, FL, 2002.
- [18] J. Zhang, R. Jin, Y. Yang, and A. Hauptmann. Modified logistic regression: An approximation to SVM and its applications in large-scale text categorization. In *International Conference on Machine Learning*, Aug. 2003.