

Differentiable Sampling with Flexible Reference Word Order for Neural Machine Translation

Weijia Xu

University of Maryland
weijia@cs.umd.edu

Xing Niu

University of Maryland
xingniu@cs.umd.edu

Marine Carpuat

University of Maryland
marine@cs.umd.edu

Abstract

Despite some empirical success at correcting exposure bias in machine translation, scheduled sampling algorithms suffer from a major drawback: they incorrectly assume that words in the reference translations and in sampled sequences are aligned at each time step. Our new differentiable sampling algorithm addresses this issue by optimizing the probability that the reference can be aligned with the sampled output, based on a soft alignment predicted by the model itself. As a result, the output distribution at each time step is evaluated with respect to the whole predicted sequence. Experiments on IWSLT translation tasks show that our approach improves BLEU compared to maximum likelihood and scheduled sampling baselines. In addition, our approach is simpler to train with no need for sampling schedule and yields models that achieve larger improvements with smaller beam sizes.¹

1 Introduction

Neural machine translation (NMT) models are typically trained to maximize the likelihood of reference translations (Sutskever et al., 2014; Bahdanau et al., 2015). While simple and effective, this objective suffers from the exposure bias problem (Ranzato et al., 2015): the model is only exposed to reference target sequences during training, but has to rely on its own predictions at inference. As a result, errors can accumulate along the generated sequence at inference time.

This is a well-known issue in sequential decision making (Langford and Zadrozny, 2005; Cohen and Carvalho, 2005; Kääriäinen and Langford, 2006, i.a.) and it has been addressed in past work by incorporating the previous decoding choices into the training scheme, using imitation learning (Daumé et al., 2009; Ross et al.,

2011; Bengio et al., 2015; Leblond et al., 2018) and reinforcement learning (Ranzato et al., 2015; Bahdanau et al., 2016) techniques. In this paper, we focus on a simple and computationally inexpensive family of approaches, known as Data as Demonstrator (Venkatraman et al., 2015) and scheduled sampling (Bengio et al., 2015; Goyal et al., 2017). The algorithms use a stochastic mixture of the reference words and model predictions with an annealing schedule controlling the mixture probability. Despite their empirical success in various sequence prediction tasks, they are based on an assumption that does not hold for machine translation: they assume that words in the reference translations and in sampled sequences are aligned at each time step, which results in weak and sometimes misleading training signals.

In this paper, we introduce a differentiable sampling algorithm that exposes machine translation models to their own predictions during training, and allows for differences in word order when comparing model outputs with reference translations. We compute the probability that the reference can be aligned with the sampled output using a soft alignment predicted based on the model states, so that the model will not be punished too severely for producing hypotheses that deviate from the reference, as long as the hypotheses can still be aligned with the reference.

Experiments on three IWSLT tasks (German-English, English-German and Vietnamese-English) show that our approach significantly improves BLEU compared to both maximum likelihood and scheduled sampling baselines. We also provide evidence that our approach addresses exposure bias by decoding with varying beam sizes, and show that our approach is simpler to train than scheduled sampling as it requires no annealing schedule.

¹The code is available at <https://github.com/Izecson/saml-nmt>

2 Approach

Our approach is designed to optimize the standard sequence-to-sequence model for translating a source sentence \mathbf{x} into a target sentence \mathbf{y} (Bahdanau et al., 2015). This model computes the probability of \mathbf{y} given \mathbf{x} as:

$$P(\mathbf{y} | \mathbf{x}) = \prod_{t=1}^T p(y_t | \mathbf{y}_{<t}, \mathbf{x}; \theta) \quad (1)$$

where θ represents the model parameters. Given \mathbf{x} , the model first produces a sequence of hidden representations $\mathbf{h}_{1..T}$: $\mathbf{h}_t = f(\mathbf{y}_{<t}, \mathbf{x})$, where T is the length of \mathbf{y} , and f is usually an encoder-decoder network. At each time step t , the hidden representation \mathbf{h}_t is fed to a linear projection layer $\mathbf{s}_t = \mathbf{W}\mathbf{h}_t + \mathbf{b}$ to obtain a vector of scores \mathbf{s}_t over all possible words in the vocabulary \mathcal{V} . Scores are then turned into a conditional probability distribution: $p(\cdot | \mathbf{y}_{<t}, \mathbf{x}; \theta) = \text{softmax}(\mathbf{s}_t)$.

The traditional maximum likelihood (ML) objective maximizes the log-likelihood of the training data $\mathcal{D} \equiv \{(\mathbf{x}^{(n)}, \mathbf{y}^{(n)})\}_{n=1}^N$ consisting of N pairs of source and target sentences:

$$\mathcal{J}_{ML}(\theta) = \sum_{n=1}^N \sum_{t=1}^T \log p(y_t^{(n)} | \mathbf{y}_{<t}^{(n)}, \mathbf{x}^{(n)}; \theta) \quad (2)$$

At test time, prefixes $\mathbf{y}_{<t}$ are subsequences generated by the model and therefore contain errors. By contrast, in ML training, prefixes $\mathbf{y}_{<t}$ are subsequences of reference translations. As a result, the model is never exposed to its own errors during training and errors accumulate at test time. This mismatch is known as the exposure bias problem (Ranzato et al., 2015).

2.1 Limitations in Scheduled Sampling

Bengio et al. (2015) introduced the scheduled sampling algorithm to address exposure bias. Scheduled sampling gradually replaces the reference words with sampled model predictions in the prefix used at training time. An annealing schedule controls the probability of using reference words vs. model predictions. The training objective remains the same as the ML objective, except for the nature of the prefix $\hat{\mathbf{y}}_{<t}$, which contains a mixture of reference and predicted words:

$$\mathcal{J}_{SS}(\theta) = \sum_{n=1}^N \sum_{t=1}^T \log p(y_t^{(n)} | \hat{\mathbf{y}}_{<t}^{(n)}, \mathbf{x}^{(n)}; \theta) \quad (3)$$

Despite the empirical success of scheduled sampling, one limitation is that the discontinuity of the argmax operation makes it impossible to penalize errors made in previous steps, which can lead to slow and unstable training. We address this issue using a continuous relaxation to the greedy search and sampling process, similarly to Goyal et al. (2017), which we describe in Section 2.2.

Another limitation of scheduled sampling is that it incorrectly assumes that the reference and predicted sequence are aligned by time indices which introduces additional noise to the training signal.² We address this problem with a novel differentiable sampling algorithm with an alignment based objective called soft aligned maximum likelihood (SAML). It is used in combination with maximum likelihood to define our training objective $\mathcal{J} = \mathcal{J}_{ML} + \mathcal{J}_{SAML}$, where \mathcal{J}_{ML} is computed based on reference translations, and \mathcal{J}_{SAML} is computed based on sampled translations of the same input sentences. We define \mathcal{J}_{SAML} in Section 2.3.

2.2 Differentiable Sampling

To backpropagate errors made in the previous decoding steps, we use a continuous relaxation of the discrete sampling operation similar to Goyal et al. (2017), except that we use the Straight-Through (ST) Gumbel-Softmax estimator (Jang et al., 2017; Bengio et al., 2013) instead of Gumbel-Softmax (Jang et al., 2017; Maddison et al., 2014) to better simulate the scenario at inference time.³

The Gumbel-Softmax is derived from the Gumbel-Max trick (Maddison et al., 2014), an algorithm for sampling one-hot vector $z \in \mathbb{R}^k$ from a categorical distribution (p_1, \dots, p_k) :

$$z = \text{one-hot}(\arg \max_i (\log p_i + \beta g_i)) \quad (4)$$

where g_i is the Gumbel noise drawn i.i.d from $\text{Gumbel}(0, 1)$ ⁴, and β is a hyperparameter controlling the scale of the noise. Here, the trick is used to approximate the discontinuous argmax function with the differentiable softmax:

$$\tilde{z} = \text{softmax}((\log p_i + \beta g_i)/\tau) \quad (5)$$

²<https://nlpers.blogspot.com/2016/03/a-dagger-by-any-other-name-scheduled.html>

³The Straight-Through estimator consistently outperforms the Gumbel-Softmax in preliminary experiments.

⁴ $g_i = -\log(-\log(u_i))$ and $u_i \sim \text{Uniform}(0, 1)$.

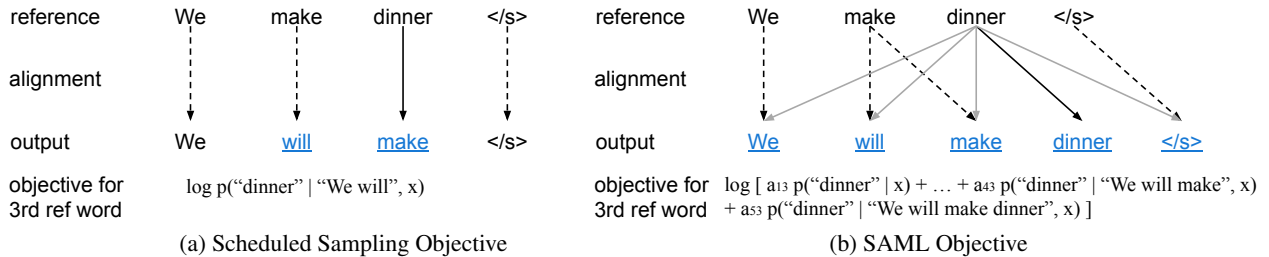


Figure 1: Difference between objectives used in scheduled sampling (left) and our approach (right), when computing the contribution to the objective of the reference word “dinner”. The schedule sampling hypothesis uses a mixture of the reference (black) and sampled (blue underlined) words, while the entire hypothesis sequence is sampled in our approach.

where τ is the temperature parameter. As τ diminishes to zero, \tilde{z} becomes the same as one-hot sample z .

The Straight-Through Gumbel-Softmax maintains the differentiability of the Gumbel-Softmax estimator while allowing for discrete sampling by taking different paths in the forward and backward pass. It uses argmax to get the one-hot sample z in the forward pass, but uses its continuous approximation \tilde{z} in the backward pass. While ST estimators are biased, they have been shown to work well in latent tree learning (Choi et al., 2018) and semi-supervised machine translation (Niu et al., 2019).

2.3 Soft Aligned Maximum Likelihood

The soft aligned maximum likelihood (SAML) is defined as the probability that the reference can be aligned with the sampled output using a soft alignment predicted by the model:

$$P_{SAML}(\mathbf{y} | \mathbf{x}) = \prod_{t=1}^T \sum_{j=1}^{T'} a_{tj} \cdot p(y_t | \tilde{\mathbf{y}}_{<j}, \mathbf{x}; \theta) \quad (6)$$

where T is the length of the reference sequence, T' is the length of the sampled sequence, a_{tj} is the predicted soft alignment between the reference word y_t and sampled prefix $\tilde{\mathbf{y}}_{<j}$.

Training with the SAML objective consists in maximizing:

$$\mathcal{J}_{SAML}(\theta) = \sum_{n=1}^N \log P_{SAML}(\mathbf{y}^{(n)} | \mathbf{x}^{(n)}) \quad (7)$$

The conditional probability of the next word $p(y_t | \tilde{\mathbf{y}}_{<j}, \mathbf{x}; \theta)$ is computed as follows:

$$p(\cdot | \tilde{\mathbf{y}}_{<j}, \mathbf{x}; \theta) = \text{softmax}(\mathbf{W}\tilde{\mathbf{h}}_j + \mathbf{b}) \quad (8)$$

where \mathbf{W} and \mathbf{b} are model parameters. $\tilde{\mathbf{h}}_j$ is the hidden representation at step j conditioned on the

Task	sentences (K)			vocab (K)	
	train	dev	test	src	tgt
de-en	153.3	7.0	6.8	113.5	53.3
vi-en	121.3	1.5	1.3	23.9	50.0

Table 1: We evaluate on two translation tasks.

source sequence \mathbf{x} and the preceding words $\tilde{\mathbf{y}}_{<j}$ sampled from the model distribution using differentiable sampling:

$$\tilde{\mathbf{h}}_j = f(\tilde{\mathbf{y}}_{<j}, \mathbf{x}) \quad (9)$$

We compute the soft alignment a_{tj} between y_t and $\tilde{\mathbf{y}}_{<j}$ based on the model’s hidden states:

$$a_{tj} = \frac{\exp(\text{score}(\tilde{\mathbf{h}}_j, \mathbf{e}_{y_t}))}{\sum_{i=1}^{T'} \exp(\text{score}(\tilde{\mathbf{h}}_i, \mathbf{e}_{y_t}))} \quad (10)$$

where \mathbf{e}_{y_t} is the embedding of the reference word y_t . The score function captures the similarity between the hidden state $\tilde{\mathbf{h}}_j$ and the embedding \mathbf{e}_{y_t} . We use the dot product here as it does not introduce additional parameters:

$$\text{score}(\mathbf{h}, \mathbf{e}) = \mathbf{h}^\top \mathbf{e} \quad (11)$$

Figure 1 illustrates how the resulting objective differs from scheduled sampling: (1) it is computed over sampled sequences as opposed to sequences that contain a mixture of sampled and reference words, and (2) each reference word is soft-aligned to the sampled sequence.

3 Experiments

Data We evaluate our approach on IWSLT 2014 German-English (de-en) as prior work (Goyal et al., 2017), as well as two additional tasks: IWSLT 2014 English-German (en-de) and IWSLT

Method	Anneal	de-en	en-de	vi-en
Baseline	No	27.41 \pm 0.26	22.64 \pm 0.13	23.59 \pm 0.13
+SS	Yes	27.47 \pm 0.28	22.56 \pm 0.17	23.97 \pm 0.39
+DSS	Yes	27.30 \pm 0.24	22.47 \pm 0.20	23.68 \pm 0.35
+SS	No	22.91 \pm 0.21	17.78 \pm 0.20	19.57 \pm 0.19
+SAML	No	27.94\pm0.12	23.30\pm0.19	24.60\pm0.35

Table 2: BLEU scores of our approach (SAML) and three baselines including the maximum likelihood (ML) baseline, scheduled sampling (SS), and differentiable scheduled sampling (DSS). The *Anneal* column indicates whether the sampling rate is annealed. For each task, we report the mean and standard deviation over 5 runs with different random seeds. SAML achieves the best BLEU scores and is simpler to train than SS and DSS, as it requires no annealing schedule.

2015 Vietnamese-English (vi-en). For de-en and en-de, we follow the preprocessing steps in Ranzato et al. (2015). For vi-en, we use the data pre-processed by Luong and Manning (2015), with test2012 for validation and test2013 for testing. Table 1 summarizes the data statistics.

Setup Our translation models are attentional RNNs (Bahdanau et al., 2015) built on Sockeye (Hieber et al., 2017). We use bi-directional LSTM encoder and single-layer LSTM decoder with 256 hidden units, embeddings of size 256, and multilayer perceptron attention with a layer size of 256. We apply layer normalization (Ba et al., 2016) and label smoothing (0.1). We add dropout to embeddings (0.1) and decoder hidden states (0.2). For ST Gumbel-Softmax, we use temperature $\gamma = 1$ and noise scale $\beta = 0.5$. The decoding beam size is 5 unless stated otherwise. We train the models using the Adam optimizer (Kingma and Ba, 2015) with a batch size of 1024 words. We checkpoint models every 1000 updates. The initial learning rate is 0.0002, and it is reduced by 30% after 4 checkpoints without validation perplexity improvement. Training stops after 12 checkpoints without improvement. For training efficiency, we first pre-train a baseline model for each task using only \mathcal{J}_{ML} and fine-tune it using different approaches. In the fine-tuning phase, we inherit all settings except that we initialize the learning rate to 0.00002 and set the minimum number of checkpoints before early stopping to 24. We fine-tune each randomly seeded model independently.

Baselines We compare our model against three baselines: (1) a standard baseline trained with the ML objective, and models fine-tuned with (2) scheduled sampling (SS) (Bengio et al., 2015) and (3) differentiable scheduled sampling (DSS) (Goyal et al., 2017). In SS and DSS, the probability of using reference words ϵ_s is annealed using inverse sigmoid decay (Bengio et al., 2015): $\epsilon_s = k/(k + \exp(i/k))$ at the i -th checkpoint with $k = 10$.

Results Table 2 shows that the SAML improves over the ML baseline by +0.5 BLEU on de-en, +0.7 BLEU on en-de, and +1.0 BLEU on vi-en task. In addition, SAML consistently improves over both the scheduled sampling and differentiable scheduled sampling on all tasks. All improvements are significant with $p < 0.002$. Interestingly, differentiable scheduled sampling performs no better than scheduled sampling in our experiments, unlike in Goyal et al. (2017).

Unlike scheduled sampling, our approach does not require an annealing schedule, and it is therefore simpler to train. We verify that the annealing schedule is needed in scheduled sampling by training a contrastive model with the same objective as scheduled sampling, but without annealing schedule (Table 2). We set the sampling rate to 0.5. The contrastive model hurts BLEU scores by at least 4.0 points compared to both the ML baseline and models fine-tuned with scheduled sampling, confirming that scheduled sampling needs the annealing schedule to work well.

We further examine the performance gain of different approaches over the baseline with varying beam sizes (Figure 2). Our approach yields larger BLEU improvements when decoding with greedy search and smaller beams, while there is no clear pattern for scheduled sampling models. These results support the hypothesis that our approach mitigates exposure bias, as it yields bigger improvements in settings where systems have fewer opportunities to recover from early errors.

4 Related Work

Daumé et al. (2009) first addressed exposure bias in an imitation learning framework by training a classifier on examples generated using a mixture of the ground truth and the model’s current predictions. DAGger (Ross et al., 2011) is a similar algorithm which differs in how the training examples are generated and aggregated. Both al-

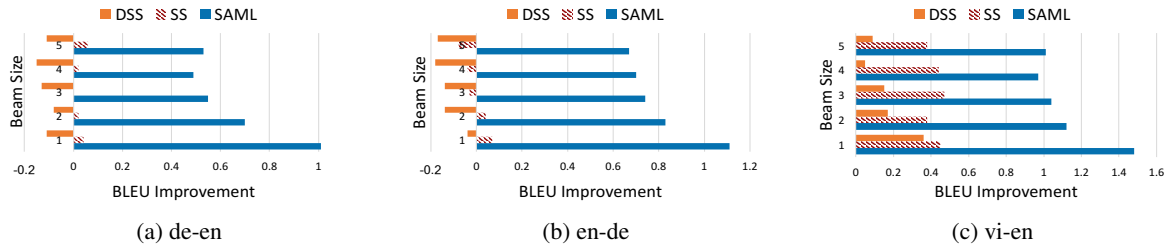


Figure 2: Improvements from our method (SAML), scheduled sampling (SS), and differentiable scheduled sampling (DSS) over the maximum likelihood (ML) baseline when decoding with varying beam sizes (average of 5 runs). The SAML model consistently yields the largest improvements with smaller beams.

gorithms require an expert policy, which produces the best next token given any model predicted prefix, and assume that policy can be efficiently computed from the reference. However, for structured prediction tasks such as machine translation with large vocabulary and complex loss functions, it is intractable to find the best next token given any prefix. For time series modeling, the Data as Demonstrator algorithm (Venkatraman et al., 2015) derives the expert policy directly from the reference sequences which are aligned with the sampled sequences at each time step. Scheduled sampling algorithms (Bengio et al., 2015; Goyal et al., 2017) use the same strategy to train neural sequence-to-sequence models for a broader range of language generation tasks, even though the time alignment between reference and sampled sequences does not hold. Leblond et al. (2018) proposed to complete a predicted prefix with all possible reference suffixes and picking the reference suffix that yields the highest BLEU-1 score. However, they found that this approach performs well only when the prefix is close to the reference.

Reinforcement learning (RL) algorithms (Bahdanau et al., 2016; Sutton and Barto, 2018; Van Hasselt et al., 2016) address exposure bias by directly optimizing a sentence-level reward for the model generated sequences. Evaluation metrics such as BLEU can be used as rewards, but they are discontinuous and hard to optimize. Techniques such as policy gradient (Williams, 1992) and actor-critic (Sutton and Barto, 2018; Degris et al., 2012) are thus required to find an unbiased estimation of the gradient to optimize the model. Due to the high variance of the gradient estimation, training with RL can be slow and unstable (Henderson et al., 2018; Wu et al., 2018). Recent alternatives use data augmentation to incorporate the sentence-level reward into the training objective more efficiently (Norouzi et al., 2016).

Finally, our SAML loss shares the idea of flexible reference word order with the bag-of-words loss introduced by Ma et al. (2018) to improve source coverage. However, their loss is computed with teacher forcing and therefore does not address exposure bias.

5 Conclusion

We introduced a differentiable sampling algorithm which exposes a sequence-to-sequence model to its own predictions during training and compares them to reference sequences flexibly to back-propagate reliable error signals. By soft aligning reference and sampled sequences, our approach consistently improves BLEU over maximum likelihood and scheduled sampling baselines on three IWSLT tasks, with larger improvements for greedy search and smaller beam sizes. Our approach is also simple to train, as it does not require any sampling schedule.

Acknowledgments

We thank the anonymous reviewers, Amr Sharaf, Naomi Feldman, Hal Daumé III and the CLIP lab at UMD for helpful comments. This research is supported in part by an Amazon Web Services Machine Learning Research Award and by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via contract #FA8650-17-C-9117. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- Lei Jimmy Ba, Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer normalization. *CoRR*, abs/1607.06450.
- Dzmitry Bahdanau, Philemon Brakel, Kelvin Xu, Anirudh Goyal, Ryan Lowe, Joelle Pineau, Aaron C. Courville, and Yoshua Bengio. 2016. An actor-critic algorithm for sequence prediction. *CoRR*, abs/1607.07086.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of the 3th International Conference on Learning Representations*.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems 28*, pages 1171–1179.
- Yoshua Bengio, Nicholas Léonard, and Aaron C. Courville. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *CoRR*, abs/1308.3432.
- Jihun Choi, Kang Min Yoo, and Sang-goo Lee. 2018. Learning to compose task-specific tree structures. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, pages 5094–5101. AAAI Press.
- William W. Cohen and Victor R. Carvalho. 2005. Stacked Sequential Learning. In *Proceedings of the IJCAI*.
- Hal Daumé, John Langford, and Daniel Marcu. 2009. Search-based structured prediction. *Machine Learning*, 75(3):297–325.
- Thomas Degris, Patrick M Pilarski, and Richard S Sutton. 2012. Model-free reinforcement learning with continuous action in practice. In *American Control Conference (ACC), 2012*, pages 2177–2182. IEEE.
- Kartik Goyal, Chris Dyer, and Taylor Berg-Kirkpatrick. 2017. Differentiable scheduled sampling for credit assignment. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 366–371. Association for Computational Linguistics.
- Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. 2018. Deep reinforcement learning that matters. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*. AAAI Press.
- Felix Hieber, Tobias Domhan, Michael Denkowski, David Vilar, Artem Sokolov, Ann Clifton, and Matt Post. 2017. Sockeye: A toolkit for neural machine translation. *CoRR*, abs/1712.05690.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical reparameterization with gumbel-softmax. In *Proceedings of the 5th International Conference on Learning Representations*.
- Matti Kääriäinen and John Langford. 2006. Lower Bounds for Reductions. In *Talk at the Atomic Learning Workshop*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of the 3th International Conference on Learning Representations*.
- John Langford and Bianca Zadrozny. 2005. Relating reinforcement learning performance to classification performance. In *Proceedings of the 22nd international conference on Machine learning*, pages 473–480. ACM.
- Rémi Leblond, Jean-Baptiste Alayrac, Anton Osokin, and Simon Lacoste-Julien. 2018. SEARNN: Training RNNs with global-local losses. In *International Conference on Learning Representations*.
- Minh-Thang Luong and Christopher D Manning. 2015. Stanford neural machine translation systems for spoken language domains. In *Proceedings of the International Workshop on Spoken Language Translation*, pages 76–79.
- Shuming Ma, Xu Sun, Yizhong Wang, and Junyang Lin. 2018. Bag-of-words as target for neural machine translation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 332–338. Association for Computational Linguistics.
- Chris J. Maddison, Daniel Tarlow, and Tom Minka. 2014. A* sampling. In *Advances in Neural Information Processing Systems 27*, pages 3086–3094. Curran Associates, Inc.
- Xing Niu, Weijia Xu, and Marine Carpuat. 2019. Bidirectional differentiable input reconstruction for low-resource neural machine translation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics.
- Mohammad Norouzi, Samy Bengio, Zhifeng Chen, Navdeep Jaitly, Mike Schuster, Yonghui Wu, and Dale Schuurmans. 2016. Reward augmented maximum likelihood for neural structured prediction. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 1723–1731. Curran Associates, Inc.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. *CoRR*, abs/1511.06732.

- Stephane Ross, Geoffrey Gordon, and Drew Bagnell. 2011. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pages 627–635. PMLR.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems 27*, pages 3104–3112.
- Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- Hado Van Hasselt, Arthur Guez, and David Silver. 2016. Deep reinforcement learning with double q-learning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, volume 2, page 5. AAAI Press.
- Arun Venkatraman, Martial Hebert, and J. Andrew Bagnell. 2015. Improving multi-step prediction of learned time series models. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 3024–3030. AAAI Press.
- Ronald J. Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3):229–256.
- Lijun Wu, Fei Tian, Tao Qin, Jianhuang Lai, and Tie-Yan Liu. 2018. A study of reinforcement learning for neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3612–3621. Association for Computational Linguistics.