

# Accelerating NMT Batched Beam Decoding with LMBR Posteriors for Deployment

Gonzalo Iglesias<sup>†</sup> William Tambellini<sup>†</sup> Adrià De Gispert<sup>†‡</sup> Eva Hasler<sup>†</sup> Bill Byrne<sup>†‡</sup>

<sup>†</sup>SDL Research

{giglesias|wtambellini|agispert|ehasler|bbyrne}@sdl.com

<sup>‡</sup>Department of Engineering, University of Cambridge, U.K.

## Abstract

We describe a batched beam decoding algorithm for NMT with LMBR n-gram posteriors, showing that LMBR techniques still yield gains on top of the best recently reported results with Transformers. We also discuss acceleration strategies for deployment, and the effect of the beam size and batching on memory and speed.

## 1 Introduction

The advent of Neural Machine Translation (NMT) has revolutionized the market. Objective improvements (Sutskever et al., 2014; Bahdanau et al., 2015; Sennrich et al., 2016b; Gehring et al., 2017; Vaswani et al., 2017) and a fair amount of neural hype have increased the pressure on companies offering Machine Translation services to shift as quickly as possible to this new paradigm.

Such a radical change entails non-trivial challenges for deployment; consumers certainly look forward to better translation quality, but do not want to lose all the good features that have been developed over the years along with SMT technology. With NMT, real time decoding is challenging without GPUs, and still an avenue for research (Devlin, 2017). Great speeds have been reported by Junczys-Dowmunt et al. (2016) on GPUs, for which batching queries to the neural model is essential. Disk usage and memory footprint of pure neural systems are certainly lower than that of SMT systems, but at the same time GPU memory is limited and high-end GPUs are expensive.

Further to that, consumers still need the ability to constrain translations; in particular, brand-related information is often as important for companies as translation quality itself, and is currently under investigation (Chatterjee et al., 2017; Hokamp and Liu, 2017; Hasler et al., 2018). It is also well known that pure neural systems

reach very high fluency, often sacrificing adequacy (Tu et al., 2017; Zhang et al., 2017; Koehn and Knowles, 2017), and have been reported to behave badly under noisy conditions (Belinkov and Bisk, 2018). Stahlberg et al. (2017) show an effective way to counter these problems by taking advantage of the higher adequacy inherent to SMT systems via Lattice Minimum Bayes Risk (LMBR) decoding (Tromble et al., 2008). This makes the system more robust to pitfalls, such as over- and under-generation (Feng et al., 2016; Meng et al., 2016; Tu et al., 2016) which is important for commercial applications.

In this paper, we describe a batched beam decoding algorithm that uses NMT models with LMBR n-gram posterior probabilities (Stahlberg et al., 2017). Batching in NMT beam decoding has been mentioned or assumed in the literature, e.g. (Devlin, 2017; Junczys-Dowmunt et al., 2016), but to the best of our knowledge it has not been formally described, and there are interesting aspects for deployment worth taking into consideration.

We also report on the effect of LMBR posteriors on state-of-the-art neural systems, for five translation tasks. Finally, we discuss how to prepare (LMBR-based) NMT systems for deployment, and how our batching algorithm performs in terms of memory and speed.

## 2 Neural Machine Translation and LMBR

Given a source sentence  $\mathbf{x}$ , a sequence-to-sequence NMT model scores a candidate translation sentence  $\mathbf{y} = y_1^T$  with  $T$  words as:

$$P_{NMT}(y_1^T|\mathbf{x}) = \prod_{t=1}^T P_{NMT}(y_t|y_1^{t-1}, \mathbf{x}) \quad (1)$$

where  $P_{NMT}(y_t|y_1^{t-1}, \mathbf{x})$  uses a neural function  $f_{NMT}(\cdot)$ . To account for batching  $B$  neu-

ral queries together, our abstract function takes the form of  $f_{NMT}(\mathbf{S}_{t-1}, \mathbf{y}_{t-1}, \mathbf{A})$  where  $\mathbf{S}_{t-1}$  is the previous batch state with  $B$  state vectors in rows,  $\mathbf{y}_{t-1}$  is a vector with the  $B$  preceding generated target words, and  $\mathbf{A}$  is a matrix with the annotations (Bahdanau et al., 2015) of a source sentence. The model has a vocabulary size  $V$ .

The implementation of this function is determined by the architecture of specific models. The most successful ones in the literature typically share in common an attention mechanism that determines which source word to focus on, informed by  $\mathbf{A}$  and  $\mathbf{S}_{t-1}$ . Bahdanau et al. (2015) use recurrent layers to both compute  $\mathbf{A}$  and the next target word  $y_t$ . Gehring et al. (2017) use convolutional layers instead, and Vaswani et al. (2017) prescind from GRU or LSTM layers, relying heavily on multi-layered attention mechanisms, stateful only on the translation side. Finally, this function can also represent an ensemble of neural models.

Lattice Minimum Bayes Risk decoding computes n-gram posterior probabilities from an evidence space and uses them to score a hypothesis space (Kumar and Byrne, 2004; Tromble et al., 2008; Blackwood et al., 2010). It improves single SMT systems, and also lends itself quite nicely to system combination (Sim et al., 2007; de Gispert et al., 2009). Stahlberg et al. (2017) have recently shown a way to use it with NMT decoding: a traditional SMT system is first used to create an evidence space  $\varphi_e$ , and the NMT space is then scored left-to-right with both the NMT model(s) and the n-gram posteriors gathered from  $\varphi_e$ . More formally:

$$\hat{y} = \arg \max_y \sum_{t=1}^T \left( \Theta_0 + \sum_{n=1}^4 \Theta_n P_{LMBR}(y_{t-n}^t | \varphi_e) + \lambda \log P_{NMT}(y_t | y_1^{t-1}, \mathbf{x}) \right) \quad (2)$$

For our purposes  $\mathbf{L}$  is arranged as a matrix with each row uniquely associated to an n-gram history identified in  $\varphi_e$ : each row contains scores for any word  $y$  in the NMT vocabulary.

$\mathbf{L}$  can be precomputed very efficiently, and stored in the GPU memory. The number of distinct n-gram histories is typically no more than 500 for our phrase-based decoder producing 200 hypotheses. Notice that such a matrix only containing  $P_{LMBR}$  contributions would be very sparse, but

it turns into a dense matrix with the summation of  $\Theta_0$ . Both sparse and dense operations can be performed on the GPU. We have found it more efficient to compute first all the sparse operations on CPU, and then upload to the GPU memory and sum the constant  $\Theta_0$  in GPU<sup>1</sup>.

### 3 NMT batched beam decoding

Algorithm 1 describes NMT decoding with LMBR posteriors using beam size  $B$  equal to the batch size. Lines 2-5 initialize the decoder; the number of time steps  $T$  is usually a heuristic function of the source length.  $\mathbf{q}$  will keep track of the  $B$  best scores per time step,  $\mathbf{b}$  and  $\mathbf{y}$  are indices.

Lines 7-16 are the core of the batch decoding procedure. At each time step  $t$ , given  $\mathbf{S}_{t-1}$ ,  $\mathbf{y}_{t-1}$  and  $\mathbf{A}$ ,  $f_{NMT}$  returns two matrices:  $\mathbf{P}_t$ , with size  $B \times V$ , contains log-probabilities for all possible candidates in the vocabulary given  $B$  live hypotheses.  $\mathbf{S}_t$  is the next batch state. Each row in  $\mathbf{S}_t$  is the vector state that corresponds to any candidate in the same row of  $\mathbf{P}_t$  (line 8).

Lines 9, 10 add the n-gram posterior scores. Given the indices in  $\mathbf{b}$  and  $\mathbf{y}$  it is straightforward to read the unique histories for the  $B$  open hypotheses: the topology of the hypothesis space is that of a tree because an NMT state represents the entire live hypothesis from time step 0. Note that  $b_{tj} < B$  is the index to access the previous word in  $\mathbf{y}_{t-1}$ . In effect, indices in  $\mathbf{b}$  function as backpointers, allowing to reconstruct not only n-grams per time step, but also complete hypotheses. As discussed for Equation 2, these histories are associated to rows in our matrix  $\mathbf{L}$ . Function GETMATRIXBYROWS( $\cdot$ ) simply creates a new matrix of size  $B \times V$  by fetching those  $B$  rows from  $\mathbf{L}$ . This new matrix is summed to  $\mathbf{P}_t$  (line 10).

In line 11, we get the indices and scores in  $\mathbf{P}_t + \mathbf{q}'_{t-1}$  of the top  $B$  hypotheses. These best hypotheses could come from any row in  $\mathbf{P}_t$ . For example, all  $B$  best hypotheses could have been found in row 0. In that case, the new batch state to be used in the next time step should contain copies of row 0 in the other  $B - 1$  rows. This is achieved again with GETMATRIXBYROWS( $\cdot$ ) in line 12.

Finally, lines 13-16 identify whether there are any end-of-sentence (EOS) candidates; the corre-

<sup>1</sup> Ideally we would want to keep  $\mathbf{L}$  as a sparse matrix and sum  $\Theta_0$  on-the-fly. However this is not possible with ArrayFire 3.6.

---

**Algorithm 1** Batch decoding with LMBR n-gram posteriors

---

```
1: procedure DECODENMT( $x, L$ )
2:    $T \leftarrow$  Maximum target hypothesis length
3:    $\mathbf{b}, \mathbf{y}, \mathbf{q}$  indices and scores, with  $\mathbf{b}_0 \leftarrow \mathbf{0}, \mathbf{y}_0 \leftarrow \mathbf{0}, \mathbf{q}_0 \leftarrow \mathbf{0}$ 
4:    $\mathbf{A} \leftarrow$  Annotations for source sentence  $x$ 
5:    $\mathbf{S}_0 \leftarrow$  Initial decoder state
6:    $F = \{\}$  ▷ Set of EOS survivors
7:   for  $t = 1$  to  $T$  do
8:      $\mathbf{P}_t, \mathbf{S}_t \leftarrow f_{NMT}(\mathbf{S}_{t-1}, \mathbf{y}_{t-1}, \mathbf{A})$ 
9:      $\mathbf{h} \leftarrow B$  histories identified through  $\mathbf{b}, \mathbf{y}$  and  $t$ 
10:     $\mathbf{P}_t \leftarrow \mathbf{P}_t + \text{GETMATRIXBYROWS}(\mathbf{L}, \mathbf{h})$  ▷ Add LMBR contributions
11:     $\mathbf{b}_t, \mathbf{y}_t, \mathbf{q}_t \leftarrow \text{TOPB}(\mathbf{P}_t + \mathbf{q}'_{t-1})$ 
12:     $\mathbf{S}_t \leftarrow \text{GETMATRIXBYROWS}(\mathbf{S}_t, \mathbf{b}_t)$ 
13:    for  $j = 0$  to  $B - 1$  do
14:      if  $y_{tj} = \text{EOS}$  then
15:         $F \leftarrow F \cup \{t, j, q_{tj}\}$  ▷ Track indices and score
16:         $q_{tj} \leftarrow -\infty$  ▷ Mask out to prevent hypothesis extension
17:  return  $\text{GETBESTHYPOTHESIS}(F, \mathbf{b}, \mathbf{y})$ 
```

---

sponding indices and score are pushed into stack  $F$  and these candidates are masked out (i.e. set to  $-\infty$ ) to prevent further expansion. In line 17,  $\text{GETBESTHYPOTHESIS}(F)$  traces backwards the best hypothesis in  $F$ , again using indices in  $\mathbf{b}$  and  $\mathbf{y}$ . Optionally, normalization by hypothesis length happens in this step.

It is worth noting that:

1. If we drop lines 9, 10 we have a pure left-to-right NMT batched beam decoder.
2. Applying a constraint (e.g. for lattice rescoring or other user constraints) involves masking out scores in  $\mathbf{P}_t$  before line 11.
3. Because the batch size is tied to the beam size, the memory footprint increases with the beam.
4. Due to the beam being used for both EOS and non EOS candidates, it can be argued that this empoverishes the beam and it could be kept in addition to non EOS candidates (either by using a bigger beam, or keeping separately). Empirically we have found that this does not affect quality with real models.
5. The opposite, i.e. that EOS candidates never survive in the beam for  $T$  time steps, can happen, although very infrequently. Several pragmatic backoff strategies can be applied in this situation: for example, running the decoder for additional time steps, or tracking

all EOS candidates that did not survive in a separate stack and picking the best hypothesis from there. We chose the latter.

### 3.1 Extension to Sentence batching

In addition to batching all  $B$  queries to the neural model needed to compute the next time step for one sentence, we can do *sentence batching*: this is, we translate  $N$  sentences simultaneously, batching  $B \times N$  queries per time step.

With small modifications, Algorithm 1 can be easily extended to handle sentence batching. If the number of sentences is  $N$ ,

1. Instead of one set  $F$  to store EOS candidates, we need  $F_1 \dots F_N$  sets.
2. For every time step,  $\mathbf{b}_t, \mathbf{y}_t$  and  $\mathbf{q}_t$  need to be matrices instead of vectors, and minor changes are required in  $\text{TOPB}(\cdot)$  to fetch the best candidates per sentence efficiently.
3.  $\mathbf{P}_t$  and  $\mathbf{S}_t$  can remain as matrices, in which case the new batch size is simply  $B \cdot N$ .
4. The heuristic function used to compute  $T$  is typically sentence specific.

## 4 Experiments

### 4.1 Experimental Setup

We report experiments on English-German, German-English and Chinese-English language

	WMT17			WAT	
	ger-eng	eng-ger	chi-eng	eng-jpn	jpn-eng
<b>PBMT</b>	28.9	19.6	15.8	33.4	18.0
<b>FNMT</b>	32.8	26.1	20.8	39.1	25.3
<b>LNMT</b>	33.7	26.6	22.0	40.4	26.1
<b>TNMT</b>	35.2	28.9	24.8	44.6	29.4
<b>LTNMT</b>	35.4	29.2	25.4	44.9	30.2
<b>Best submissions</b>	35.1	28.3	26.4	43.3	28.4

Table 1: Quality assessment of our NMT systems with and without LMBR posteriors for GRU-based (FNMT, LNMT) and Transformer models (TNMT, LTNMT). Cased BLEU scores reported on 5 translation tasks. The exact PBMT systems used to compute n-gram posteriors for LNMT and LTNMT systems are also reported. The last row shows scores for the best official submissions to each task.

pairs for the WMT17 task, and Japanese-English and English-Japanese for the WAT task. For the German tasks we use *news-test2013* as a development set, and *news-test2017* as a test set; for Chinese-English, we use *news-dev2017* as a development set, and *news-test2017* as a test set. For Japanese tasks we use the ASPEC corpus (Nakazawa et al., 2016).

We use all available data in each task for training. In addition, for German we use back-translation data (Sennrich et al., 2016a). All training data for neural models is preprocessed with the byte pair encoding technique described by Sennrich et al. (2016b). We use Blocks (van Merriënboer et al., 2015) with Theano (Bastien et al., 2012) to train attention-based single GRU layer models (Bahdanau et al., 2015), henceforth called **FNMT**. The vocabulary size is 50K. Transformer models (Vaswani et al., 2017), called here **TNMT**, are trained using the Tensor2Tensor package<sup>2</sup> with a vocabulary size of 30K.

Our proprietary translation system is a modular homegrown tool that supports pure neural decoding (FNMT and TNMT) and with LMBR posteriors (henceforth called **LNMT** and **LTNMT** respectively), and flexibly uses other components (phrase-based decoding, byte pair encoding, etcetera) to seamlessly deploy an end-to-end translation system.

FNMT/LNMT systems use ensembles of 3 neural models unless specified otherwise; TNMT/LTNMT systems decode with 1 to 2 models, each averaging over the last 20 checkpoints.

The Phrase-based decoder (**PBMT**) uses standard features with one single 5-gram language

<sup>2</sup><https://github.com/tensorflow/tensor2tensor>

model (Heafield et al., 2013), and is tuned with standard MERT (Och, 2003); n-gram posterior probabilities are computed on-the-fly over rich translation lattices, with size bounded by the PBMT stack and distortion limits. The parameter  $\lambda$  in Equation 2 is set as 0.5 divided by the number of models in the ensemble. Empirically we have found this to be a good setting in many tasks.

Unless noted otherwise, the beam size is set to 12 and the NMT beam decoder always batches queries to the neural model. The beam decoder relies on an early preview of ArrayFire 3.6 (Yalamanchili et al., 2015)<sup>3</sup>, compiled with CUDA 8.0 libraries. For speed measurements, the decoder uses one single CPU thread. For hardware, we use an Intel Xeon CPU E5-2640 at 2.60GHz. The GPU is a GeForce GTX 1080Ti. We report cased BLEU scores (Papineni et al., 2002), strictly comparable to the official scores in each task<sup>4</sup>.

#### 4.2 The effect of LMBR n-gram posteriors

Table 1 shows contrastive experiments for all five language pair/tasks. We make the following observations:

1. LMBR posteriors show consistent gains on top of the GRU model (LNMT vs FNMT rows), ranging from +0.5BLEU to +1.2BLEU. This is consistent with the findings reported by Stahlberg et al. (2017).
2. The TNMT system boasts improvements across the board, ranging from +1.5BLEU in German-English to an impressive +4.2BLEU in English-Japanese WAT

<sup>3</sup><http://arrayfire.org>

<sup>4</sup><http://matrix.statmt.org/> and <http://lotus.kuee.kyoto-u.ac.jp/WAT/evaluation/index.html>

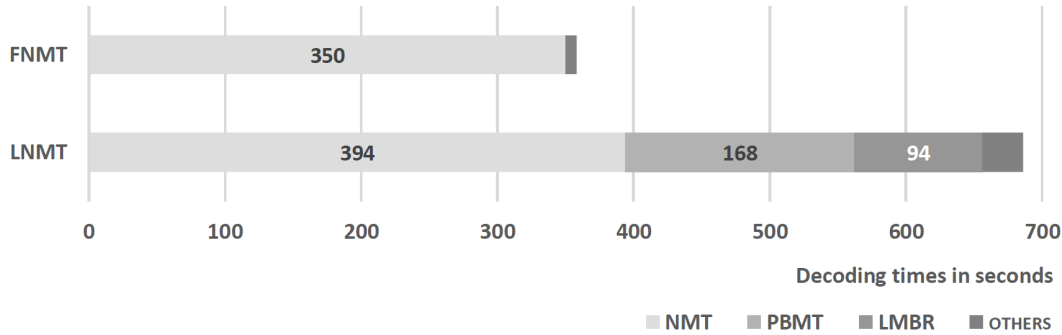


Figure 1: Accelerated FNMT and LNMT decoding times for newstest-2017 test set.

(TNMT vs LNMT). This is in line with findings by Vaswani et al. (2017) and sets new very strong baselines to improve on.

- Further, applying LMBR posteriors along with the Transformer model yields gains in all tasks (LTNMT vs TNMT), up to +0.8BLEU in Japanese-English. Interestingly, while we find that rescoring PBMT lattices (Stahlberg et al., 2016) with GRU models yields similar improvements to those reported by Stahlberg et al. (2017), we did not find gains when rescoring with the stronger TNMT models instead.

#### 4.3 Accelerating FNMT and LNMT systems for deployment

There is no particular constraint on speed for the research systems reported in Table 1. We now address the question of deploying NMT systems so that MT users get the best quality improvements at real-time speed and with acceptable memory footprint. As an example, we analyse in detail the English-German FNMT and LNMT case and discuss the main trade-offs if one wanted to accelerate them. Although the actual measurements vary across all our productised NMT engines, the trends are similar to the ones reported here.

In this particular case we specify a beam width of 0.01 for early pruning (Wu et al., 2016; Delaney et al., 2006) and reduce the beam size to 4. We also shrink the ensemble into one single big model<sup>5</sup> using the data-free shrinking method described by Stahlberg and Byrne (2017), an inexpensive way to improve both speed and GPU memory footprint.

<sup>5</sup>The file size of each 3 individual models of the ensemble is 510MB; the size of the shrunken model is 1.2GB.

In addition, for LNMT systems we tune phrase-based decoder parameters such as the distortion limit, the number of translations per source phrase and the stack limit. To compute n-gram posteriors we now only take a 200-best from the phrase-based translation lattice.

Table 2 shows a contrast of our English-German WMT17 research systems versus the respective accelerated ones.

	Research		Accelerated	
	BLEU	speed	BLEU	speed
FNMT	26.1	2207	25.2	9449
LNMT	26.6	263	25.7	4927

Table 2: Cased BLEU scores for *research* vs *accelerated* English-to-German WMT17 systems. Speed reported in words per minute.

In the process, both accelerated systems have lost 0.9 BLEU relative to the baseline. As an example, let us break down the effects of accelerating the LNMT system: using only 200-best hypotheses from the phrase-based translation lattice reduces 0.3 BLEU. Replacing the ensemble with a data-free shrunken model reduces another 0.2 BLEU and decreasing the beam size reduces 0.4 BLEU. The impact of reducing the beam size varies from system to system, although often does not result in substantial quality loss for NMT models (Britz et al., 2017).

It is worth noting that these two systems share exactly the same neural model and parameter values. However, LNMT runs 4500 words per minute (wpm) slower than FNMT. Figure 1 breaks down the decoding times for both the accelerated FNMT and LNMT systems. The LNMT pipeline also requires a phrase-based decoder and the extra component to compute the n-gram posterior probabil-

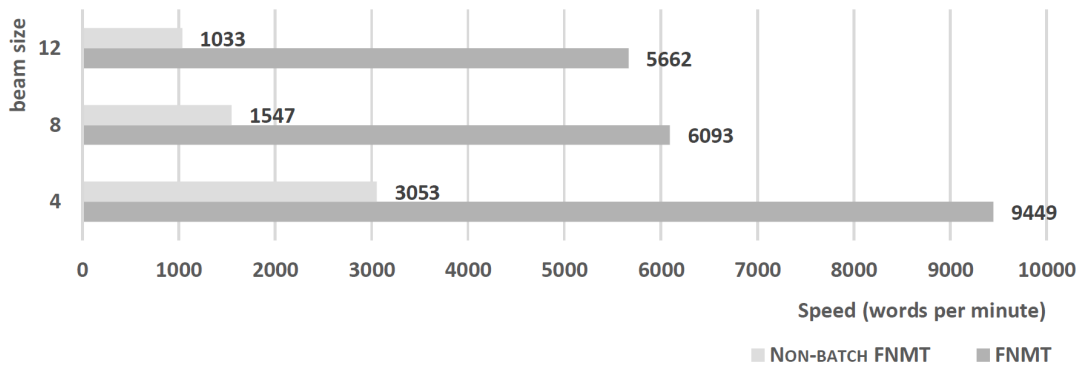


Figure 2: Batch beam decoder speed measured over newstest-2017 test set, using the accelerated FNMT system (25.2 BLEU for beam size = 4).

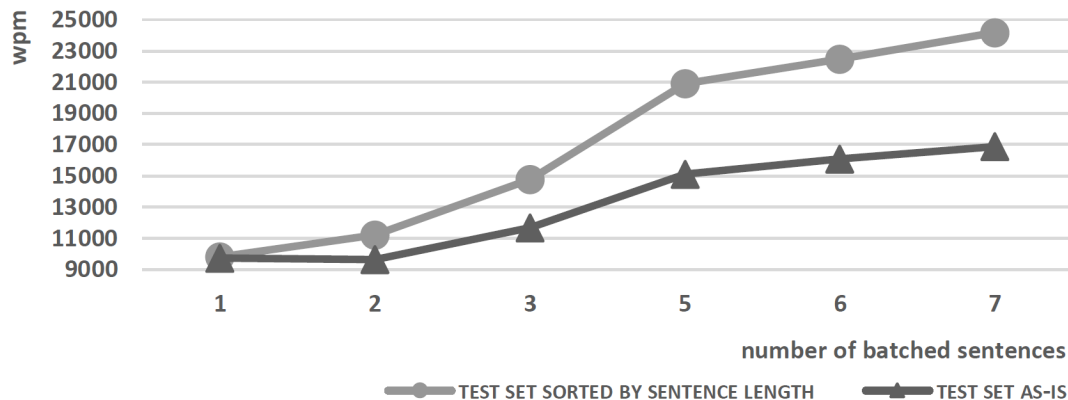


Figure 3: Batch beam decoder speed measured over newstest-2017 test set, using the accelerated eng-ger-wmt17 FNMT system (26.1 BLEU) with additional sentence batching, up to 7 sentences.

ities. In effect, while both are remarkably fast by themselves (e.g. the phrase-based decoder is running at 20000 wpm), these extra contributions explain most of the speed reduction for the accelerated LNMT system. In addition, the beam decoder itself is slightly slower for LNMT than for FNMT. This is mainly due to the computation of  $L$  as explained in Section 2. Finally, the respective GPU memory footprints for FNMT and LNMT are 4.1 and 4.8 GB.

#### 4.4 Batched beam decoding and beam size

We next discuss the impact of using batch decoding and the beam size. To this end we use the accelerated FNMT system (25.2 BLEU, 9449 wpm) to decode with and without batching; we also widen the beam. Figure 2 shows the results.

The accelerated system itself with batched beam decoding and beam size of 4 is 3 times faster than without batching (3053 wpm). The GPU memory footprint is 1 GB bigger when batching

(4.1 vs 3.1 GB). As can be expected, widening the beam decreases the speed of both decoders. The relative speed-up ratio favours the batch decoder for wider beams, i.e. it is 5 times faster for beam size 12. However, because the batch size is tied to the beam size, this comes at a cost in GPU memory footprint (under 8 GB).

#### 4.5 Sentence batching

As described in Section 3.1, it is straightforward to extend beam batching to sentence batching. Figure 3 shows the effect of sentence batching up to 7 sentences on our accelerated FNMT system.

Whilst the speed-up of our implementation is sub-linear, when batching 5 sentences the decoder runs at almost 21000 wpm, and goes beyond 24000 for 7 sentences. Thus, our implementation of sentence batching is 2.5 times faster on top of beam batching. Again, this comes at a cost: the GPU memory footprint increases as we batch more and more sentences together, up to 11 GB for

7 sentences, which approaches the limit of GPU memory.

Note that sentence batching does not change translation quality. For example, when translating 7 sentences, we are effectively batching 28 neural queries per time step. Indeed, each individual sentence is still being translated with a beam size of 4.

Figure 3 also shows the effect of sorting the test set by sentence length. Because sentences have similar lengths, less padding is required and hence we have less wasteful GPU computation. With 7 batched sentences the decoder would run at barely 17000 wpm, this is, 7000 wpm less due to not sorting by sentence length. A similar strategy is common for neural training (Sutskever et al., 2014; Morishita et al., 2017).

## 5 Conclusions

We have described a left-to-right batched beam NMT decoding algorithm that is transparent to the neural model and can be combined with LMBR n-gram posteriors. Our quality assessment with Transformer models (Vaswani et al., 2017) has shown that LMBR posteriors can still improve such a strong baseline in terms of BLEU. Finally, we have also discussed our acceleration strategy for deployment and the effect of batching and the beam size on memory and speed.

## References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). *ICLR*.
- Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, James Bergstra, Ian J. Goodfellow, Arnaud Bergeron, Nicolas Bouchard, and Yoshua Bengio. 2012. Theano: new features and speed improvements. Deep Learning and Unsupervised Feature Learning NIPS 2012 Workshop.
- Yonatan Belinkov and Yonatan Bisk. 2018. [Synthetic and natural noise both break neural machine translation](#). *ICLR*.
- Graeme Blackwood, Adrià Gispert, and William Byrne. 2010. [Efficient path counting transducers for minimum bayes-risk decoding of statistical machine translation lattices](#). In *Proceedings of ACL*, pages 27–32.
- Denny Britz, Anna Goldie, Minh-Thang Luong, and Quoc Le. 2017. [Massive exploration of neural machine translation architectures](#). In *Proceedings of EMNLP*, pages 1442–1451.
- Rajen Chatterjee, Matteo Negri, Marco Turchi, Marcello Federico, Lucia Specia, and Frédéric Blain. 2017. [Guiding neural machine translation decoding with external knowledge](#). In *Proceedings of WMT*, pages 157–168.
- Brian Delaney, Wade Shen, and Timothy Anderson. 2006. [An efficient graph search decoder for phrase-based statistical machine translation](#). In *Proceedings of IWSLT*.
- Jacob Devlin. 2017. [Sharp models on dull hardware: Fast and accurate neural machine translation decoding on the cpu](#). In *Proceedings of EMNLP*, pages 2820–2825.
- Shi Feng, Shujie Liu, Nan Yang, Mu Li, Ming Zhou, and Kenny Q. Zhu. 2016. [Improving attention modeling with implicit distortion and fertility for machine translation](#). In *Proceedings of COLING*, pages 3082–3092.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. [Convolutional sequence to sequence learning](#). *CoRR*, abs/1705.03122.
- Adrià de Gispert, Sami Virpioja, Mikko Kurimo, and William Byrne. 2009. [Minimum bayes risk combination of translation hypotheses from alternative morphological decompositions](#). In *Proceedings of NAACL-HLT*, pages 73–76.
- Eva Hasler, Adrià de Gispert, Gonzalo Iglesias, and Bill Byrne. 2018. [Neural machine translation decoding with terminology constraints](#). In *Proceedings of NAACL-HLT*.
- Kenneth Heafield, Ivan Pouzyrevsky, Jonathan H. Clark, and Philipp Koehn. 2013. [Scalable modified kneser-ney language model estimation](#). In *Proceedings of ACL*, pages 690–696.
- Chris Hokamp and Qun Liu. 2017. [Lexically constrained decoding for sequence generation using grid beam search](#). In *Proceedings of ACL*, pages 1535–1546.
- Marcin Junczys-Dowmunt, Tomasz Dwojak, and Hieu Hoang. 2016. [Is neural machine translation ready for deployment? A case study on 30 translation directions](#). *CoRR*, abs/1610.01108.
- Philipp Koehn and Rebecca Knowles. 2017. [Six challenges for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39.
- Shankar Kumar and William Byrne. 2004. [Minimum bayes-risk decoding for statistical machine translation](#). In *Proceedings of NAACL-HLT*, pages 169–176.
- Fandong Meng, Zhengdong Lu, Hang Li, and Qun Liu. 2016. [Interactive attention for neural machine translation](#). In *Proceedings of COLING*, pages 2174–2185.

- Bart van Merriënboer, Dzmitry Bahdanau, Vincent Dumoulin, Dmitriy Serdyuk, David Warde-Farley, Jan Chorowski, and Yoshua Bengio. 2015. [Blocks and fuel: Frameworks for deep learning](#). *CoRR*, abs/1506.00619.
- Makoto Morishita, Yusuke Oda, Graham Neubig, Koichiro Yoshino, Katsuhito Sudoh, and Satoshi Nakamura. 2017. [An empirical study of mini-batch creation strategies for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 61–68.
- Toshiaki Nakazawa, Manabu Yaguchi, Kiyotaka Uchimoto, Masao Utiyama, Eiichiro Sumita, Sadao Kurohashi, and Hitoshi Isahara. 2016. [ASPEC: Asian scientific paper excerpt corpus](#). In *Proceedings of LREC*.
- Franz Josef Och. 2003. [Minimum error rate training in statistical machine translation](#). In *Proceedings of ACL*, pages 160–167.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of ACL*, pages 311–318.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. [Improving neural machine translation models with monolingual data](#). In *Proceedings of ACL*, pages 86–96.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. [Neural machine translation of rare words with subword units](#). In *Proceedings of ACL*, pages 1715–1725.
- K. C. Sim, W. J. Byrne, M. J. F. Gales, H. Sahbi, and P. C. Woodland. 2007. [Consensus network decoding for statistical machine translation system combination](#). In *Proceedings of ICASSP*, volume 4, pages 105–108.
- Felix Stahlberg and Bill Byrne. 2017. [Unfolding and shrinking neural machine translation ensembles](#). In *Proceedings of EMNLP*, pages 1946–1956.
- Felix Stahlberg, Adrià de Gispert, Eva Hasler, and Bill Byrne. 2017. [Neural machine translation by minimising the bayes-risk with respect to syntactic translation lattices](#). In *Proceedings of EACL*, volume 2, pages 362–368.
- Felix Stahlberg, Eva Hasler, Aurelien Waite, and Bill Byrne. 2016. [Syntactically guided neural machine translation](#). In *Proceedings of ACL*, pages 299–305.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. [Sequence to sequence learning with neural networks](#). In *Proceedings of NIPS*, volume 2, pages 3104–3112, Cambridge, MA, USA. MIT Press.
- Roy Tromble, Shankar Kumar, Franz Och, and Wolfgang Macherey. 2008. [Lattice Minimum Bayes-Risk decoding for statistical machine translation](#). In *Proceedings of EMNLP*, pages 620–629.
- Zhaopeng Tu, Yang Liu, Zhengdong Lu, Xiaohua Liu, and Hang Li. 2017. [Context gates for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 5:87–99.
- Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. [Modeling coverage for neural machine translation](#). In *Proceedings of ACL*, pages 76–85.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). *CoRR*, abs/1706.03762.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Lukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s neural machine translation system: Bridging the gap between human and machine translation](#). *CoRR*, abs/1609.08144.
- Pavan Yalamanchili, Umar Arshad, Zakiuddin Mohammed, Pradeep Garigipati, Peter Entschew, Brian Kloppenborg, James Malcolm, and John Melonakos. 2015. [ArrayFire - A high performance software library for parallel computing with an easy-to-use API](#).
- Jingyi Zhang, Masao Utiyama, Eiichiro Sumita, Graham Neubig, and Satoshi Nakamura. 2017. [Improving neural machine translation through phrase-based forced decoding](#). In *Proceedings of IJCNLP*, pages 152–162.