
Online and Linear-Time Attention by Enforcing Monotonic Alignments

Colin Raffel¹ Minh-Thang Luong¹ Peter J. Liu¹ Ron J. Weiss¹ Douglas Eck¹

Abstract

Recurrent neural network models with an attention mechanism have proven to be extremely effective on a wide variety of sequence-to-sequence problems. However, the fact that soft attention mechanisms perform a pass over the entire input sequence when producing each element in the output sequence precludes their use in online settings and results in a quadratic time complexity. Based on the insight that the alignment between input and output sequence elements is monotonic in many problems of interest, we propose an end-to-end differentiable method for learning monotonic alignments which, at test time, enables computing attention online and in linear time. We validate our approach on sentence summarization, machine translation, and online speech recognition problems and achieve results competitive with existing sequence-to-sequence models.

1. Introduction

Recently, the “sequence-to-sequence” framework (Sutskever et al., 2014; Cho et al., 2014) has facilitated the use of recurrent neural networks (RNNs) on sequence transduction problems such as machine translation and speech recognition. In this framework, an input sequence is processed with an RNN to produce an “encoding”; this encoding is then used by a second RNN to produce the target sequence. As originally proposed, the encoding is a single fixed-length vector representation of the input sequence. This requires the model to effectively compress all important information about the input sequence into a single vector. In practice, this often results in the model having difficulty generalizing to longer sequences than those seen during training (Bahdanau et al., 2015).

An effective solution to these shortcomings are attention

mechanisms (Bahdanau et al., 2015). In a sequence-to-sequence model with attention, the encoder produces a sequence of hidden states (instead of a single fixed-length vector) which correspond to entries in the input sequence. The decoder is then allowed to refer back to any of the encoder states as it produces its output. Similar mechanisms have been used as soft addressing schemes in memory-augmented neural network architectures (Graves et al., 2014; Sukhbaatar et al., 2015) and RNNs used for sequence generation (Graves, 2013). Attention-based sequence-to-sequence models have proven to be extremely effective on a wide variety of problems, including machine translation (Bahdanau et al., 2015; Luong et al., 2015), image captioning (Xu et al., 2015), speech recognition (Chorowski et al., 2015; Chan et al., 2016), and sentence summarization (Rush et al., 2015). In addition, attention creates an implicit soft alignment between entries in the output sequence and entries in the input sequence, which can give useful insight into the model’s behavior.

A common criticism of soft attention is that the model must perform a pass over the entire input sequence when producing each element of the output sequence. This results in the decoding process having complexity $\mathcal{O}(TU)$, where T and U are the input and output sequence lengths respectively. Furthermore, because the entire sequence must be processed prior to outputting any symbols, soft attention cannot be used in “online” settings where output sequence elements are produced when the input has only been partially observed.

The focus of this paper is to propose an alternative attention mechanism which has linear-time complexity and can be used in online settings. To achieve this, we first note that in many problems, the input-output alignment is roughly monotonic. For example, when transcribing an audio recording of someone saying “good morning”, the region of the speech utterance corresponding to “good” will always precede the region corresponding to “morning”. Even when the alignment is not strictly monotonic, it often only contains local input-output reorderings. Separately, despite the fact that soft attention allows for assignment of focus to multiple disparate entries of the input sequence, in many cases the attention is assigned mostly to a single entry. For examples of alignments with these characteristics, we refer to e.g. (Chorowski et al. 2015 Figure

¹Google Brain, Mountain View, California, USA. Correspondence to: Colin Raffel <crffel@gmail.com>.

2; Chan et al. 2016 Figure 2; Rush et al. 2015 Figure 1; Bahdanau et al. 2015 Figure 3), etc. Of course, this is not true in all problems; for example, when using soft attention for image captioning, the model will often change focus arbitrarily between output steps and will spread attention across large regions of the input image (Xu et al., 2015).

Motivated by these observations, we propose using *hard monotonic* alignments for sequence-to-sequence problems because, as we argue in section 2.2, they enable computing attention online and in linear time. Towards this end, we show that it is possible to train such an attention mechanism with a quadratic-time algorithm which computes its expected output. This allows us to continue using standard backpropagation for training while still facilitating efficient online decoding at test-time. On all problems we studied, we found these added benefits only incur a small decrease in performance compared to softmax-based attention.

The rest of this paper is structured as follows: In the following section, we develop an interpretation of soft attention as optimizing a stochastic process in expectation and formulate a corresponding stochastic process which allows for online and linear-time decoding by relying on hard monotonic alignments. In analogy with soft attention, we then show how to compute the expected output of the monotonic attention process and elucidate how the resulting algorithm differs from standard softmax attention. After giving an overview of related work, we apply our approach to the tasks of sentence summarization, machine translation, and online speech recognition, achieving results competitive with existing sequence-to-sequence models. Finally, we present additional derivations, experimental details, and ideas for future research in the appendix.

2. Online and Linear-Time Attention

To motivate our approach, we first point out that softmax-based attention is computing the expected output of a simple stochastic process. We then detail an alternative process which enables online and linear-time decoding. Because this process is nondifferentiable, we derive an algorithm for computing its expected output, allowing us to train a model with standard backpropagation while applying our online and linear-time process at test time. Finally, we propose an alternative energy function motivated by the differences between monotonic attention and softmax-based attention.

2.1. Soft Attention

To begin with, we review the commonly-used form of soft attention proposed originally in (Bahdanau et al., 2015). Broadly, a sequence-to-sequence model produces a sequence of outputs based on a processed input sequence. The model consists of two RNNs, referred to

as the “encoder” and “decoder”. The encoder RNN processes the input sequence $\mathbf{x} = \{x_1, \dots, x_T\}$ to produce a sequence of hidden states $\mathbf{h} = \{h_1, \dots, h_T\}$. We refer to \mathbf{h} as the “memory” to emphasize its connection to memory-augmented neural networks (Graves et al., 2014; Sukhbaatar et al., 2015). The decoder RNN then produces an output sequence $\mathbf{y} = \{y_1, \dots, y_U\}$, conditioned on the memory, until a special end-of-sequence token is produced.

When computing y_i , a soft attention-based decoder uses a learnable nonlinear function $a(\cdot)$ to produce a scalar value $e_{i,j}$ for each entry h_j in the memory based on h_j and the decoder’s state at the previous timestep s_{i-1} . Typically, $a(\cdot)$ is a single-layer neural network using a tanh nonlinearity, but other functions such as a simple dot product between s_{i-1} and h_j have been used (Luong et al., 2015; Graves et al., 2014). These scalar values are normalized using the softmax function to produce a probability distribution over the memory, which is used to compute a context vector c_i as the weighted sum of \mathbf{h} . Because items in the memory have a sequential correspondence with items in the input, these attention distributions create a soft alignment between the output and input. Finally, the decoder updates its state to s_i based on s_{i-1} and c_i and produces y_i . In total, producing y_i involves

$$e_{i,j} = a(s_{i-1}, h_j) \quad (1)$$

$$\alpha_{i,j} = \exp(e_{i,j}) / \sum_{k=1}^T \exp(e_{i,k}) \quad (2)$$

$$c_i = \sum_{j=1}^T \alpha_{i,j} h_j \quad (3)$$

$$s_i = f(s_{i-1}, y_{i-1}, c_i) \quad (4)$$

$$y_i = g(s_i, c_i) \quad (5)$$

where $f(\cdot)$ is a recurrent neural network (typically one or more LSTM (Hochreiter & Schmidhuber, 1997) or GRU (Chung et al., 2014) layers) and $g(\cdot)$ is a learnable nonlinear function which maps the decoder state to the output space (e.g. an affine transformation followed by a softmax when the target sequences consist of discrete symbols).

To motivate our monotonic alignment scheme, we observe that eqs. (2) and (3) are computing the expected output of a simple stochastic process, which can be formulated as follows: First, a probability $\alpha_{i,j}$ is computed independently for each entry h_j of the memory. Then, a memory index k is sampled by $k \sim \text{Categorical}(\alpha_i)$ and c_i is set to h_k . We visualize this process in fig. 1. Clearly, eq. (3) shows that soft attention replaces sampling k and assigning $c_i = h_k$ with direct computation of the expected value of c_i .

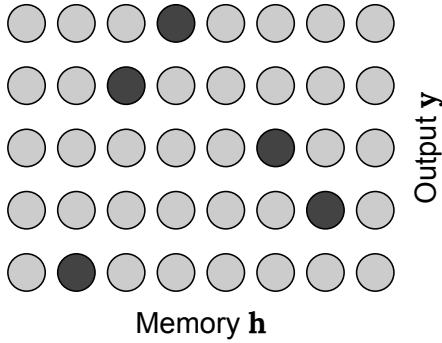


Figure 1. Schematic of the stochastic process underlying softmax-based attention decoders. Each node represents a possible alignment between an entry of the output sequence (vertical axis) and the memory (horizontal axis). At each output timestep, the decoder inspects all memory entries (indicated in gray) and attends to a single one (indicated in black). A black node indicates that memory element h_j is aligned to output y_i . In terms of which memory entry is chosen, there is no dependence across output timesteps or between memory entries.




Figure 2. Schematic of our novel monotonic stochastic decoding process. At each output timestep, the decoder inspects memory entries (indicated in gray) from left-to-right starting from where it left off at the previous output timestep and chooses a single one (indicated in black). A black node indicates that memory element h_j is aligned to output y_i . White nodes indicate that a particular input-output alignment was not considered because it violates monotonicity. Arrows indicate the order of processing and dependence between memory entries and output timesteps.

2.2. A Hard Monotonic Attention Process

The discussion above makes clear that softmax-based attention requires a pass over the entire memory to compute the terms $\alpha_{i,j}$ required to produce each element of the output sequence. This precludes its use in online settings, and results in a complexity of $\mathcal{O}(TU)$ for generating the output sequence. In addition, despite the fact that \mathbf{h} represents a transformation of a sequence (which ostensibly exhibits dependencies between subsequent elements), the attention probabilities are computed independent of temporal order and the attention distribution at the previous timestep.

We address these shortcomings by first formulating a stochastic process which explicitly processes the memory in a left-to-right manner. Specifically, for output timestep i we begin processing memory entries from index t_{i-1} , where t_i is the index of the memory entry chosen at output timestep i (for convenience, letting $t_0 = 1$). We sequentially compute, for $j = t_{i-1}, t_{i-1} + 1, t_{i-1} + 2, \dots$

$$e_{i,j} = a(s_{i-1}, h_j) \quad (6)$$

$$p_{i,j} = \sigma(e_{i,j}) \quad (7)$$

$$z_{i,j} \sim \text{Bernoulli}(p_{i,j}) \quad (8)$$

where $a(\cdot)$ is a learnable deterministic “energy function” and $\sigma(\cdot)$ is the logistic sigmoid function. As soon as we sample $z_{i,j} = 1$ for some j , we stop and set $c_i = h_j$ and $t_i = j$, “choosing” memory entry j for the context vector. Each $z_{i,j}$ can be seen as representing a discrete choice of whether to ingest a new item from the memory ($z_{i,j} = 0$) or produce an output ($z_{i,j} = 1$). For all sub-

sequent output timesteps, we repeat this process, always starting from t_{i-1} (the memory index chosen at the previous timestep). If for any output timestep i we have $z_{i,j} = 0$ for $j \in \{t_{i-1}, \dots, T\}$, we simply set c_i to a vector of zeros. This process is visualized in fig. 2 and is presented more explicitly in algorithm 1 (appendix A).

Note that by construction, in order to compute $p_{i,j}$, we only need to have computed h_k for $k \in \{1, \dots, j\}$. It follows that our novel process can be computed in an online manner; i.e. we do not need to wait to observe the entire input sequence before we start producing the output sequence. Furthermore, because we start inspecting memory elements from where we left off at the previous output timestep (i.e. at index t_{i-1}), the resulting process only computes at most $\max(T, U)$ terms $p_{i,j}$, giving it a linear runtime. Of course, it also makes the strong assumption that the alignment between the input and output sequence is strictly monotonic.

2.3. Training in Expectation

The online alignment process described above involves sampling, which precludes the use of standard backpropagation. In analogy with softmax-based attention, we therefore propose training with respect to the expected value of c_i , which can be computed straightforwardly as follows. We first compute $e_{i,j}$ and $p_{i,j}$ exactly as in eqs. (6) and (7), where $p_{i,j}$ are interpreted as the probability of choosing memory element j at output timestep i . The attention distribution over the memory is then given by (see appendix C

for a derivation)

$$\alpha_{i,j} = p_{i,j} \sum_{k=1}^j \left(\alpha_{i-1,k} \prod_{l=k}^{j-1} (1 - p_{i,l}) \right) \quad (9)$$

$$= p_{i,j} \left((1 - p_{i,j-1}) \frac{\alpha_{i,j-1}}{p_{i,j-1}} + \alpha_{i-1,j} \right) \quad (10)$$

We provide a solution to the recurrence relation of eq. (10) which allows computing $\alpha_{i,j}$ for $j \in \{1, \dots, T\}$ in parallel with cumulative sum and cumulative product operations in appendix C.1. Defining $q_{i,j} = \alpha_{i,j}/p_{i,j}$ gives the following procedure for computing $\alpha_{i,j}$:

$$e_{i,j} = a(s_{i-1}, h_j) \quad (11)$$

$$p_{i,j} = \sigma(e_{i,j}) \quad (12)$$

$$q_{i,j} = (1 - p_{i,j-1})q_{i,j-1} + \alpha_{i-1,j} \quad (13)$$

$$\alpha_{i,j} = p_{i,j}q_{i,j} \quad (14)$$

where we define the special cases of $q_{i,0} = 0, p_{i,0} = 0$ to maintain equivalence with eq. (9). As in softmax-based attention, the $\alpha_{i,j}$ values produce a weighting over the memory, which are then used to compute the context vector at each timestep as in eq. (3). However, note that α_i may not be a valid probability distribution because $\sum_j \alpha_{i,j} \leq 1$. Using α_i as-is, without normalization, effectively associates any additional probability not allocated to memory entries to an additional all-zero memory location. Normalizing α_i so that $\sum_{j=1}^T \alpha_{i,j} = 1$ has two issues: First, we can't perform this normalization at test time and still achieve online decoding because the normalization depends on $\alpha_{i,j}$ for $j \in \{1, \dots, T\}$, and second, it would result in a mismatch compared to the probability distribution induced by the hard monotonic attention process which sets c_i to a vector of zeros when $z_{i,j} = 0$ for $j \in \{t_{i-1}, \dots, T\}$.

Note that computing c_i still has a quadratic complexity because we must compute $\alpha_{i,j}$ for $j \in \{1, \dots, T\}$ for each output timestep i . However, because we are training directly with respect to the expected value of c_i , we will train our decoders using eqs. (11) to (14) and then use the online, linear-time attention process of section 2.2 at test time. Furthermore, if $p_{i,j} \in \{0, 1\}$ these approaches are equivalent, so in order for the model to exhibit similar behavior at training and test time, we need $p_{i,j} \approx 0$ or $p_{i,j} \approx 1$. We address this in section 2.5.

2.4. Modified Energy Function

While various ‘‘energy functions’’ $a(\cdot)$ have been proposed, the most common to our knowledge is the one proposed in (Bahdanau et al., 2015):

$$a(s_{i-1}, h_j) = v^\top \tanh(Ws_{i-1} + Vh_j + b) \quad (15)$$

where W and V are weight matrices, b is a bias vector,¹ and v is a weight vector. We make two modifications to eq. (15) for use with our monotonic decoder: First, while the softmax is invariant to offset,² the logistic sigmoid is not. As a result, we make the simple modification of adding a scalar variable r after the tanh function, allowing the model to learn the appropriate offset for the pre-sigmoid activations. Note that eq. (13) tends to exponentially decay attention over the memory because $1 - p_{i,j} \in [0, 1]$; we therefore initialized r to a negative value prior to training so that $1 - p_{i,j}$ tends to be close to 1. Second, the use of the sigmoid nonlinearity in eq. (12) implies that our mechanism is particularly sensitive to the scale of the energy terms $e_{i,j}$, or correspondingly, the scale of the energy vector v . We found an effective solution to this issue was to apply weight normalization (Salimans & Kingma, 2016) to v , replacing it by $gv/\|v\|$ where g is a scalar parameter. Initializing g to the inverse square root of the attention hidden dimension worked well for all problems we studied.

The above produces the energy function

$$a(s_{i-1}, h_j) = g \frac{v^\top}{\|v\|} \tanh(Ws_{i-1} + Vh_j + b) + r \quad (16)$$

The addition of the two scalar parameters g and r prevented the issues described above in all our experiments while incurring a negligible increase in the number of parameters.

2.5. Encouraging Discreteness

As mentioned above, in order for our mechanism to exhibit similar behavior when training in expectation and when using the hard monotonic attention process at test time, we require that $p_{i,j} \approx 0$ or $p_{i,j} \approx 1$. A straightforward way to encourage this behavior is to add noise before the sigmoid in eq. (12), as was done e.g. in (Frey, 1997; Salakhutdinov & Hinton, 2009; Foerster et al., 2016). We found that simply adding zero-mean, unit-variance Gaussian noise to the pre-sigmoid activations was sufficient in all of our experiments. This approach is similar to the recently proposed Gumbel-Softmax trick (Jang et al., 2016; Maddison et al., 2016), except we did not find it necessary to anneal the temperature as suggested in (Jang et al., 2016).

Note that once we have a model which produces $p_{i,j}$ which are effectively discrete, we can eschew the sampling involved in the process of section 2.2 and instead simply set $z_{i,j} = \mathbb{I}(p_{i,j} > \tau)$ where \mathbb{I} is the indicator function and τ is a threshold. We used this approach in all of our experiments, setting $\tau = 0.5$. Furthermore, at test time we do not add pre-sigmoid noise, making decoding purely deter-

¹ b is occasionally omitted, but we found it often improves performance and only incurs a modest increase in parameters, so we include it.

²That is, $\text{softmax}(e) = \text{softmax}(e + r)$ for any $r \in \mathbb{R}$.

ministic. Combining all of the above, we present our differentiable approach to training the monotonic alignment decoder in algorithm 2 (appendix A).

3. Related Work

(Luo et al., 2016) and (Zaremba & Sutskever, 2015) both study a similar framework in which a decoder RNN can decide whether to ingest another entry from the input sequence or emit an entry of the output sequence. Instead of training in expectation, they maintain the discrete nature of this decision while training and use reinforcement learning (RL) techniques. We initially experimented with RL-based training methods but were unable to find an approach which worked reliably on the different tasks we studied. Empirically, we also show superior performance to (Luo et al., 2016) on online speech recognition tasks; we did not attempt any of the tasks from (Zaremba & Sutskever, 2015). (Aharoni & Goldberg, 2016) also study hard monotonic alignments, but their approach requires target alignments computed via a separate statistical alignment algorithm in order to be trained.

As an alternative approach to monotonic alignments, Connectionist Temporal Classification (CTC) (Graves et al., 2006) and the RNN Transducer (Graves, 2012) both assume that the output sequences consist of symbols, and add an additional “null” symbol which corresponds to “produce no output”. More closely to our model, (Yu et al., 2016b) similarly add “shift” and “emit” operations to an RNN. Finally, the Segmental RNN (Kong et al., 2015) treats a segmentation of the input sequence as a latent random variable. In all cases, the alignment path is marginalized out via a dynamic program in order to obtain a conditional probability distribution over label sequences and train directly with maximum likelihood. These models either require conditional independence assumptions between output symbols or don’t condition the decoder (language model) RNN on the input sequence. We instead follow the framework of attention and marginalize out alignment paths when computing the context vectors c_i which are subsequently fed into the decoder RNN, which allows the decoder to condition on its past output as well as the input sequence. Our approach can therefore be seen as a marriage of these CTC-style techniques and attention. Separately, instead of performing an approximate search for the most probable output sequence at test time, we use hard alignments which facilitates linear-time decoding.

A related idea is proposed in (Raffel & Lawson, 2017), where “subsampling” probabilities are assigned to each entry in the memory and a stochastic process is formulated which involves keeping or discarding entries from the input sequence according to the subsampling probabilities. A dynamic program similar to the one derived in section 2.3 is

then used to compute the expected output which allows for training with standard backpropagation. Our approach differs in that we utilize an RNN decoder to construct the output sequence, and furthermore allows for output sequences which are longer than the input.

Some similar ideas to those in section 2.3 were proposed in the context of speech recognition in (Chorowski et al., 2015): First, the prior attention distributions are convolved with a bank of one-dimensional filters and then included in the energy function calculation. Second, instead of computing attention over the entire memory they only compute it over a sliding window. This reduces the runtime complexity at the expense of the strong assumption that memory locations attended to at subsequent output timesteps fall within a small window of one another. Finally, they also advocate replacing the softmax function with a sigmoid, but they then normalize by the sum of these sigmoid activations across the memory window instead of interpreting these probabilities in the left-to-right framework we use. While these modifications encourage monotonic attention, they do not explicitly enforce it, and so the authors do not investigate online decoding.

In a similar vein, (Luong et al., 2015) explore only computing attention over a small window of the memory. In addition to simply monotonically increasing the window location at each output timestep, they also consider learning a policy for producing the center of the memory window based on the current decoder state.

(Kim et al., 2017) also make the connection between soft attention and selecting items from memory in expectation. They consider replacing the softmax in standard soft attention with an elementwise sigmoid nonlinearity, but do not formulate the interpretation of addressing memory from left-to-right and the corresponding probability distributions as we do in section 2.3.

(Jaitly et al., 2015) apply standard softmax attention in online settings by splitting the input sequence into chunks and producing output tokens using the attentive sequence-to-sequence framework over each chunk. They then devise a dynamic program for finding the approximate best alignment between the model output and the target sequence. In contrast, our ingest/emit probabilities $p_{i,j}$ can be seen as adaptively chunking the input sequence (rather than providing a fixed setting of the chunk size) and we instead train by exactly computing the expectation over alignment paths.

4. Experiments

To validate our proposed approach for learning monotonic alignments, we applied it to a variety of sequence-to-sequence problems: sentence summarization, machine translation, and online speech recognition. In the follow-

ing subsections, we give an overview of the models used and the results we obtained; for more details about hyperparameters and training specifics please see appendix D. Incidentally, all experiments involved predicting discrete symbols (e.g. phonemes, characters, or words); as a result, the output of the decoder in each of our models was fed into an affine transformation followed by a softmax non-linearity with a dimensionality corresponding to the number of possible symbols. At test time, we performed a beam search over softmax predictions on all problems except machine translation. All networks were trained using standard cross-entropy loss with teacher forcing against target sequences using the Adam optimizer (Kingma & Ba, 2014). All of our decoders used the monotonic attention mechanism of section 2.3 during training to address the hidden states of the encoder. For comparison, we report test-time results using both the hard linear-time decoding method of section 2.2 and the “soft” monotonic attention distribution. We also present the results of a synthetic benchmark we used to measure the potential speedup offered by our linear-time decoding process in appendix F.

Online Speech Recognition Online speech recognition involves transcribing the words spoken in a speech utterance in real-time, i.e. as a person is talking. This problem is a natural application for monotonic alignments because online decoding is an explicit requirement. In addition, this precludes the use of bidirectional RNNs, which degrades performance somewhat (Graves et al., 2013). We tested our approach on two datasets: TIMIT (Garofolo et al., 1993) and the Wall Street Journal corpus (Paul & Baker, 1992).

Speech recognition on the TIMIT dataset involves transcribing the phoneme sequence underlying a given speech utterance. Speech utterances were represented as sequences of 40-filter (plus energy) mel-filterbank spectra, computed every 10 milliseconds, with delta- and delta-delta-features. Our encoder RNN consisted of three unidirectional LSTM layers. Following (Chan et al., 2016), after the first and second LSTM layer we placed time reduction layers which skip every other sequence element. Our decoder RNN was a single unidirectional LSTM. Our output softmax had 62 dimensions, corresponding to the 60 phonemes from TIMIT plus special start-of-sequence and end-of-sequence tokens. At test time, we utilized a beam search over softmax predictions, with a beam width of 10. We report the phone error rate (PER) after applying the standard mapping to 39 phonemes (Graves et al., 2013). We used the standard train/validation/test split and report results on the test set.

Our model’s performance, with a comparison to other online approaches, is shown in table 1. We achieve better performance than recently proposed sequence-to-sequence models (Luo et al., 2016; Jaitly et al., 2015), though the

Table 1. Phone error rate on the TIMIT dataset for different online methods.

Method	PER
(Luo et al., 2016) (stacked LSTM)	21.5%
(Jaitly et al., 2015) (end-to-end)	20.8%
(Luo et al., 2016) (grid LSTM)	20.5%
Hard Monotonic Attention (ours)	20.4%
Soft Monotonic Attention (ours, offline)	20.1%
(Graves et al., 2013) (CTC)	19.6%

small size of the TIMIT dataset and the resulting variability of results precludes making substantive claims about one approach being best. We note that (Jaitly et al., 2015) were able to improve performance by precomputing alignments using an HMM system and providing them as a supervised signal to their decoder; we did not experiment with this idea. CTC (Graves et al., 2013) still outperforms all sequence-to-sequence models. In addition, there remains a substantial gap between these online results and offline results using bidirectional LSTMs, e.g. (Chorowski et al., 2015) achieves a 17.6% phone error rate using a softmax-based attention mechanism and (Graves et al., 2013) achieved 17.7% using a pre-trained RNN transducer model. We are interested in investigating ways to close this gap in future work.

Because of the size of the dataset, performance on TIMIT is often highly dependent on appropriate regularization. We therefore also evaluated our approach on the Wall Street Journal (WSJ) speech recognition dataset, which is about 10 times larger. For the WSJ corpus, we present speech utterances to the network as 80-filter mel-filterbank spectra with delta- and delta-delta features, and normalized using per-speaker mean and variance computed offline. The model architecture is a variation of that from (Zhang et al., 2016), using an 8 layer encoder including: two convolutional layers which downsample the sequence in time, followed by one unidirectional convolutional LSTM layer, and finally a stack of three unidirectional LSTM layers interleaved with linear projection layers and batch normalization. The encoder output sequence is consumed by the proposed online attention mechanism which is passed into a decoder consisting of a single unidirectional LSTM layer followed by a softmax layer.

Our output softmax predicted one of 49 symbols, consisting of alphanumeric characters, punctuation marks, and start-of sequence, end-of-sequence, “unknown”, “noise”, and word delimiter tokens. We utilized label smoothing during training (Chorowski & Jaitly, 2017), replacing the targets at time y_t with a convex weighted combination of the surrounding five labels (full details in appendix D.1.2). Performance was measured in terms of word error rate (WER) on the test set after segmenting the model’s predic-

Table 2. Word error rate on the WSJ dataset. All approaches used a unidirectional encoder; results in grey indicate offline models.

Method	WER
CTC (our model)	33.4%
(Luo et al., 2016) (hard attention)	27.0%
(Wang et al., 2016) (CTC)	22.7%
Hard Monotonic Attention (our model)	17.4%
Soft Monotonic Attention (our model)	16.5%
Softmax Attention (our model)	16.0%

tions according to the word delimiter tokens. We used the standard dataset split of si284 for training, dev93 for validation, and eval92 for testing. We did not use a language model to improve decoding performance.

Our results on WSJ are shown in table 2. Our model, with hard monotonic decoding, achieved a significantly lower WER than the other online methods. While these figures show a clear advantage to our approach, our model architecture differed significantly from those of (Luo et al., 2016; Wang et al., 2016). We therefore additionally measured performance against a baseline model which was identical to our model except that it used softmax-based attention (which makes it quadratic-time and offline) instead of a monotonic alignment decoder. This resulted in a small decrease of 1.4% WER, suggesting that our hard monotonic attention approach achieves competitive performance while being substantially more efficient. To get a qualitative picture of our model’s behavior compared to the softmax-attention baseline, we plot each model’s input-output alignments for two example speech utterances in fig. 4 (appendix B). Both models learn roughly the same alignment, with some minor differences caused by ours being both hard and strictly monotonic.

Sentence Summarization Speech recognition exhibits a strictly monotonic input-output alignment. We are interested in testing whether our approach is also effective on problems which only exhibit approximately monotonic alignments. We therefore ran a “sentence summarization” experiment using the Gigaword corpus, which involves predicting the headline of a news article from its first sentence.

Overall, we used the model of (Liu & Pan, 2016), modifying it only so that it used our monotonic alignment decoder instead of a soft attention decoder. Because online decoding is not important for sentence summarization, we utilized bidirectional RNNs in the encoder for this task (as is standard). We expect that the bidirectional RNNs will give the model local context which may help allow for strictly monotonic alignments. The model both took as input and produced as output one-hot representations of the word IDs, with a vocabulary of the 200,000 most common words in the training set. Our encoder consisted of

Table 3. ROUGE F-measure scores for sentence summarization on the Gigaword test set of (Rush et al., 2015). (Rush et al., 2015) reports ROUGE recall scores, so we report the F-1 scores computed for that approach from (Chopra et al., 2016). As is standard, we report unigram, bigram, and longest common subsequence metrics as R-1, R-2, and R-L respectively.

Method	R-1	R-2	R-L
(Zeng et al., 2016)	27.82	12.74	26.01
(Rush et al., 2015)	29.76	11.88	26.96
(Yu et al., 2016b)	30.27	13.68	27.91
(Chopra et al., 2016)	33.78	15.97	31.15
(Miao & Blunsom, 2016)	34.17	15.94	31.92
(Nallapati et al., 2016)	34.19	16.29	32.13
(Yu et al., 2016a)	34.41	16.86	31.83
(Suzuki & Nagata, 2017)	36.30	17.31	33.88
Hard Monotonic (ours)	37.14	18.00	34.87
Soft Monotonic (ours)	38.03	18.57	35.70
(Liu & Pan, 2016)	38.22	18.70	35.74




Figure 3. Example sentence-summary pair with attention alignments for our hard monotonic model and the softmax-based attention model of (Liu & Pan, 2016). Attention matrices are displayed so that black corresponds to 1 and white corresponds to 0. The ground-truth summary is “greece pumps more money and personnel into bird flu defense”.

a word embedding matrix (which was initialized randomly and trained as part of the model) followed by four bidirectional LSTM layers. We used a single LSTM layer for the decoder. For data preparation and evaluation, we followed the approach of (Rush et al., 2015), measuring performance using the ROUGE metric.

Our results, along with the scores achieved by other approaches, are presented in table 3. While the monotonic alignment model outperformed existing models by a substantial margin, it fell slightly behind the model of (Liu & Pan, 2016) which we used as a baseline. The higher performance of our model and the model of (Liu & Pan, 2016) can be partially explained by the fact that their encoders have roughly twice as many layers as most models proposed in the literature.

For qualitative evaluation, we plot an example input-output

pair and alignment matrices for our hard monotonic attention model and the softmax-attention baseline of (Liu & Pan, 2016) in fig. 3 (an additional example is shown in fig. 6, appendix B). Most apparent is that a given word in the summary is not always aligned to the most obvious word in the input sentence; the hard monotonic decoder aligns the first four words in the summary reasonably (greek \leftrightarrow greek, government \leftrightarrow finance, approves \leftrightarrow approved, more \leftrightarrow more), but the latter four words have unexpected alignments (funds \leftrightarrow in, to \leftrightarrow for, bird \leftrightarrow measures, bird \leftrightarrow flu). We believe this is due to the ability of the multilayer bidirectional RNN encoder to reorder words in the input sequence. This effect is also apparent in fig. 6/ (appendix B), where the monotonic alignment decoder is able to produce the phrase “human rights criticism” despite the fact that the input sentence has the phrase “criticism of human rights”. Separately, we note that the softmax attention model’s alignments are extremely “soft” and non-monotonic; this may be advantageous for this problem and partially explain its slightly superior performance.

Machine Translation We also evaluated our approach on machine translation, another task which does not exhibit strictly monotonic alignments. In fact, for some language pairs (e.g. English and Japanese, English and Korean), we do not expect monotonicity at all. However, for other pairs (e.g. English and French, English and Vietnamese) only local word reorderings are required. Our translation experiments therefore involved English to Vietnamese translation using the parallel corpus of TED talks (133K sentence pairs) provided by the IWSLT 2015 Evaluation Campaign (Cettolo et al., 2015). Following (Luong & Manning, 2015), we tokenize the corpus with the default Moses tokenizer, preserve casing, and replace words whose frequencies are less than 5 by `<unk>`. As a result, our vocabulary sizes are 17K and 7.7K for English and Vietnamese respectively. We use the TED tst2012 (1553 sentences) as a validation set for hyperparameter tuning and TED tst2013 (1268 sentences) as a test set. We report results in both perplexity and BLEU.

Our baseline neural machine translation (NMT) system is the softmax attention-based sequence-to-sequence model described in (Luong et al., 2015). From that baseline, we substitute the softmax-based attention mechanism with our proposed monotonic alignment decoder. The model utilizes two-layer unidirectional LSTM networks for both the encoder and decoder.

In (Luong et al., 2015), the authors demonstrated that under their proposed architecture, a dot product-based energy function worked better than eq. (15). Since our architecture is based on that of (Luong et al., 2015), to facilitate comparison we also tested the following variant:

$$a(s_{i-1}, h_j) = g(s_{i-1}^\top W h) + r \quad (17)$$

Table 4. Performance on the IWSLT 2015 English-Vietnamese TED talks for our monotonic alignment model and the baseline softmax-attention model of (Luong & Manning, 2015).

Method	BLEU
(Luong & Manning, 2015)	23.3
Hard Monotonic, energy function eq. (16)	22.6
Hard Monotonic, energy function eq. (17)	23.0

where g and r are scalars (initialized as in section 2.4) and W is a weight matrix.

Our results are shown in Table 4. To get a better picture of each model’s behavior, we plot input-output alignments in fig. 5 (appendix B). Most noticeable is that the monotonic alignment model tends to focus attention later in the input sequence than the baseline softmax-attention model. We hypothesize that this is a way to compensate for non-monotonic alignments when a unidirectional encoder is used; i.e. the model has effectively learned to focus on words at the end of phrases which require reordering, at which point the unidirectional encoder has observed the whole phrase. This can be seen most clearly in the example on the right, where translating “a huge famine” to Vietnamese requires reordering (as suggested by the softmax-attention model’s alignment), so the hard monotonic alignment model focuses attention on the final word in the phrase (“famine”) while producing its translation. We suspect our model’s small decrease in BLEU compared to the baseline model may be due in part to this increased modeling burden.

5. Discussion

Our results show that our differentiable approach to enforcing monotonic alignments can produce models which, following the decoding process of section 2.2, provide efficient online decoding at test time without sacrificing substantial performance on a wide variety of tasks. We believe our framework presents a promising environment for future work on online and linear-time sequence-to-sequence models. We are interested in investigating various extensions to this approach, which we outline in appendix E. To facilitate experimentation with our proposed attention mechanism, we have made an example TensorFlow (Abadi et al., 2016) implementation of our approach available online³ and added a reference implementation to TensorFlow’s `tf.contrib.seq2seq` module. We also provide a “practitioner’s guide” in appendix G.

³<https://github.com/craffel/mad>

Acknowledgements

We thank Jan Chorowski, Mark Daoust, Pietro Kreitlon Carolino, Dieterich Lawson, Navdeep Jaitly, George Tucker, Quoc V. Le, Kelvin Xu, Cinjon Resnick, Melody Guan, Matthew D. Hoffman, Jeffrey Dean, Kevin Swersky, Ashish Vaswani, and members of the Google Brain team for helpful discussions and insight.

References

- Abadi, Martin, Barham, Paul, Chen, Jianmin, Chen, Zhifeng, Davis, Andy, Dean, Jeffrey, Devin, Matthieu, Ghemawat, Sanjay, Irving, Geoffrey, Isard, Michael, Kudlur, Manjunath, Levenberg, Josh, Monga, Rajat, Moore, Sherry, Murray, Derek G., Steiner, Benoit, Tucker, Paul, Vasudevan, Vijay, Warden, Pete, Wicke, Martin, Yu, Yuan, and Zheng, Xiaoqiang. TensorFlow: A system for large-scale machine learning. In *Operating Systems Design and Implementation*, 2016.
- Aharoni, Roei and Goldberg, Yoav. Sequence to sequence transduction with hard monotonic attention. *arXiv preprint arXiv:1611.01487*, 2016.
- Bahdanau, Dzmitry, Cho, Kyunghyun, and Bengio, Yoshua. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations*, 2015.
- Cettolo, Mauro, Niehues, Jan, Stüker, Sebastian, Bentivogli, Luisa, Cattoni, Roldano, and Federico, Marcello. The IWSLT 2015 evaluation campaign. In *International Workshop on Spoken Language Translation*, 2015.
- Chan, William, Jaitly, Navdeep, Le, Quoc V., and Vinyals, Oriol. Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In *International Conference on Acoustics, Speech and Signal Processing*, 2016.
- Cho, Kyunghyun, van Merriënboer, Bart, Gülçehre, Çağlar, Bahdanau, Dzmitry, Bougares, Fethi, Schwenk, Holger, and Bengio, Yoshua. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Conference on Empirical Methods in Natural Language Processing*, 2014.
- Chopra, Sumit, Auli, Michael, and Rush, Alexander M. Abstractive sentence summarization with attentive recurrent neural networks. *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2016.
- Chorowski, Jan and Jaitly, Navdeep. Towards better decoding and language model integration in sequence to sequence models. *arXiv preprint arXiv:1612.02695*, 2017.
- Chorowski, Jan, Bahdanau, Dzmitry, Serdyuk, Dmitriy, Cho, Kyunghyun, and Bengio, Yoshua. Attention-based models for speech recognition. In *Conference on Neural Information Processing Systems*, 2015.
- Chung, Junyoung, Gulcehre, Caglar, Cho, Kyunghyun, and Bengio, Yoshua. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- Foerster, Jakob, Assael, Yannis M., de Freitas, Nando, and Whiteson, Shimon. Learning to communicate with deep multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems*, 2016.
- Frey, Brendan J. Continuous sigmoidal belief networks trained using slice sampling. *Advances in neural information processing systems*, 1997.
- Garofolo, John S., Lamel, Lori F., Fisher, William M., Fiscus, Jonathon G., and Pallett, David S. DARPA TIMIT acoustic-phonetic continuous speech corpus. 1993.
- Graves, Alex. Sequence transduction with recurrent neural networks. *arXiv preprint arXiv:1211.3711*, 2012.
- Graves, Alex. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013.
- Graves, Alex, Fernández, Santiago, Gomez, Faustino, and Schmidhuber, Jürgen. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *International conference on Machine learning*, 2006.
- Graves, Alex, Mohamed, Abdel-rahman, and Hinton, Geoffrey. Speech recognition with deep recurrent neural networks. In *International Conference on Acoustics, Speech and Signal Processing*, 2013.
- Graves, Alex, Wayne, Greg, and Danihelka, Ivo. Neural Turing machines. *arXiv preprint arXiv:1410.5401*, 2014.
- Hochreiter, Sepp and Schmidhuber, Jürgen. Long short-term memory. *Neural computation*, 9(8), 1997.
- Jaitly, Navdeep, Sussillo, David, Le, Quoc V., Vinyals, Oriol, Sutskever, Ilya, and Bengio, Samy. A neural transducer. *arXiv preprint arXiv:1511.04868*, 2015.
- Jang, Eric, Gu, Shixiang, and Poole, Ben. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- Kim, Yoon, Denton, Carl, Hoang, Luong, and Rush, Alexander M. Structured attention networks. *arXiv preprint arXiv:1702.00887*, 2017.

- Kingma, Diederik and Ba, Jimmy. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kong, Lingpeng, Dyer, Chris, and Smith, Noah A. Segmental recurrent neural networks. *arXiv preprint arXiv:1511.06018*, 2015.
- Liu, Peter J. and Pan, Xin. Text summarization with TensorFlow. <http://goo.gl/16RNEu>, 2016.
- Luo, Yuping, Chiu, Chung-Cheng, Jaitly, Navdeep, and Sutskever, Ilya. Learning online alignments with continuous rewards policy gradient. *arXiv preprint arXiv:1608.01281*, 2016.
- Luong, Minh-Thang and Manning, Christopher D. Stanford neural machine translation systems for spoken language domain. In *International Workshop on Spoken Language Translation*, 2015.
- Luong, Minh-Thang, Pham, Hieu, and Manning, Christopher D. Effective approaches to attention-based neural machine translation. In *Conference on Empirical Methods in Natural Language Processing*, 2015.
- Maddison, Chris J., Mnih, Andriy, and Teh, Yee Whye. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.
- Miao, Yishu and Blunsom, Phil. Language as a latent variable: Discrete generative models for sentence compression. *arXiv preprint arXiv:1609.07317*, 2016.
- Nallapati, Ramesh, Zhou, Bowen, dos Santos, Cicero Nogueira, Gülçehre, Çağlar, and Xiang, Bing. Abstractive text summarization using sequence-to-sequence RNNs and beyond. In *Conference on Computational Natural Language Learning*, 2016.
- Paul, Douglas B. and Baker, Janet M. The design for the Wall Street Journal-based CSR corpus. In *Workshop on Speech and Natural Language*, 1992.
- Raffel, Colin and Lawson, Dieterich. Training a sub-sampling mechanism in expectation. *arXiv preprint arXiv:1702.06914*, 2017.
- Rush, Alexander M., Chopra, Sumit, and Weston, Jason. A neural attention model for abstractive sentence summarization. In *Conference on Empirical Methods in Natural Language Processing*, 2015.
- Salakhutdinov, Ruslan and Hinton, Geoffrey. Semantic hashing. *International Journal of Approximate Reasoning*, 50(7), 2009.
- Salimans, Tim and Kingma, Diederik P. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In *Advances in Neural Information Processing Systems*, 2016.
- Sukhbaatar, Sainbayar, Szlam, Arthur, Weston, Jason, and Fergus, Rob. End-to-end memory networks. In *Advances in neural information processing systems*, 2015.
- Sutskever, Ilya, Vinyals, Oriol, and Le, Quoc V. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, 2014.
- Suzuki, Jun and Nagata, Masaaki. Cutting-off redundant repeating generations for neural abstractive summarization. *arXiv preprint arXiv:1701.00138*, 2017.
- Wang, Chong, Yogatama, Dani, Coates, Adam, Han, Tony, Hannun, Awni, and Xiao, Bo. Lookahead convolution layer for unidirectional recurrent neural networks. In *Workshop Extended Abstracts of the 4th International Conference on Learning Representations*, 2016.
- Xu, Kelvin, Ba, Jimmy, Kiros, Ryan, Cho, Kyunghyun, Courville, Aaron, Salakhudinov, Ruslan, Zemel, Rich, and Bengio, Yoshua. Show, attend and tell: Neural image caption generation with visual attention. In *International Conference on Machine Learning*, 2015.
- Yu, Lei, Blunsom, Phil, Dyer, Chris, Grefenstette, Edward, and Kocisky, Tomas. The neural noisy channel. *arXiv preprint arXiv:1611.02554*, 2016a.
- Yu, Lei, Buys, Jan, and Blunsom, Phil. Online segment to segment neural transduction. In *Conference on Empirical Methods in Natural Language Processing*, 2016b.
- Zaremba, Wojciech and Sutskever, Ilya. Reinforcement learning neural Turing machines. *arXiv preprint arXiv:1505.00521*, 362, 2015.
- Zeng, Wenyuan, Luo, Wenjie, Fidler, Sanja, and Urtasun, Raquel. Efficient summarization with read-again and copy mechanism. *arXiv preprint arXiv:1611.03382*, 2016.
- Zhang, Yu, Chan, William, and Jaitly, Navdeep. Very deep convolutional networks for end-to-end speech recognition. *arXiv preprint arXiv:1610.03022*, 2016.