

Adapting End-to-End Speech Recognition for Readable Subtitles

Danni Liu, Jan Niehues, Gerasimos Spanakis

Department of Data Science and Knowledge Engineering,
Maastricht University

{danni.liu, jan.niehues, jerry.spanakis}@maastrichtuniversity.nl

Abstract

Automatic speech recognition (ASR) systems are primarily evaluated on transcription accuracy. However, in some use cases such as subtitling, verbatim transcription would reduce output readability given limited screen size and reading time. Therefore, this work focuses on ASR with output compression, a task challenging for supervised approaches due to the scarcity of training data. We first investigate a cascaded system, where an unsupervised compression model is used to post-edit the transcribed speech. We then compare several methods of end-to-end speech recognition under output length constraints. The experiments show that with limited data far less than needed for training a model from scratch, we can adapt a Transformer-based ASR model to incorporate both transcription and compression capabilities. Furthermore, the best performance in terms of WER and ROUGE scores is achieved by explicitly modeling the length constraints within the end-to-end ASR system.

1 Introduction

Automatic speech recognition (ASR) has become ubiquitous in human interaction with digital devices, such as keyboard voice inputs (He et al., 2019) and virtual home assistants (Li et al., 2017). While transcription accuracy is often the primary goal when designing ASR systems, in some use cases the readability of outputs is crucial to user experience. A prominent example is subtitling for TV. In this case, the audience need to multitask, i.e. simultaneously watch video contents, listen to speech utterances, and read subtitles. To avoid a visual overload, not every spoken word needs to be displayed. Meanwhile, the shortened subtitles must still retain the meaning of the spoken content. Moreover, large deviations from the original utterance are also undesirable, as the disagreement with auditory input would create a distraction.

To compress the subtitles, one straightforward approach is to post-process ASR transcriptions. The task of sentence compression has been well-studied (Knight and Marcu, 2002; Clarke and Lapata, 2006; Rush et al., 2015). In extractive compression (Filippova et al., 2015; Angerbauer et al., 2019), only deletion operations are performed on the input. Despite the simplicity, this approach tends to produce outputs that are less grammatical (Knight and Marcu, 2002). On the other hand, abstractive compression (Cohn and Lapata, 2008; Rush et al., 2015; Chopra et al., 2016; Yu et al., 2018) involves more sophisticated input reformulation, such as word reordering and paraphrasing (Clarke and Lapata, 2006). For the task of compressing subtitles, however, the extent of rewriting must be controlled in order to retain consistency with spoken utterances.

From a practical point of view, building a sentence compression system typically requires training corpora where the target sequences are summarized. For most languages and domains, there exists scarcely any resource suitable for supervised training. This low-resource condition is even more severe for audio inputs. To the best of our knowledge, currently there is no publicly available spoken language compression corpora.

Given the challenges outlined above, this work investigates ASR with output compression. We test our approaches on German TV subtitles. The combination of this task and the use case is to the extent of our knowledge previously unexplored.

The first contribution of this work is a comparison of cascaded and end-to-end approaches to generating compressed ASR transcriptions, where the former consists of separate ASR and compression modules, and the latter integrates transcription and compression. The experiments show that our sentence compression module trained in an unsupervised fashion tends to excessively paraphrase,

whereas the end-to-end model can be better adapted to the task of interest. Secondly, we show that, after fine-tuning on a small adaption corpus, an ASR model can perform transcription and compression simultaneously. Without being given explicit length constraints, the adapted model shows increased recognition accuracy on rare words as well as paraphrasing capabilities to produce shorter outputs. Furthermore, by explicitly encoding the length constraints, we achieve further performance gains in addition to those brought by adaptation.

2 Task

The task of creating readable subtitles for video contents has several unique properties. First, due to limited screen size and reading time, not every spoken word needs to be transcribed, especially when utterances are spoken fast. A full transcription could even hamper user experience due to poor readability. Second, although output shortening is typically realized by deleting non-essential words, the output is not only deletion-based. A real-life example from the German TV program Tagesschau¹ shown in Table 1 contain rephrasing (from “freed from” to “without”) in addition to word removal (dropping the word “ethically”). A further requirement is that the subtitles should stay reasonably authentic to the spoken contents, only modifying them when necessary. Otherwise, the disagreement with audio contents could become distracting to the audience.

Within the framework of common NLP tasks, the task of generating readable subtitles combines ASR and abstractive compression, while being subjected to the additional requirements as outlined above.

Spoken: Befreit vom fraktionszwang soll das Parlament wohl nach der Sommerpause die ethisch schwierige Frage debattieren. (Freed from pressure from the coalition party, the parliamentary should debate the ethically difficult question after the summer break.)

Subtitle: Ohne Fraktionszwang soll das Parlament wohl nach der Sommerpause die schwierige Frage debattieren. (Without pressure from the coalition party, the parliamentary should debate the difficult question after the summer break.)

Table 1: Examples of TV subtitles compared to actual spoken words. Underlined words are the differences between the spoken words and the subtitles.

¹<https://www.tagesschau.de/>

3 ASR with Output Length Constraints

3.1 Baseline ASR Model with Adaptation

For the baseline ASR model, we use the Transformer architecture (Vaswani et al., 2017) similar to that by Pham et al. (2019). As there is no spoken language compression corpus available to us that is large enough for training an end-to-end model from scratch, we first train an ASR model without output compression, and then adapt it to our task of interest using a small web-scraped corpus. In the first training stage, the model is solely trained for transcribing speech. In the fine-tuning stage, we let the model continue training at a reduced learning rate on the adaptation corpus with shortened transcriptions. The intended goal of adaptation is to let the model learn the compression task on the basis of the transcription capability acquired before.

3.2 End-to-End Length-Constrained ASR

With the baseline introduced above, the ASR model is not aware of the compression task until the adaptation step. If the model already has a sense of output length constraints earlier, i.e. when training for the ASR task, it could better utilize the abundance of training data. Motivated by this hypothesis, we inject information about the allowable output lengths using a count-down at each decoding step, as illustrated in Figure 1. In a vanilla decoder, the hidden state at position i would ingest an embedding of the previously generated token. With the count-down for target length t , decoder state y_i additionally ingests a representation of the number of allowed output tokens, $t - i$.




Figure 1: An illustration of the length countdown during decoding. The values are represented with learned embeddings or trigonometric encoding.

We explore two ways to represent the length count-down. The first one utilizes length embeddings learned during training, motivated by the approach Kikuchi et al. (2016) proposed. Given a

target sequence of length t , at decoding time step i , the input to decoder hidden state \mathbf{y}_i is based on a concatenation of the previous state \mathbf{y}_{i-1} and an embedding of the remaining length

$$\mathbf{y}_{i-1} \oplus \text{emb}(t - i), \quad (1)$$

where $\text{emb}(t - i)$ is an embedding of the number of allowed tokens. To keep the same dimensionality as that of the original word embedding, the output from Equation 1 further undergoes a linear transformation followed by the ReLU activation.

With the length embedding approach, the model learns representations of different length values during training. Therefore, learning to represent rarely-encountered lengths may be difficult.

The second method modifies the trigonometric positional encoding from the Transformer (Vaswani et al., 2017) to represent the remaining length rather than the current position. This method has been applied in summarization (Takase and Okazaki, 2019) and machine translation (Lakew et al., 2019; Niehues, 2020) to limit output lengths. Motivated by these examples from related sequence generation tasks, we explore the “backward” positional encoding in ASR models.

With the original positional encoding, for input dimension $d \in \{0, 1, \dots, D - 1\}$, the encoding at position i is defined as

$$\begin{cases} \sin(i/10000^{d/D}), & \text{if } d \text{ is even} \\ \cos(i/10000^{(d-1)/D}), & \text{if } d \text{ is odd.} \end{cases} \quad (2)$$

The backward positional encoding is the same as Equation (2), except that the current position i is replaced by the remaining length $t - i$. Given a target sequence of length t , the length encoding at decoding step i becomes

$$\begin{cases} \sin((t - i)/10000^{d/D}), & \text{if } d \text{ is even} \\ \cos((t - i)/10000^{(d-1)/D}), & \text{if } d \text{ is odd.} \end{cases} \quad (3)$$

Like the positional encoding, the length encoding is also summed together with the input embedding to decoder hidden states. Moreover, since the encoding is based on sinusoids, it can be easily extrapolated to lengths unseen during training. This is a potential advantage over the learned length embeddings.

3.3 Unsupervised Sentence Compression

Compared to training ASR models to jointly perform transcription and compression, a more

straightforward approach is to post-edit ASR outputs using a compression model. However, training such a model in a supervised fashion requires reliable target sequences. Due to the scarcity of suitable training corpora, we choose an unsupervised approach inspired by multilingual translation (Ha et al., 2017; Johnson et al., 2017).

Similar to in Niehues (2020), this approach relies on a multilingual translation system that is trained on several language pairs. At training time, language tokens are embedded together with the source and target sentences. At test time, the model is given the same source and target language token, which is a translation direction unseen in training. Since the multilingual training enables zero-shot translation, the model is able to reformulate the input in the same language. To achieve output compression, the length constraints introduced in Section 3.2 are applied in the decoder.

4 Experiment Setup

4.1 Datasets

Table 2 provides an overview of audio corpora we use. The baseline ASR model is trained on the German part of LibriVoxDeEn (Beilharz et al., 2020), a recently released corpus consisting of open-domain German audio books. Since the corpus creators did not suggest a train-dev-test partition, we split the dataset ourselves. The test set contains the following books: *Jonathan Frock*², *Jolanthes Hochzeit* and *Kammacher*.

For the spoken language compression adaptation corpus, we collect spoken utterances and subtitles from the German news program Tagesschau from 1 January to 15 August 2019. To control for recording condition and disfluency, we exclude interviews or press conferences and only keep utterances from the news anchors. The utterances are segmented based on the start and end time of subtitles. Since the timestamps do not always precisely correspond to utterance boundaries, we manually verify the test set and edit when necessary.

For the unsupervised compression system, we use the multilingual translation corpus from the IWSLT 2017 evaluation campaign (Cettolo et al., 2017). It consists of English, German, Dutch, Italian and Romanian parallel sentences based on TED

²In the recordings of *Jonathan Frock*, we found some sections with misaligned utterances therefore incorrect transcriptions. The following sections are excluded from our test set: 00006_jonathanfrock, 00007_jonathanfrock.

Dataset	Total length (h:m)	Total utterances	Average length (s)	Total words
LibriVoxDeEn (train)	469:21	206,490	8.18	3,622,560
LibriVoxDeEn (test)	5:17	2,446	7.78	51,314
Tagesschau (adapt)	37:28	11,559	11.67	243,728
Tagesschau (test)	46	213	13.01	4,864

Table 2: Corpus statistics.

talks. All 10×2 translation directions are used in training. At test time, the model is given the same source and target language tag (German in our case) in order to generate summarization in the same language. We use the positional embedding introduced in Section 3.2 for length control.

4.2 Preprocessing

We use the Kaldi toolkit (Povey et al., 2011) to preprocess the raw audio utterances into 23-dimensional filter banks. We choose not to apply any utterance-level normalization to allow for future work towards online processing. For text materials, i.e. audio transcriptions and the translation source and target sentences, we use byte-pair encoding (BPE) (Sennrich et al., 2016) to create subword-based dictionaries.

4.3 Hyperparameters

For the ASR model, we adopt many reported values in the work of Pham et al. (2019), including the optimizer choice, learning rate, warmup steps, dropout rate, label smoothing rate, and embedding dimension. There are several parameters that we choose differently. The size of the inner feed forward layer is 2048. Moreover, we use 32 encoder and 12 decoder layers, and BPE of size 10,000. For the compression model, we use a Transformer with 8 encoder and decoder layers each.³

5 Experiments

5.1 Post-Editing with Compression Model

To gain an initial understanding of the task, we start with a more controlled setup, where the test utterances are transcribed by a commercial off-the-shelf ASR system. The transcriptions are then post-processed with our compression model.

First, we analyze the level of desired output compression by contrasting the lengths of the off-the-shelf transcriptions against those of the references.

³The code is available at <https://github.com/quanpn90/NMTGMinor/tree/DbMajor>.




Figure 2: Histogram of compression rate on the test set. Ratios lower than 1 are outliers where the off-the-shelf ASR model terminates decoding prematurely.

Ground-truth: (17 words) Allerdings können Ausnahmeanträge gestellt werden, von Pendlern und von Stuttgartern, was Tausende auch bereits getan haben. (However, commuters and Stuttgarter can apply for exception-permits, which thousands have already done.)

Reference: (6 words) Es können aber Ausnahmeanträge gestellt werden. (However, exception-permits can be applied for.)

Ground-truth: (13 words) In der Debatte über neue Regeln für Organspenden gibt es einen ersten Gesetzentwurf. (There is a first draft law in the debate about new rules for organ donation.)

Reference: (12 words) In der Debatte über neue Regeln für Organspenden gibt es einen Gesetzentwurf. (There is a draft law in the debate about new rules for organ donation.)

Table 3: Two examples of various levels of compression, where the first shortens from 17 to 6 words, and the second only removes one word.

In Figure 2, we plot the distribution of the ratio between transcription lengths and target lengths over the test set. The first observation is that most of the transcriptions require shortening, as shown by the high frequencies of ratios over 1.

Moreover, the compression ratio varies across different utterances. Table 3 shows two examples, where the first compressed from 17 to 6 words, while the second only deletes one word. When inspecting the original videos, we notice that the first example contains many other visual contents, some also in text form, whereas the second example only involves the new anchorwoman speaking. The examples showcase that the level of desired compression depends on various factors, such as

Model	Ratio (output to desired length)	WER	R-1	R-2	R-L
Off-the-shelf ASR	1.09	41.5	75.9	58.9	72.6
+ compression	0.95	69.1	55.7	29.5	51.6
Baseline ASR	1.21	57.7	65.2	43.6	61.9
+ compression	1.00	74.3	48.5	23.3	44.7

Table 4: Word error rate and summarization quality on the test set by cascading two separate models for ASR and sentence compression.

Ground-truth	Es ist kurz nach Mitternacht, als plötzlich ein Auto in eine Gruppe von Menschen steuert, die ausgelassen ins neue Jahr feiern. It is just after midnight, when a car suddenly drives into a group of people who joyfully celebrate the new year.
Reference	Kurz nach Mitternacht steuert ein Auto in eine Gruppe von Menschen, die ins neue Jahr feiern. Just after midnight a car drives into a group of people who celebrate the new year.
Unsupervised compression	Kurz nach Mitternacht <u>fährt</u> ein Auto plötzlich in eine Gruppe von <u>Leuten</u> , die das <u>nächste</u> Jahr feiern. Just after midnight a car suddenly drives into a group of people who celebrate the <u>next</u> year.
Ground-truth	Unter dem Eindruck der Massenproteste hatten sich zuletzt auch hochrangige Militärs von ihm abgewandt. Under the impression of mass protests, senior military officials have finally also turned away from him.
Reference	Unter dem Eindruck der Proteste wandten sich zuletzt auch hochrangige Militärs ab. Under the impression of protests, senior military officials finally also turned away.
Unsupervised compression	Unter dem Eindruck <u>von Massenprotestieren</u> <u>waren</u> auch hochrangige Militärs von ihm <u>entfernt</u> . Under the impression of mass <u>protesting</u> , senior military officials <u>were</u> also <u>distanced</u> from him.

Table 5: Examples outputs of the unsupervised compression model, where undesired paraphrasing is underlined. English translations are in gray.

the amount of visual information simultaneously shown on the screen. Therefore, a globally fixed compression rate would not be suitable.

To comply with length constraints, the ASR outputs are shortened by the unsupervised sentence compression model. As the system is trained based on subwords, we use the number of BPE-tokens in the reference as target length. The first two rows in Table 4 contrast the output quality before and after compression, as measured in case-insensitive word error rate (WER), and ROUGE scores (Lin, 2004).⁴ To our surprise, compression has a large negative impact on the outputs in all four metrics, creating a gap of over 20% absolute. Via an exhaustive manual inspection over the test set, we find that the unsupervised compression model tends to paraphrase much more frequently than the references. While

⁴Ideally, these metrics should be accompanied by an independent human evaluation, which we could not perform due to resource constraints.

the paraphrased output is often valid both grammatically and semantically, the deviation from the references leads to higher WER and lower ROUGE scores. Two examples are given in Table 5, where several synonym replacements appear in the compression outputs, e.g. “fährt” for “steuert” (both “drives”), “Leute” for “Menschen” (both “people”), “nächste Jahr” for “neue Jahr” (“next year” for “new year”). In all these places, the references keep the original spoken words unchanged. Given the nature of our task, it is indeed undesirable to paraphrase excessively, as subtitles that are too different from the original spoken utterances could create a cognitive overload to users.

Considering these downsides, an ASR system is trained from scratch to provide more flexibility of structural modification. On our self-partitioned LibriVoxDeEn test set, we achieve a WER of 9.2%. The compression performance is reported in the lower section of Table 4. As this model is only

Model	Ratio (output to desired length)	WER	R-1	R-2	R-L
Baseline ASR	1.12	57.7	65.2	43.6	61.9
+adapt	1.05	38.5	76.8	58.7	74.4
+adapt+stop decoding	0.96	39.9	74.6	57.0	72.6

Table 6: Word error rate and summarization quality on the test set by the baseline ASR model. Fine-tuning on in-domain data is beneficial for both shortening outputs and increasing transcription quality.

Reference	In Brasilien ist Präsident Bolsonaro vereidigt worden. In Brazil new president Bolsonaro has been inaugurated.
Before adaptation	In Brasilien <u>ist</u> der neue Präsident <i>voll zu Narro</i> vereidigt worden. In Brazil the new president “ <i>voll zu Narro</i> ” has been inaugurated.
After adaptation	In Brasilien <u>wurde</u> der neue Präsident <i>Bolsonaro</i> vereidigt. In Brazil the new president <i>Bolsonaro</i> was inaugurated.

Table 7: Example outputs of the ASR model before and after adapting to the compression task. After adaptation, the model can perform shortening by changing tense, and correctly recognize the proper noun “Bolsonaro”. English translations are below each sentence in gray.

trained on LibriVoxDeEn audio books, its performance suffers from the train-test domain mismatch. This is exhibited by the gap of nearly 10% to the off-the-shelf system, which is trained on larger volume of data from various domains. Moreover, the transcription errors by the ASR system carry over as the input of the compression model, which is further disadvantageous to the final output quality. Lastly, similar to previous observations, post-processing by the compression model has a negative effect in terms of the evaluation metrics.

Overall, the performance of the unsupervised compression model suffers from paraphrasing. As an anonymous reviewer suggested, the aggressive paraphrasing can be remedied during decoding. For example, given a large beam size, we can select candidates that contain less paraphrasing. Otherwise, training with a paraphrasing penalty could also alleviate the problem. While we have not explored these methods here, they would indeed provide a more complete picture when comparing the cascaded and end-to-end approach.

5.2 Fine-Tuning ASR Model for Compression

Our baseline ASR model is trained on a different domain than the test set, and only for the transcription task. To improve performance on our task of interest, we apply fine-tuning on the adaptation corpus. The results are in Table 6. For easy visual comparison, the pre-adaptation performance in the first row is repeated from Table 4. Contrasting the performance before and after adaptation, we see no-

ticeable gains brought by adaptation in terms of all four quality metrics. Moreover, as evidenced by the reduced ratio of output to desired length, the model already performs shortening, despite not having received explicit length constraints. Table 7 shows an example where the adapted model changes verb tense to reduce output length. Specifically, by using the simple past (“wurde vereidigt / was inaugurated”) instead of the present perfect tense (“ist vereidigt worden / has been inaugurated”), the output becomes shorter. Meanwhile, we also observe the correct transcription of the proper noun “Bolsonaro”, which was mistakenly transcribed to the phonetically-similar “voll zu Narro” before adaptation. This illustrates that the adaptation step enables the model to improve recognition quality and compress its outputs simultaneously.

Despite the positive observations, the desired length constraints are not yet fully satisfied, as shown by the ratio of 1.05 between output to desired lengths. To examine the scenario of fully obeying the length constraints, we stop decoding once the number of allowed tokens runs out. The result is reported in the last row of Table 6. The ratio of 0.96 is lower than 1 because of a few instances where decoding stops before the countdown reaches zero. As the same output sequence can be constructed by different BPE-units, choosing longer subword units earlier on can lead to reaching the end-of-sequence token before depleting the number of allowed tokens. Lastly, from the quality metrics in the last row of Table 6, we see

Model	Ratio (output to desired length)	WER	R-1	R-2	R-L
(1): Baseline (satisfy length)	0.97	55.1	62.5	41.5	60.4
(2): + adapt	0.96	39.9	74.6	57.0	72.6
(3): Length embedding	1.00	57.8	61.9	40.5	59.8
(4): + adapt	0.96	39.3	74.3	55.2	72.5
(5): Length encoding	1.00	57.4	62.6	40.7	60.1
(6): + adapt	0.96	38.6	75.1	56.4	73.2

Table 8: Word error rate and summarization quality by the models with length count-down. Adaptation results in large gains. The model with length encoding outperforms the baseline and the one with length embedding.

Reference	[... 75 tokens] Sea-Watch spricht von einer politisch motivierten Blockade, um Rettungsaktionen zu verhindern. Sea-Watch speaks of a politically motivated blockade to prevent bailouts.
Len. Encoding	[... 83 tokens] Die Hilfsorganisation Sea-Watch spricht von einer politisch motivierten Blockade. The aid organization Sea-Watch speaks of a politically motivated blockade.
Len. Embedding	[... 82 tokens] Die Hilfsorganisation Sea-Watch spricht von einer politisch motivierten Blockade zu verhinder. The aid organization Sea-Watch speaks of a politically motivated blockade to prevent.

Table 9: Comparison of length encoding and length embedding models (after adaptation). When facing exceptionally long outputs, the length embedding model tends to stop abruptly, producing non-grammatical half-sentences.

that the forced termination of decoding comes with higher WER and lower ROUGE scores, indicating reduced output quality when fully satisfying the target length constraints.

5.3 Models with Explicit Length Constraints

While the baseline ASR model achieves some degree of compression after adaptation, it cannot fully comply with length constraints. Therefore, the following experiments examine the effects of training with explicit length count-downs. In Table 8, we report the performance of the ASR models with length embedding or encoding, as introduced in Section 3.2. For a complete comparison, in the first two rows of Table 8, we also include the baseline performance with forced termination of decoding.

Rows (3) and (5) show the performance of the two length count-down methods before adaptation. As the ratios of output to desired length are equal to 1, the models are always faithful to the given length constraints. This shows the effectiveness of injecting allowable length in training. However, we also observe that there is no quality gain over the baseline in row (1). To investigate the reason for this, we experimented by decoding one sample sequence with different allowable lengths. As

we gradually reduce the target length, the models first shorten their outputs by removing punctuation marks. Afterwards, instead of shortening the outputs by summarization, they stop decoding when running out of allowed number of tokens. Indeed, during training, the models are only incentivized to accurately transcribe the spoken utterance and to stop decoding when the count-down reaches zero. The same behavior therefore carries over to test time. Contrary to abstractive summarization (Takase and Okazaki, 2019) and machine translation (Lakew et al., 2019), in ASR, an input sequence has one single ground-truth transcription rather than multiple viable outputs. This could lead to a different level of abstraction than required in summarization or translation models. These observations with the unadapted model also highlight the importance of the subsequent fine-tuning step.

The results after adaptation are reported in rows (4) and (6). The large improvement in WER and ROUGE scores is in line with the previous finding when adapting the baseline. When decoding with length count-down, we define the target output length as the minimum of the baseline output length and the reference length. This ensures the same output-to-desired length ratio as the baseline

in row (1). Here, the first observation is that the length encoding model in row (5) outperforms the baseline in three of the four evaluation metrics. This suggests that it is beneficial to represent the constraints explicitly during training. Moreover, the length encoding model also consistently outperforms its embedding-based counterpart. This could be because the length encoding can extrapolate to any length values, and is equipped with a sense of relative differences between numerical values at initialization. On the other hand, the length embedding would need to learn the representation for different lengths during training. When inspecting the outputs for long utterances, we found that the embedding model is more likely to abruptly stop, such as the example shown in Table 9.

6 Related Work

6.1 Length-Controlled Text Generation

Controlling output length of natural language generation systems has been studied for several tasks.

For abstractive summarization, [Kikuchi et al. \(2016\)](#) proposed two methods to incorporate length constraints into LSTM-based encoder-decoder models. The first method uses length embedding at every decoding step, while the second adds the desired length in the first decoder state. For convolutional models, [Fan et al. \(2018\)](#) used special tokens to represent quantized length ranges, and provides the desired token to the decoder before output generation. [Liu et al. \(2018\)](#) adopted a more general approach, where the decoder directly ingests the desired length. More recently, [Takase and Okazaki \(2019\)](#) modified the positional encoding from the Transformer ([Vaswani et al., 2017](#)) to encode allowable lengths. [Makino et al. \(2019\)](#) proposed a loss function that encourages summaries within desired lengths. [Saito et al. \(2020\)](#) introduced a model that controls both output length and informativeness.

For machine translation, [Lakew et al. \(2019\)](#) used both the length range token and reverse length encoding. [Niehues \(2020\)](#) used the length embedding, encoding, as well as a combination of the original positional encoding and length count-down.

6.2 Sentence Compression

Our task of length-controlled ASR outputs is related to sentence compression, as the transcriptions can be compressed in post-processing. An early approach of supervised extractive sentence com-

pression was by [Filippova et al. \(2015\)](#), who proposed to predict the delete-or-keep choice for each output symbol. [Angerbauer et al. \(2019\)](#) extended this approach by integrating the desired compression ratio as part of the prediction label. [Yu et al. \(2018\)](#) proposed to combine the merits of extractive and abstractive approaches by first deleting on non-essential words and then generating new words. For unsupervised compression, [Fevry and Phang \(2018\)](#) trained a denoising auto-encoder to reconstruct original sentences, and in this way circumvented the need for supervised corpora.

7 Conclusion

In this work, we explored the task of compressing ASR outputs to enhance subtitle readability. This task has several unique properties. First, the compression is not solely deletion-based. Moreover, unnecessary paraphrasing must be limited to maintain a consistent user experience between hearing and reading.

We first investigated cascading an ASR module with a sentence compression model. Due to the absence of supervised corpora, the compression model is trained in an unsupervised fashion. Experiments showed that the outputs generated this way do not suit our task requirements because of unnecessary paraphrasing. We then adapted an end-to-end ASR model on a small corpus with compressed transcriptions. Via adaptation, the model learned to both shorten its outputs and improve transcription quality. Nevertheless, the given length constraints were not fully satisfied. Lastly, by explicitly injecting length constraints via reverse positional encoding, we achieved further performance gain, while completely adhering to length constraints.

A direction for future work is to incorporate more diverse measurements of output length as well as complexity. In this work, we measured length by the number of BPE-tokens. While this typically corresponds to the output length perceived visually, a more direct metric would be the number of characters. Moreover, output complexity, such as the proportion of long words, is also important for readability and therefore worth exploring. In a broader scope, as an anonymous reviewer suggested, a way to alleviate the resource scarcity for end-to-end ASR compression is to augment the training data with synthesized utterances from summarization corpora. We expect the augmentation to be complementary to our approaches in this work.

References

- Katrin Angerbauer, Heike Adel, and Ngoc Thang Vu. 2019. [Automatic compression of subtitles with neural networks and its effect on user experience](#). In [Proceedings of Interspeech'2019](#), pages 594–598.
- Benjamin Beilharz, Xin Sun, Sariya Karimova, and Stefan Riezler. 2020. [LibriVoxDeEn: A corpus for German-to-English speech translation and speech recognition](#). In [Proceedings of LREC'2020](#).
- Mauro Cettolo, Marcello Federico, Luisa Bentivogli, Jan Niehues, Sebastian Stüker, Katsutho Sudoh, Koichiro Yoshino, and Christian Federmann. 2017. [Overview of the IWSLT 2017 evaluation campaign](#). In [Proceedings of IWSLT'2017](#), pages 2–14.
- Sumit Chopra, Michael Auli, and Alexander M. Rush. 2016. [Abstractive sentence summarization with attentive recurrent neural networks](#). In [Proceedings of NAACL-HLT'2016](#), pages 93–98.
- James Clarke and Mirella Lapata. 2006. [Models for sentence compression: A comparison across domains, training requirements and evaluation measures](#). In [Proceedings of COLING/ACL'2006](#), pages 377–384.
- Trevor Cohn and Mirella Lapata. 2008. [Sentence compression beyond word deletion](#). In [Proceedings of COLING'2008](#), pages 137–144.
- Angela Fan, David Grangier, and Michael Auli. 2018. [Controllable abstractive summarization](#). In [Proceedings of WnMT'2018](#), pages 45–54, Melbourne, Australia. Association for Computational Linguistics.
- Thibault Fevry and Jason Phang. 2018. [Unsupervised sentence compression using denoising auto-encoders](#). In [Proceedings of CoNLL'2018](#), pages 413–422.
- Katja Filippova, Enrique Alfonseca, Carlos A. Colmenares, Lukasz Kaiser, and Oriol Vinyals. 2015. [Sentence compression by deletion with LSTMs](#). In [Proceedings of EMNLP'2015](#), pages 360–368.
- Thanh-Le Ha, Jan Niehues, and Alexander Waibel. 2017. [Toward multilingual neural machine translation with universal encoder and decoder](#). In [Proceedings of IWSLT'2016](#).
- Y. He, T. N. Sainath, R. Prabhavalkar, I. McGraw, R. Alvarez, D. Zhao, D. Rybach, A. Kannan, Y. Wu, R. Pang, Q. Liang, D. Bhatia, Y. Shangguan, B. Li, G. Pundak, K. C. Sim, T. Bagby, S. Chang, K. Rao, and A. Gruenstein. 2019. [Streaming end-to-end speech recognition for mobile devices](#). In [Proceedings of ICASSP'2019](#), pages 6381–6385.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. [Google's multilingual neural machine translation system: Enabling zero-shot translation](#). [Transactions of the Association for Computational Linguistics](#), 5:339–351.
- Yuta Kikuchi, Graham Neubig, Ryohei Sasano, Hiroya Takamura, and Manabu Okumura. 2016. [Controlling output length in neural encoder-decoders](#). In [Proceedings of EMNLP'2016](#), pages 1328–1338.
- Kevin Knight and Daniel Marcu. 2002. [Summarization beyond sentence extraction: A probabilistic approach to sentence compression](#). [Artificial Intelligence](#), 139(1):91–107.
- Surafel Melaku Lakew, Mattia Di Gangi, and Marcello Federico. 2019. [Controlling the output length of neural machine translation](#). In [Proceedings of IWSLT'2019](#).
- Bo Li, Tara N Sainath, Arun Narayanan, Joe Caroselli, Michiel Bacchiani, Ananya Misra, Izhak Shafran, Hasim Sak, Golan Pundak, Kean K Chin, et al. 2017. [Acoustic modeling for google home](#). In [Proceedings of Interspeech'2017](#), pages 399–403.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In [Text Summarization Branches Out](#), pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Yizhu Liu, Zhiyi Luo, and Kenny Zhu. 2018. [Controlling length in abstractive summarization using a convolutional neural network](#). In [Proceedings of EMNLP'2018](#), pages 4110–4119.
- Takuya Makino, Tomoya Iwakura, Hiroya Takamura, and Manabu Okumura. 2019. [Global optimization under length constraint for neural text summarization](#). In [Proceedings of ACL'2019](#), pages 1039–1048.
- Jan Niehues. 2020. [Machine translation with unsupervised length-constraints](#). [arXiv preprint arXiv:2004.03176](#).
- Ngoc-Quan Pham, Thai-Son Nguyen, Jan Niehues, Markus Müller, and Alex Waibel. 2019. [Very deep self-attention networks for end-to-end speech recognition](#). In [Proceedings of Interspeech'2019](#), pages 66–70.
- Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, Jan Silovsky, Georg Stemmer, and Karel Vesely. 2011. [The Kaldi speech recognition toolkit](#). In [Proceedings of ASRU'2011](#).
- Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. [A neural attention model for abstractive sentence summarization](#). In [Proceedings EMNLP'2015](#), pages 379–389.

- Itsumi Saito, Kyosuke Nishida, Kosuke Nishida, Atsushi Otsuka, Hisako Asano, Junji Tomita, Hiroyuki Shindo, and Yuji Matsumoto. 2020. [Length-controllable abstractive summarization by guiding with summary prototype](#). [arXiv preprint arXiv:2001.07331](#).
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In [Proceedings of ACL'2016](#), pages 1715–1725.
- Sho Takase and Naoaki Okazaki. 2019. [Positional encoding to control output sequence length](#). In [Proceedings NAACL-HLT'2019](#), pages 3999–4004.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, undefin-dukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In [Proceedings of NIPS'2017](#), pages 6000–6010.
- Naitong Yu, Jie Zhang, Minlie Huang, and Xiaoyan Zhu. 2018. [An operation network for abstractive sentence compression](#). In [Proceedings of COLING'2018](#), pages 1065–1076.