

Online optimisation of log-linear weights in interactive machine translation

Mara Chinae Rios, Germán Sanchis-Trilles, Daniel Ortiz-Martínez, Francisco Casacuberta

Pattern Recognition and Human Language Technologies Group

Universitat Politècnica de València

Camino de Vera s/n

Valencia, Spain

mchinae@iti.upv.es, gsanchis@dsic.upv.es, dortiz@iti.upv.es, fcn@dsic.upv.es

Abstract

Whenever the quality provided by a machine translation system is not enough, a human expert is required to correct the sentences provided by the machine translation system. In such a setup, it is crucial that the system is able to learn from the errors that have already been corrected. In this paper, we analyse the applicability of discriminative ridge regression for learning the log-linear weights of a state-of-the-art machine translation system underlying an interactive machine translation framework, with encouraging results.

Keywords: Interactive translation prediction, online learning, adaptive systems

1. Introduction

Adaptability is an important feature whenever statistical machine translation (SMT) systems are to be used within a computer assisted translation (CAT) framework. In such cases, the user expects the system to learn dynamically from its own errors, so that errors corrected once do not need to be corrected over and over again. Hence, the models need to be adapted *online*, i.e. without a complete retraining of the model parameters, since such retraining would be too costly. In this paper we will focus on adapting only the log-linear weights λ present in every state-of-the-art SMT system. Hence, the standard SMT equation (Brown et al., 1994) is complemented with a superindex denoting the current instant t :

$$\begin{aligned}\hat{\mathbf{y}} &= \operatorname{argmax}_{\mathbf{y}} \sum_{m=1}^M \lambda_m^t h_m(\mathbf{x}, \mathbf{y}) \\ &= \operatorname{argmax}_{\mathbf{y}} \lambda_t \cdot \mathbf{h}(\mathbf{x}, \mathbf{y}) = \operatorname{argmax}_{\mathbf{y}} g(\mathbf{x}, \mathbf{y})\end{aligned}\quad (1)$$

To simplify notation, we will omit subindex t from input sentence \mathbf{x} and output sentence $\hat{\mathbf{y}}$, although it is always assumed. $h_m(\mathbf{x}, \mathbf{y})$ represents an important feature for the translation of \mathbf{x} into \mathbf{y} , M is the number of models (or features) and λ_m are the weights acting as scaling factors of the score functions. $g(\mathbf{x}, \mathbf{y})$ represents the score of a hypothesis \mathbf{y} given an input sentence \mathbf{x} , and is not treated as a probability since the normalisation term has been omitted. Common feature functions $h_m(\mathbf{x}, \mathbf{y})$ include translation models, re-ordering models or the target language model. \mathbf{h} and λ are estimated by means of training and development sets, respectively. However, the domain of such sets has an important impact on the final translation

quality (Callison-Burch et al., 2011), and adaptation arises as an efficient way of alleviating this fact by using very limited amounts of in-domain data. We focus on two different CAT scenarios: standard post-edition (PE), and the more sophisticated interactive machine translation (IMT) (Barrachina et al., 2009) scenario. The main difference between PE and IMT is that in IMT the system provides improved completions after each user interaction, while in PE the system remains passive after providing the initial translation of the source sentence.

Similar work was performed in (Ortiz-Martínez et al., 2010), where an incremental version of the Expectation-Maximisation algorithm is used. However, such work focuses in the feature functions, while the present one focuses on the log-linear weights. Online adaptation of such weights is also studied in (Martínez-Gómez et al., 2012), although only applied to a conventional PE setup. As we will see later, the IMT setup yields challenges of its own.

2. Discriminative ridge regression

The main purpose of discriminative Ridge regression (Martínez-Gómez et al., 2012) (DRR) is that *good* hypothesis within a given N -best list score *higher*, and *bad* hypotheses score *lower*. It implements the estimation of λ as a regression problem between $g(\mathbf{x}, \mathbf{y})$, with $\mathbf{y} \in nbest(\mathbf{x})$ (i.e., the set of n best hypotheses which can be derived from \mathbf{x}), and the translation quality of \mathbf{y} , $\mu(\mathbf{y})$. Let \mathbf{y}^* be the hypothesis with the highest quality, but which might yield a lower score in Eq. 1¹. Our purpose is to adapt the

¹ \mathbf{y}^* does not necessarily match the reference translation \mathbf{y}^r due to eventual coverage problems.

model parameters so that \mathbf{y}^* is rewarded and achieves a higher score according to Eq. 1.

We define the *difference* in translation quality between the proposed hypothesis $\hat{\mathbf{y}}$ and the best hypothesis \mathbf{y}^* in terms of a given quality measure $\mu(\cdot)$:

$$l(\hat{\mathbf{y}}) = |\mu(\hat{\mathbf{y}}) - \mu(\mathbf{y}^*)|, \quad (2)$$

where the absolute value has been introduced in order to preserve generality. We also define the score difference between $\hat{\mathbf{y}}$ and \mathbf{y}^* as:

$$\phi(\hat{\mathbf{y}}) = g(\mathbf{x}, \mathbf{y}^*) - g(\mathbf{x}, \hat{\mathbf{y}}). \quad (3)$$

Ideally, we would like differences in $l(\cdot)$ to correspond to differences in $\phi(\cdot)$: if hypothesis \mathbf{y} has a translation quality $\mu(\mathbf{y})$ that is very similar to the translation quality of $\mu(\mathbf{y}^*)$, we would like this to be reflected in translation score g , i.e., $g(\mathbf{x}, \mathbf{y})$ is very similar to $g(\mathbf{x}, \mathbf{y}^*)$.

For computing the new scaling factors λ_t , the previously learnt λ_{t-1} is combined, for a certain learning rate α , with an appropriate update step $\check{\lambda}_t$, yielding (Martínez-Gómez et al., 2012):

$$\lambda_t = (1 - \alpha)\lambda_{t-1} + \alpha\check{\lambda}_t. \quad (4)$$

2.1. Discriminative ridge regression in PE

In a conventional post-editing scenario where the hypotheses are provided by a regular SMT system, the DRR algorithm requires an N -best list of hypotheses in decreasing order of score. Let $nbest(\mathbf{x})$ be such a list computed by our models for sentence \mathbf{x} . For adapting λ , we define an $N \times M$ matrix $H_{\mathbf{x}}$, where M is the number of features in Eq. 1, containing the feature functions \mathbf{h} of every hypothesis:

$$H_{\mathbf{x}} = [\mathbf{h}(\mathbf{x}, \mathbf{y}_1), \dots, \mathbf{h}(\mathbf{x}, \mathbf{y}_N)]'. \quad (5)$$

Additionally, let $H_{\mathbf{x}}^*$ be a matrix such that

$$H_{\mathbf{x}}^* = [\mathbf{h}(\mathbf{x}, \mathbf{y}^*), \dots, \mathbf{h}(\mathbf{x}, \mathbf{y}^*)]', \quad (6)$$

where all rows are identical and equal to the feature vector of the best hypothesis \mathbf{y}^* within the N -best list. Then, $R_{\mathbf{x}}$ is defined as

$$R_{\mathbf{x}} = H_{\mathbf{x}}^* - H_{\mathbf{x}}. \quad (7)$$

The key idea is to find a vector $\check{\lambda}_t$ such that differences in scores are reflected as differences in the quality of the hypotheses. That is,

$$R_{\mathbf{x}} \cdot \check{\lambda}_t \propto \mathbf{1}_{\mathbf{x}}, \quad (8)$$

where $\mathbf{1}_{\mathbf{x}}$ is a column vector of N rows such that

$$\mathbf{1}_{\mathbf{x}} = [l(\mathbf{y}_1) \dots l(\mathbf{y}_n) \dots l(\mathbf{y}_N)]', \quad \forall \mathbf{y}_i \in nbest(\mathbf{x}). \quad (9)$$

The objective is to find $\check{\lambda}_t$ such that

$$\check{\lambda}_t = \underset{\lambda}{\operatorname{argmin}} |R_{\mathbf{x}} \cdot \lambda - \mathbf{1}_{\mathbf{x}}| \quad (10)$$

$$= \underset{\lambda}{\operatorname{argmin}} \|R_{\mathbf{x}} \cdot \lambda - \mathbf{1}_{\mathbf{x}}\|^2, \quad (11)$$

where $\|\cdot\|^2$ is the Euclidean norm. Although Eqs. 10 and 11 are equivalent, Eq. 11 allows for a direct implementation thanks to the ridge regression², such that $\check{\lambda}_t$ can be computed as the solution to the overdetermined system $R_{\mathbf{x}} \cdot \check{\lambda}_t = \mathbf{1}_{\mathbf{x}}$, given by

$$\check{\lambda}_t = (R_{\mathbf{x}}' \cdot R_{\mathbf{x}} + \beta I)^{-1} R_{\mathbf{x}}' \cdot \mathbf{1}_{\mathbf{x}}, \quad (12)$$

where a small β is used as a regularisation term to stabilise $R_{\mathbf{x}}' \cdot R_{\mathbf{x}}$. $\beta = 0.01$ was used in the experiments described in this paper.

2.2. Discriminative ridge regression in IMT

In an IMT setting the quality metric to be used is no longer inherent to a single hypothesis, but to a complete wordgraph. In fact, it is quite common to measure the quality of a given IMT system by computing the amount of interactions required in order to modify the system's hypothesis so that it matches the reference. Once a single word has been introduced, the IMT system modifies the suffix, which implies that the number of interactions cannot be computed as a function of the hypothesis, but must be computed by first simulating the interaction procedure and is a function of a given wordgraph. Hence, since the metric to be optimised by online learning does not depend on a single-best hypothesis, the formulation of DRR needs to be reviewed.

At this stage, it would be reasonable to consider instead of a list of N -best hypotheses a list of N -best wordgraphs. However, the concept of N -best wordgraph is somewhat fuzzy. For this reason, instead of computing a true list of N -best wordgraphs we will obtain a set of N scaling factors λ , $\Lambda = \{\lambda^1, \dots, \lambda^n, \dots, \lambda^N\}$, and compute the wordgraph $W_{\lambda^n}(\mathbf{x})$ associated to a given input sentence \mathbf{x} and obtained for a certain set of scaling factors λ^n . The resulting wordgraphs will not constitute a true list of N -best wordgraphs, but since the purpose of DRR is to reward those hypotheses (in this case wordgraphs) that score well, and penalise those that score worse, what is really important is to have wordgraphs with good quality, and others with worse. Hence, $\mathbf{1}_{\mathbf{y}}$ will be a column vector of N rows such that

$$\mathbf{1}_{\mathbf{y}} = [l(W_{\lambda^1}(\mathbf{x})) \dots l(W_{\lambda^n}(\mathbf{x})) \dots l(W_{\lambda^N}(\mathbf{x}))]' \quad (13)$$

²Also known as Tikhonov regularisation.

where $l(W_{\lambda^n}(\mathbf{x}))$ is the quality metric associated to wordgraph $W_{\lambda^n}(\mathbf{x})$.

In addition, when considering DRR within an IMT setting matrix $H_{\mathbf{x}}$ also needs to be redefined, since the features that need to be considered in this case no longer correspond to hypotheses in the N -best list, but to the wordgraphs generated with Λ . Since a certain wordgraph $W_{\lambda^n}(\mathbf{x})$ does not have a single set of features, but rather one feature vector for each one of the paths through the wordgraph, we will consider for building $H_{\mathbf{x}}$ the feature vector \mathbf{h} of the best path in $W_{\lambda^n}(\mathbf{x})$, i.e., the feature vector of the best hypothesis in $W_{\lambda^n}(\mathbf{x})$. Abusing notation and with the purpose of keeping notation unclogged, let \mathbf{h}_{λ^n} be such feature vector. Then, $H_{\mathbf{x}}$ is defined for the IMT case as

$$H_{\mathbf{x}} = [\mathbf{h}_{\lambda^1}, \dots, \mathbf{h}_{\lambda^N}]'. \quad (14)$$

Equivalently, $H_{\mathbf{x}}^*$ is defined in this case as

$$H_{\mathbf{x}}^* = [\mathbf{h}_{\lambda^*}, \dots, \mathbf{h}_{\lambda^*}], \quad (15)$$

with \mathbf{h}_{λ^*} being the feature vector of the best hypothesis of wordgraph $W_{\lambda^*}(\mathbf{x})$, and $W_{\lambda^*}(\mathbf{x})$ the wordgraph with the best performance from among those derived from Λ .

In the present work, we explored two different strategies for computing Λ :

1. Sampling Λ from a Gaussian distribution with the mean centred on the set of weights obtained in training time.
2. Running a Simplex optimisation procedure to compute the best set of weights for each sentence. Since the Simplex algorithm is an iterative algorithm, Λ will be composed of those weights that arise in each one of the iterations.

3. Experiments

The experiments conducted in this work were performed by using the English-Spanish data provided for the 2013 Workshop on Machine Translation³. The initial set of log-linear weights was adjusted by means of MERT (Och, 2003) on the development sets of 2008-2010. This initial system will be referred to as *baseline*. Since a true IMT experiment is too costly for experimentation purposes because it would require a human evaluation, we simulated this procedure by evaluating the system with the 2011 test set, and adapting the weights online after the system's performance was assessed for each bilingual sentence. System performance was evaluated by means of the

Table 1: Results in KSMR of the different online learning strategies studied

Method	α	N	KSMR
baseline	—	—	40.6
original	0.01	5000	42.8
gaussian	0.001	201	40.9
simplex	0.0001	70	40.4

Key Stroke Mouse Action (Barrachina et al., 2009) (KSMR) ratio, which measures the amount of key strokes and mouse actions that a user would need to perform in order to transform the original hypothesis provided by the system into the reference.

The open-source SMT toolkit Moses (Koehn et al., 2007) was used for building the initial SMT system, and the wordgraphs produced by the decoder were then used for producing the completion hypotheses by using an in-house IMT engine. Note that, by tuning the SMT system on the same domain as the test data, the improvements reported in this paper are not a product of a topic-adaptation process, but rather are inherent to the online process itself, seen as the ability of the system to adapt itself to the current test set being translated.

The results of the different online learning approaches can be seen in Table 1, together with the optimal learning rates for each one of them and the size of the N -best list used ($|\Lambda|$ in the case of the Gaussian and simplex strategies). In this table, *baseline* displays the KSMR achieved without any online learning procedure enabled. *original* refers to the online learning of the log-linear weights as originally defined for PE (see Section 2.1.). *Gaussian* refers to the sampling strategy based on sampling from a Gaussian distribution (Section 2.2.). Finally, *simplex* refers to the Simplex strategy described in Section 2.2.. The learning rates were established previously in preliminary investigation. Note that the size of N is different in each one of the methods. The reason for this is that, while generating a large amount of N -best translation hypotheses and assessing their quality by means of conventional SMT metrics is relatively cheap, building N -best word-graphs, as described in Section 2.2., and assessing their quality is fairly expensive, since it requires a full IMT simulation for finding out the amount of corrections required. Furthermore, the size of Λ in simplex is not a meta-parameter which can be fine-tuned, since the simplex algorithm was run until convergence. This means that the value of 70 was an average, i.e., not the actual amount of Λ considered for every sentence.

³<http://www.statmt.org/wmt13>

Concerning the results shown in Table 1, it can be seen that the Simplex strategy is the one that yields the best results and the only one that is able to improve over the initial baseline. Note that previous work (López-Salcedo et al., 2012) also reported improvements with the Gaussian strategy. However, such improvements were obtained by considering $|\Lambda| = 500$. In the present work, similar improvements are obtained by considering only $|\Lambda| = 70$.

4. Conclusions

In the present paper we have presented two possible adaptations of the discriminative Ridge regression algorithm for its application in an interactive machine translation framework. First, we have empirically shown that the original definition is not valid in such a framework, and then we have evaluated two possible alternatives. Among them, the one providing the best results and with the least amount of computational resources is a strategy that relies on the iterations performed by the simplex algorithm when computing the optimum set of weights for each individual sentence. The results obtained do not provide large gains, but prove that there is room for improvement in this direction.

5. Acknowledgments

Work supported by the European Union 7th Framework Program (FP7/2007-2013) under the CasMaCat project (grant agreement nr. 287576), by Spanish MICINN under grant TIN2009-14205-C04-02 and by the Generalitat Valenciana under grant ALMPR (Prometeo/2009/014).

6. References

- S. Barrachina et al. 2009. Statistical approaches to computer-assisted translation. *Computational Linguistics*, 35(1):3–28.
- Peter F. Brown, Stephen Della Pietra, Vincent J. Della Pietra, and Robert L. Mercer. 1994. The mathematics of statistical machine translation: Parameter estimation. *Computational Linguistics*, 19(2):263–311.
- Chris Callison-Burch, Philipp Koehn, Christof Monz, and Omar F. Zaidan, editors. 2011. *Proceedings of the Sixth Workshop on SMT*. Association for Computational Linguistics, Edinburgh, Scotland, July.
- Philipp Koehn et al. 2007. Moses: Open source toolkit for statistical machine translation. In *Proc. of the ACL Demo and Poster Sessions, 2007*, pages 177–180, June 25–27.
- Francisco-Javier López-Salcedo, Germán Sanchis-Trilles, and Francisco Casacuberta. 2012. Online learning of log-linear weights in interactive machine translation. *Advances in Speech and Language Technologies for Iberian Languages*, 328(2012):277–286.
- P. Martínez-Gómez, G. Sanchis-Trilles, and F. Casacuberta. 2012. Online adaptation strategies for statistical machine translation in post-editing scenarios. *Pattern Recognition*, 45(9):3193–3203.
- Franz J. Och. 2003. Minimum error rate training for statistical machine translation. In *Proc. of the 41st annual conf. of the ACL*, pages 160–167, July 7–12.
- Daniel Ortiz-Martínez, Ismael García-Varea, and Francisco Casacuberta. 2010. Online learning for interactive statistical machine translation. In *Proc. of NAACL*, pages 546–554, June 2–4.